Pulse Connect Secure

Administration Guide

Product Release 8.1R5.1

Document Revision 1.0

Published: 2015-07-09

Pulse Connect Secure Administration Guide

The information in this document is current as of the date on the title page.

END USER LICENSE AGREEMENT

Abbreviated Table of Contents

Table of Contents

List of Figures

List of Tables

About This Guide

Objective

This guide describes basic configuration procedures for Pulse Connect Secure. This document was formerly titled *Pulse Connect Secure Administration Guide*. This document is now part of the Pulse Secure documentation set.

Audience

This guide is designed for network administrators who are configuring and maintaining a Pulse Connect Secure device. To use this guide, you need a broad understanding of networks in general and the Internet in particular, networking principles, and network configuration. Any detailed discussion of these concepts is beyond the scope of this guide.

Document Conventions

Table 1 defines notice icons used in this guide.

Table1: Notice Icons

Icon	Meaning	Description
	Informational note	Indicates important features or instructions.
	Caution	Indicates a situation that might result in loss of data or hardware damage.

	Warning	Alerts you to the risk of personal injury or death.
	Laser warning	Alerts you to the risk of personal injury from a laser.
	Tip	Indicates helpful information.
	Best practice	Alerts you to a recommended use or implementation.

Documentation

For a list of related Connect Secure documentation, see http://www.pulsesecure.net/support. If the information in the latest Release Notes differs from the information in the documentation, follow the Release Notes.

Obtaining Documentation

To obtain the most current version of all Pulse Secure technical documentation, see the products documentation page at http://www.juniper.net/techpubs.

Documentation Feedback

We encourage you to provide feedback, comments, and suggestions so that we can improve the documentation. You can provide feedback by using either of the following methods:

- Document name
- Page number
- Software release version

Requesting Technical Support

Technical product support is available through the Pulse Secure Global Support Center (PSGSC). If you have a support contract, then file a ticket with PSGSC.

- Product warranties—For product warranty information, visit http://www.pulsesecure.net.

Self-Help Online Tools and Resources

For quick and easy problem resolution, Pulse Secure, LLC has designed an online self-service portal called the Customer Support Center (CSC) that provides you with the following features:

- Find CSC offerings: http://www.pulsesecure.net/support
- Search for known bugs: http://www.pulsesecure.net/support

- Find product documentation: http://www.juniper.net/techpubs/

- Find solutions and answer questions using our Knowledge Base:
  http://www.pulsesecure.net/support

- Download the latest versions of software and review release notes:
  http://www.pulsesecure.net/support

- Search technical bulletins for relevant hardware and software notifications:
  http://www.pulsesecure.net/support

- Join and participate in the Pulse Secure, LLC Community Forum:
  http://www.pulsesecure.net/support

- Open a case online in the CSC Case Management tool: http://www.pulsesecure.net/support

To verify service entitlement by product serial number, use our Serial Number Entitlement (SNE) Tool: http://www.pulsesecure.net/support

Opening a Case with PSGSC

You can open a case with PSGSC on the Web or by telephone.

- Use the Case Management tool in the PSGSC at http://www.pulsesecure.net/support.

- Call 1-888-314-5822 (toll-free in the USA, Canada, and Mexico).

For international or direct-dial options in countries without toll-free numbers, seehttp://www.pulsesecure.net/support.

PART 1 Getting Started

- Initial Verification and Key Concepts
- Introduction to Connect Secure

CHAPTER 1 Initial Verification and Key Concepts

Verifying User Accessibility

You can easily create a user account in the system authentication server for use in verifying user accessibility to your device. After creating the account through the admin console, sign in as the user on the user sign-in page.

To verify user accessibility:

1. From the admin console, choose **Authentication > Auth. Servers**.

2. Select the **System Local** link.

3. Select the **Users** tab.

4. Click **New**.

5. Type **testuser1** as the username and enter a password, and then click **Save Changes**. The testuser1 account is now created.

6. Use another browser window to enter the machine's URL to access the user sign-in page. The URL is in the format: https://*a.b.c.d*, where *a.b.c.d* is the machine IP address you entered in the serial console when you initially configured your device.

7. Click **Yes** when prompted with the security alert to proceed without a signed certificate. The user sign-in page appears, indicating that you have successfully connected to your device.

8. Enter the username and password you created for the user account and then click Sign In to access the home page for users.

9. Enter the URL to an internal Web server in the Address box and click **Browse**. The system opens the Web page in the same browser window, so to return to the system home page, click the center button on the toolbar that appears on the target Web page.

10.   Enter the URL to your external corporate site on the system home page, and click **Browse**. A Web page opens in the same browser window, so use the button on the toolbar to return to the system home page.

11.   Click **Browsing > Windows Files** on the system home page to browse through available Windows file shares or **Browsing > UNIX/NFS Files** to browse through available UNIX NFS file shares.

Related Documentation

- Creating a Test Scenario to Learn Concepts and Best Practices

- Defining a User Role

- Defining a Resource Profile

- Defining an Authentication Server

- Defining an Authentication Realm

- Defining a Sign-In Policy

- Using the Test Scenario

Creating a Test Scenario to Learn Concepts and Best Practices

The system provides a flexible access management system that makes it easy to customize a user's remote access experience through the use of roles, resource policies, authentication servers, authentication realms, and sign-in policies. To enable you to quickly begin working with these entities, your device ships with system defaults for each. you can create each access management entity by performing the following tasks:

- Define a user role

- Define a resource policy

- Define an authentication server

- Define an authentication realm

- Define a sign-in policy

The system supports two types of users:

- Administrators—An administrator is a person who may view or modify system configuration settings. You create the first administrator account through the serial console.

- Users—A user is a person who uses the system to gain access to corporate resources as configured by an administrator.

Related Documentation

- Verifying User Accessibility

- Defining a User Role

- Defining a Resource Profile

- Defining an Authentication Server

- Defining an Authentication Realm

- Defining a Sign-In Policy

- Using the Test Scenario

Defining a User Role

Your device is preconfigured with one user role called "Users." This predefined role enables the Web and file browsing access features, enabling any user mapped to the Users role to access the Internet, corporate Web servers, and any available Windows and UNIX NFS file servers. You can view this role on the User Roles page.

After you enable an access feature for a role, configure the appropriate corresponding options that are accessible from the access feature's configuration tab.

To define a user role:

1. In the admin console, choose **Users > User Roles**.

2. Click **New Role**.

3. Enter **Test Role** in the Name box and then click **Save Changes**.

   Wait for the system to display the Test Role page with the General tab and Overview link selected.

4. Select the **Web** check box under Access features and then click **Save Changes**.

5. Select **Web > Options**.

6. Select the **User can type URLs in the IVE browser bar** check box, and then click **Save Changes**.

After completing these steps, you have defined a user role. When you create resource profiles, you can apply them to this role. You can also map users to this role through role mapping rules defined for an authentication realm.

To quickly create a user role that enables Web and file browsing, duplicate the Users role, and then enable additional access features as desired.

Related Documentation

- Verifying User Accessibility

- Creating a Test Scenario to Learn Concepts and Best Practices

- Defining a Resource Profile

- Defining an Authentication Server

- Defining an Authentication Realm

- Defining a Sign-In Policy

- Using the Test Scenario

Defining a Resource Profile

A resource profile is a set of configuration options that contains all of the resource policies, role assignments, and end-user bookmarks required to provide access to an individual resource.

Within a resource profile, a resource policy specifies the resources to which the policy applies (such as URLs, servers, and files) and whether the system grants access to a resource or performs an action. Note that the system is preconfigured with two types of resource policies:

- Web Access—The predefined Web Access resource policy, Initial Policy for Local Resources, allows access only to hosts belonging to domains within the secured network.

- Windows Access—The predefined Windows Access resource policy enables all users mapped to the Users role to access all corporate Windows file servers. By default, this resource policy applies to the Users role.

> NOTE: Delete the Windows Access resource policies if you are concerned about users having access to all of your Web and file content.

To define a resource profile:

1. In the admin console, choose **Users > Resource Profiles > Web**.

2. Click **New Profile**.

    The Web Applications Resource Profile page appears.

3. Fill in the following information:

    1. In the Type box, keep the default option (Custom).

    2. In the Name box, type **Test Web Access**.

    3. In the Base URL box, type **http://www.google.com**.

    4. Under Autopolicy: Web Access Control, select the check box next to the default policy (http://www.google.com:80/*) and choose **Delete**.

    5. In Resource box, type **http://www.google.com**, select **Deny** from the Action list, and click **Add**.

    6. Click **Save and Continue**.

        The Test Web Access page appears.

7. Click the **Roles** tab.

    1. Select **Test Role** in the Available Roles box and click **Add** to move it to the Selected Roles box.

    2. Click **Save Changes**.

The system adds Test Web Access to the Web Application Resource Policies page and automatically creates a corresponding bookmark that links to google.com.

After completing these steps, you have configured a Web Access resource profile. Even though the system comes with a resource policy that enables access to all Web resources, users mapped to Test Role are still prohibited from accessing http://www.google.com. These users

are denied access because the autopolicy you created during the resource profile configuration takes precedence over the default Web access policy that comes with the system.

Related Documentation

- Verifying User Accessibility

- Creating a Test Scenario to Learn Concepts and Best Practices

- Defining a User Role

- Defining an Authentication Server

- Defining an Authentication Realm

- Defining a Sign-In Policy

- Using the Test Scenario

Defining an Authentication Server

An authentication server is a database that stores user credentials—username and password—and typically group and attribute information. When a user signs in to the host, the user specifies an authentication realm, which is associated with an authentication server. The system forwards the user's credentials to this authentication server to verify the user's identity.

The system supports the most common authentication servers, including Windows NT Domain, Active Directory, RADIUS, LDAP, NIS, RSA ACE/Server, SAML Server, and eTrust SiteMinder, and enables you to create one or more local databases of users who are authenticated. The system is preconfigured with one local authentication server for users called "System Local." This predefined local authentication server is a system database that enables you to quickly create user accounts for user authentication. This ability provides flexibility for testing purposes and for providing third-party access by eliminating the need to create user accounts in an external authentication server.

You can view the default local authentication server on the Authentication Servers page.

> NOTE: The system also supports authorization servers. An authorization server (or directory server) is a database that stores user attribute and group information. You can configure an authentication realm to use a directory server to retrieve user attribute or group information for use in role mapping rules and resource policies.

To define an authentication server:

1. In the admin console, choose **Authentication > Auth. Servers**.

2. Select **Local Authentication** from the New list and then click **New Server**.

   The New Local Authentication page appears.

3. Enter **Test Server** in the Name box and then click **Save Changes**.

   Wait for the system to notify you that the changes are saved, after which additional configuration tabs appear.

4. Click the **Users** tab and then click **New**.

The New Local User page appears.

5. Enter **testuser2** in the Username box, enter a password, and then click **Save Changes** to create the user's account in the Test Server authentication server.

After completing these steps, you have created an authentication server that contains one user account. This user can sign in to an authentication realm that uses the Test Server authentication server.

The admin console provides last access statistics for each user account on the respective authentication servers pages, on the Users tab under a set of columns titled Last Sign-in Statistic. The statistics reported include the last successful sign-in date and time for each user, the user's IP address, and the agent or browser type and version.

Related Documentation

- Verifying User Accessibility

- Creating a Test Scenario to Learn Concepts and Best Practices

- Defining a User Role

- Defining a Resource Profile

- Defining an Authentication Realm

- Defining a Sign-In Policy

- Using the Test Scenario

Defining an Authentication Realm

An authentication realm is a grouping of authentication resources, including:

- An authentication server, which verifies a user's identity. The system forwards credentials submitted on a sign-in page to an authentication server.

- An authentication policy, which specifies realm security requirements that need to be met before the system submits credentials to an authentication server for verification.

- A directory server, which is an LDAP server that provides user and group attribute information to the system for use in role mapping rules and resource policies (optional).

- Role mapping rules, which are conditions a user must meet for the system to map a user to one or more roles. These conditions are based on information returned by the realm's directory server, the person's username, or certificate attributes.

Your system is preconfigured with one user realm called "Users." This predefined realm uses the System Local authentication server, an authentication policy that requires a minimum password length of four characters, no directory server, and contains one role mapping rule that maps all users who sign in to the Users realm to the Users role. The "testuser1" account you created is part of the Users realm, because this account is created in the System Local authentication server. The "testuser2" account you created is not part of the Users realm, because you create the user account in the new "Test Server" authentication server, which is not used by the Users realm.

You can view the default user authentication realm on the User Authentication Realms page.

To define an authentication realm:

1. In the admin console, choose **Users > User Realms**.

    The User Authentication Realms page appears.

2. Click **New**.

    The New Authentication Realm page appears.

3. Enter **Test Realm** in the Name box.

4. Select **Test Server** from the Authentication list.

5. Click **Save Changes**.

    Wait for the system to notify you that the changes are saved and to display the realm's configuration tabs.

6. Click the **Role Mapping** tab if it is not already selected, and then click **New Rule**.

    The Role Mapping Rule page appears.

7. Enter **testuser2** in the text box.

8. Under "...then assign these roles", select **Test Role** from the Available Roles list and click **Add** to move it to the Selected Roles box.

9. Click **Save Changes**.

After completing these steps, you have finished creating an authentication realm. This realm uses Test Server to authenticate users and a role mapping rule to map testuser2 to Test Role. Because the Test Web Access resource policy applies to Test Role, any user mapped to this role cannot access http://www.google.com.

Related Documentation

- Verifying User Accessibility

- Creating a Test Scenario to Learn Concepts and Best Practices

- Defining a User Role

- Defining a Resource Profile

- Defining an Authentication Server

- Defining a Sign-In Policy

- Using the Test Scenario

Defining a Sign-In Policy

A sign-in policy is a system rule that specifies:

- A URL where a user may sign in to the host.

- A sign-in page to display to the user.

- Whether or not the user needs to type or select an authentication realm to which the system submits credentials.

- The authentication realms where the sign-in policy applies.

All FIPS and non-FIPS devices are preconfigured with one sign-in policy that applies to users: */. This default user sign-in policy (*/) specifies that when a user enters the URL to the host, it displays the default sign-in page for the user and requires the user to select an authentication realm (if more than one realm exists). The */ sign-in policy is configured to apply to the Users authentication realm, therefore this sign-in policy does not apply to the authentication realm you created.

You can view the default user sign-in policy on the Signing In page. The */meeting sign-in policy is also listed on this page. This policy enables you to customize the sign-in page for Pulse Collaboration meetings.

To define a sign-in policy:

1. In the admin console, choose **Authentication > Signing in > Sign-in Policies**.

   The Signing In page appears.

2. Click ***/** under User URLs.

   The */ page appears.

3. Enter **test** after */ in the Sign-in URL box.

4. Under Authentication realm, select the **User picks from a list of authentication realms** option button, and then select **Test Realm** from the Available Realms list and click **Add** to move it to the Selected Realms box. (Repeat this process for the Users role if it is not already in the Selected Realms box.)

5. Click **Save Changes**.

After completing these steps, you have finished modifying the default users sign-in policy.

Optional Steps

You can perform these following optional steps to define a new sign-in page that is associated with the */test/ sign-in policy.

1. Select **Authentication > Signing In > Sign In Pages**, and then click **New Page.**

2. Enter **Test Sign-in Page** in the Name field, type **#FF0000** (red) in the Background color box, and then click **Save Changes**.

3. Select **Authentication > Signing In > Signing In Policies**, and then click **New URL**.

   The New Sign-In Policy page appears.

4. Type ***/test/** in the Sign-in URL box, select **Default Sign-in Page** from the Sign-in Page list, and click **Save Changes**.

5. Select **Authentication > Signing In > Sign In Policies**, and then click ***/test/** under User URLs.

   The */test/ page appears.

6. Select **Test Sign-in Page** from the Sign-in page list and then click **Save Changes**.

Related Documentation

- Verifying User Accessibility

- Creating a Test Scenario to Learn Concepts and Best Practices

- Defining a User Role

- Defining a Resource Profile

- Defining an Authentication Server

- Defining an Authentication Realm

- Using the Test Scenario

Using the Test Scenario

The test scenario enables you to do the following tasks:

- Access the user console using the modified default sign-in policy.

- Sign in as the user created in the Test Server to map to the Test Realm.

- Test your Web browsing capabilities, which are dependent upon the proper configuration of Test Role and Test Web Access.

To use the test scenario:

1. In a browser, enter the machine's URL followed by **/test** to access the user sign-in page. The URL is in the format: https://*a.b.c.d*/test, where *a.b.c.d* is the machine IP address you entered in the serial console during initial configuration.

2. Click **Yes** when prompted with the security alert to proceed without a signed certificate. If the user sign-in page appears, you have successfully connected to your device.

   > ⓘ NOTE: If you performed the optional configuration steps in "Defining a Sign-In Policy", the header color is red.

3. Enter the username and password you created for the user account in Test Server, type Test Realm in the Realm box, and then click Sign In to access the system home page for users.

   The system forwards the credentials to Test Realm, which is configured to use Test Server. Upon successful verification by this authentication server, the system processes the role mapping rule defined for Test Realm, which maps testuser2 to Test Role. Test Role enables Web browsing for users.

4. In the browser Address box, enter the URL to your corporate web site and click Browse. The Web page opens in the same browser window, so to return to the system home page, click the center icon in the browsing toolbar that appears on the target Web page.

5. On the system home page, type **www.google.com** and click **Browse**. An error message appears because the Test Web Access resource policy denies access to this site for users mapped to Test Role.

6. Return to the system home page, click **Sign Out**, and then return to the user sign-in page.

7. Enter the credentials for testuser1, specify the Users realm, and then click **Sign In**.

8. On the system home page, type **www.google.com** and click **Browse**. The Web page opens in the same browser window.

9. The test scenario demonstrates the basic access management mechanisms. You can create very sophisticated role mapping rules and resource policies that control user access depending on factors such as a realm's authentication policy, a user's group membership, and other variables. To learn more about access management, we recommend that you take a few minutes to review the online Help to familiarize yourself with its contents.

   When you configure your device for your enterprise, we recommend that you perform user access configuration. Before you make your device available from external locations, we recommend that you import a signed digital certificate from a trusted certificate authority (CA).

Related Documentation

- Verifying User Accessibility

- Creating a Test Scenario to Learn Concepts and Best Practices

- Defining a User Role

- Defining a Resource Profile

- Defining an Authentication Server

- Defining an Authentication Realm

- Defining a Sign-In Policy

Default Settings for Administrators

Just like for users, the system provides default settings that enable you to quickly configure accounts for administrators. This list summarizes the system default settings for administrators:

- Administrator roles—There are two built-in administrator roles.

  - .Administrators — This built-in role permits administrators to manage all aspects of the device. The administrator user you create through the serial console is mapped to this role.

  - .Read-Only Administrators — This built-in role permits users mapped to the role to view (but not configure) all settings. You need to map administrators to this role if you want to restrict their access.

- Administrators local authentication server — The Administrators local authentication server is a database that stores administrator accounts. You create the first administrator account in this server through the serial console. (All administrator accounts created through the serial console are added to this server.) You cannot delete this local server.

- Admin Users authentication realm — The Admin Users authentication realm uses the default Administrators local authentication server, an authentication policy that requires a minimum password length of four characters, no directory server, and one role mapping rule that maps all users who sign in to the Admin Users realm to the .Administrators role. The administrator account you create through the serial console is part of the Admin Users realm.

- */admin sign-in policy — The default administrator sign-in policy (*/admin) specifies that when a user enters the URL to the host followed by /admin, the system displays the default sign-in page for administrators. This policy also requires the administrator to select an authentication realm (if more than one realm exists). The */admin sign-in policy is configured to apply to the Admin Users authentication realm, therefore this sign-in policy applies to the administrator account you create through the serial console.

Related Documentation

- Defining a User Role

CHAPTER 2 Introduction to Connect Secure

- Connect Secure Overview
- Securing Traffic with Connect Secure
- Authenticating Users with Existing Servers
- Fine-Tuning Access to Connect Secure and the Resources It Intermediates
- Creating a Seamless Integration Between Connect Secure and the Resources It Intermediates
- Protecting Against Infected Computers and Other Security Concerns
- Ensuring Redundancy in the Connect Secure Environment
- Making the Interface Match a Company's Look-and-Feel
- Enabling Users on a Variety of Computers and Devices to Use Connect Secure
- Providing Secure Access for International Users
- Configuring Connect Secure
- Network and Security Manager and Connect Secure
- Configuring Connect Secure for the Initial DMI Connection
- Managing Large Binary Data Files
- Uploading and Linking Large Binary Data Files with NSM
- Importing Custom Sign-In Pages with NSM
- Importing Antivirus LiveUpdate Settings with NSM
- Importing Endpoint Security Assessment Plug-in (ESAP) Packages with NSM
- Uploading a Third-Party Host Checker Policy with NSM
- Linking to a Third-Party Host Checker Policy Shared Object with NSM
- Importing Hosted Java Applets with NSM
- Importing a Custom Citrix Client .cab File with NSM
- Introducing Pulse Secure Client
- Pulse Secure Client Configuration Overview
- Configuring a Role for Pulse Connect Secure
- Pulse Connection Set Options for Pulse Connect Secure
- Creating a Client Connection Set
- Configuring Connection Rules for Location Awareness
- Creating a Client Component Set for Pulse Connect Secure
- Pulse Secure Client Installation Overview

Connect Secure Overview

The Pulse Connect Secure enables you to give employees, partners, and customers secure and controlled access to your corporate data and applications including file servers, Web servers, native messaging and e-mail clients, hosted servers, and more from outside your trusted network using just a Web browser.

Connect Secure provides robust security by intermediating the data that flows between external users and your company's internal resources. Users gain authenticated access to authorized resources through an extranet session hosted by the appliance. During intermediation, the system receives secure requests from the external, authenticated users and then makes requests to the internal resources on behalf of those users. By intermediating content in this way, it eliminates the need to deploy extranet toolkits in a traditional DMZ or provision a remote access VPN for employees.

To access the intuitive home page, your employees, partners, and customers need only a Web browser that supports SSL and an Internet connection. This page provides the window from which your users can securely browse Web or file servers, use HTML-enabled enterprise applications, start the client/server application proxy, begin a Windows, Citrix, or Telnet/SSH terminal session, access corporate e-mail servers, start a secured Layer 3 tunnel, or schedule or attend a secure online meeting.

NOTE: These capabilities depend upon the product and upgrade options you have purchased.

Figure 1 shows an example of Connect Security within a LAN environment.

Figure1: Connect Secure Working Within a LAN
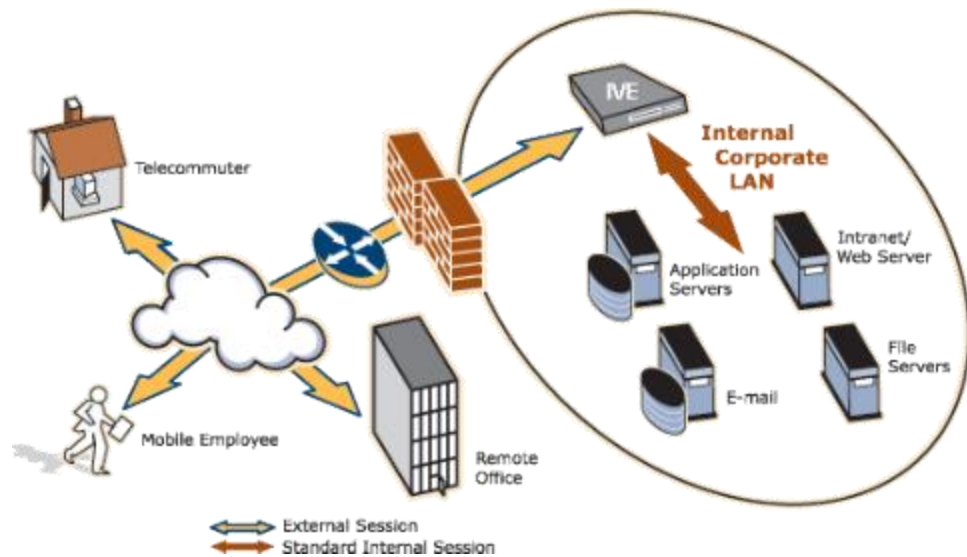
You can configure the system in the following ways:

- Provide users with secure access to a variety of resources. Connect Secure intermediates access to multiple types of applications and resources such as Web-based enterprise applications, Java applications, file shares, terminal hosts, and other client/server applications such as Microsoft Outlook, Lotus Notes, the Citrix ICA Client, and pcAnywhere. Additionally, administrators can provision an access method that allows full Layer 3 connectivity, providing the same level of access that a user would get if they were on the corporate LAN.

- Fine-tune user access to the appliance, resource types, or individual resources based on factors such as group membership, source IP address, certificate attributes, and endpoint security status. For instance, you can use dual-factor authentication and client-side digital certificates to authenticate users and use LDAP group membership to authorize users to access individual applications.

- Assess the security status of your users' computers by checking for endpoint defense tools such as current antivirus software, firewalls, and security patches. You can then allow or deny users access to the appliance, resource types, or individual resources based on the computer's security status.

Connect Secure acts as a secure, Application Layer gateway intermediating all requests between the public Internet and internal corporate resources. All requests that enter the system are already encrypted by the end user's browser, using SSL/HTTPS 128-bit or 168-bit encryption—unencrypted requests are dropped. Because the system provides a robust security layer between the public Internet and internal resources, administrators do not need to constantly manage security policies and patch security vulnerabilities for numerous different application and Web servers deployed in the public-facing DMZ.

Related Documentation

- Securing Traffic with Connect Secure
- Authenticating Users with Existing Servers

Securing Traffic with Connect Secure

Connect Secure enables you to secure access to a wide variety of applications, servers, and other resources through its remote access mechanisms. Once you have chosen which resource you want to secure, you can then choose the appropriate access mechanism.

For instance, if you want to secure access to Microsoft Outlook, you can use the Secure Application Manager (SAM). The Secure Application Manager intermediates traffic to client/server applications including Microsoft Outlook, Lotus Notes, and Citrix. Or, if you want to secure access to your company Intranet, you can use the Web rewriting feature. This feature uses the system's Content Intermediation Engine to intermediate traffic to Web-based applications and Web pages.

The system includes remote access mechanisms that intermediate the following types of traffic:

- Web-based traffic, including Web pages and Web-based applications—Use the Web rewriting feature to intermediate this type of content. The Web rewriting feature includes templates that enable you to easily configure access to applications such as Citrix, OWA, Lotus iNotes, and Sharepoint. In addition, you can use the Web rewriting custom configuration option to intermediate traffic from a wide variety of additional Web-based applications and Web pages, including custom-built Web applications.

- Java applets, including Web applications that use Java applets—Use the hosted Java applets feature to intermediate this type of content. This feature enables you to host Java applets and the HTML pages that they reference directly on Connect Secure rather than maintaining a separate Java server.

- File traffic, including file servers and directories—Use the file rewriting feature to intermediate and dynamically "webify" access to file shares. The file rewriting feature enables you to secure traffic to a variety of Windows and UNIX based servers, directories, and file shares.

- Client/server applications—Use the Secure Application Manager (SAM) feature to intermediate this type of content. SAM comes in two varieties (Windows and Java versions, or WSAM and JSAM). The WSAM and JSAM features include templates that enable you to easily configure access to applications such as Lotus Notes, Microsoft Outlook, NetBIOS file browsing, and Citrix. In addition, you can use the WSAM and JSAM custom configuration options to intermediate traffic from a wide variety of additional client/server applications and destination networks.

- Telnet and SSH terminal emulation sessions—Use the Telnet/SSH feature to intermediate this type of content. This feature enables you to easily configure access to a variety of networked devices that utilize terminal sessions including UNIX servers, networking devices, and other legacy applications.

- Windows Terminal Servers and Citrix server terminal emulation sessions— Use the Terminal Services feature to intermediate this type of content. This feature enables you to easily configure access to Windows Terminal Servers, Citrix MetaFrame Servers, and Citrix Presentation Servers (formerly known as Nfuse servers). You can also use this feature to deliver the terminal services clients directly from the system, eliminating the need to use another Web server to host the clients.

- E-mail clients based on the IMAP4, POP3, and SMTP protocols—Use the e-mail client feature to intermediate this type of content. This feature enables you to easily configure access to any corporate mail server based on the IMAP4, POP3, and SMTP protocols, such as Microsoft Exchange Server and Lotus Notes Mail servers. In addition, for iOS devices that have the Pulse Secure client, the Secure Mail option can synchronize the device with an Exchange server (ActiveSync) and encrypt the e-mail body and attachments.

- All network traffic—Use the VPN Tunneling feature to create a secure, Layer 3 tunnel over the SSL connection, allowing access to any type of application available on the corporate network. This feature enables you to easily connect remote users into your network by tunneling network traffic over port 443, enabling users full access to all of your network resources without configuring access to individual servers, applications, and resources.

Related Documentation

- Connect Secure Overview

Authenticating Users with Existing Servers

You can easily configure Connect Secure to use your company's existing servers to authenticate your end users—Users do not need to learn a new username and password to access the device. The system supports integration with LDAP, RADIUS, NIS, Windows NT Domain, Active Directory, eTrust SiteMinder, SAML, and RSA ACE/Servers.

Or, if you do not want to use one of these standard servers, you can store usernames and credentials directly on the system and use it as an authentication server. In addition, you can choose to authenticate users based on attributes contained in authentication assertions generated by SAML authorities or client-side certificates. Or, if you do not want to require your users to sign into the device, you can use the anonymous authentication server, which allows users to access the device without providing a username or password.

> NOTE: Pulse Secure Mobile client supports only one case of dual-factor authentication, where the client certificate is the primary and local auth is the secondary.

Related Documentation

- Connect Secure Overview

- AAA Server Overview

Fine-Tuning Access to Connect Secure and the Resources It Intermediates

In addition to using authentication servers to control access to Connect Secure, you can control access to the system and the resources it intermediates using a variety of additional client-side checks. Connect Secure enables you to create a multilayered approach to protect it and your resources:

1. First, you can perform preauthentication checks that control user access to the system's sign-in page. For instance, you might configure the system to check whether or not the user's computer is running a particular version of Norton Antivirus. If it is not running, you can determine that the user's computer is unsecure and disable access to the the system's sign-in page until the user has updated the computer's antivirus software.

2. Once a user has successfully accessed the system's sign-in page, you can perform realm-level checks to determine whether he can access the system's end-user home page. The most common realm-level check is performed by an authentication server. (The server determines whether the user enters a valid username and password.) You can perform other types of realm-level checks, however, such as checking that the user's IP address is in your network or that the user is using the Web browser type that you specify.

   If a user passes the realm-level checks that you specify, the user can access the system's end-user home page. Otherwise, it does not enable the user to sign in, or displays a "stripped down" version of the home page that you create. Generally, this stripped down version contains significantly less functionality than is available to your standard users because the user has not passed all of your authentication criteria. The system provides extremely flexible policy definitions, enabling you to dynamically alter end-user resource access based on corporate security policies.

3. After the system successfully assigns a user to a realm, the appliance maps the user to a role based on your selection criteria. A role specifies which access mechanisms a selected group of users can access. It also controls session and UI options for that group of users. You can use a wide variety of criteria to map users to roles. For instance, you can map users to different roles based on endpoint security checks or on attributes obtained from an LDAP server or client-side certificate.

4. In most cases, a user's role assignments control which individual resources the user can access. For instance, you might configure access to your company's Intranet page using a Web resource profile and then specify that all members of the Employees role can access that resource.

However, you can choose to further fine-tune access to individual resources. For instance, you may enable members of the Employees role to access your company's Intranet (as described earlier), but add a resource policy detailed rule that requires users to meet additional criteria to access the resource. For example, you may require users to be members of the Employees role and to sign into the device during business hours to access your company Intranet.

Related Documentation

- Connect Secure Overview

- Access Management Overview

Creating a Seamless Integration Between Connect Secure and the Resources It Intermediates

In a typical configuration, you could add bookmarks directly to the system's end-user home page. These bookmarks are links to the resources that you configure the system to intermediate. Adding these bookmarks enables users to sign into a single place (Connect Secure) and find a consolidated list of all of the resources available to them.

Within this typical configuration, you can streamline the integration between the system and the intermediated resources by enabling single sign-on (SSO). SSO is a process that allows preauthenticated users to access other applications or resources that are protected by another access management system without having to re-enter their credentials. During system configuration, you can enable SSO by specifying user credentials that you want the system to pass to the intermediated resources.

Or, if you do not want to centralize user resources on the system's end-user home page, you could create links to the -intermediated resources from another Web page. For instance, you can configure bookmarks on Connect Secure, and then add links to those bookmarks from your company's Intranet. Your users can then sign into your company Intranet and click the links there to access the intermediated resources without going through the system's home page. As with standard Connect Secure bookmarks, you can enable SSO for these external links.

Related Documentation

- Connect Secure Overview
- About Single Sign-On

Protecting Against Infected Computers and Other Security Concerns

Connect Secure enables you to protect against viruses, attacks, and other security concerns using the Host Checker feature. Host Checker performs security checks on the clients that connect to the device. For instance, you can use Host Checker to verify that end-user systems contain up-to-date antivirus software, firewalls, critical software hotfixes, and other applications that protect your users' computers. You can then enable or deny users access to the system's sign-in pages, realms, roles, and resources based on the results that Host Checker returns. Or, you can display remediation instructions to users so they can bring their computers into compliance

You can also secure your network from hostile outside intrusion by integrating your device with a Juniper Networks Intrusion Detection and Prevention (IDP) sensor. You can use IDP devices to detect and block most network worms based on software vulnerabilities, non-file-based Trojan horses, the effects of Spyware, Adware, and Key Loggers, many types of malware, and zero day attacks through the use of anomaly detection.

Related Documentation

- Connect Secure Overview
- Configuring Connect Secure to Interoperate with IDP

Ensuring Redundancy in the Connect Secure Environment

You can ensure redundancy in your environment using the clustering feature. With this feature, you can deploy two or more appliances as a cluster, ensuring no user downtime in the rare event of failure and stateful peering that synchronizes user settings, system settings, and user session data.

These appliances support active/passive or active/active configurations across a LAN. In Active/Passive mode, one device actively serves user requests while the other device runs passively in the background to synchronize state data. If the active device goes offline, the passive device automatically starts servicing user requests. In active/active mode, all the machines in the cluster actively handle user requests sent by an external load balancer. The load balancer hosts the cluster VIP and routes user requests to a device defined in its cluster group based on source-IP routing. If a device goes offline, the load balancer adjusts the load on the other active device.

> NOTE: WAN clustering is not supported on the MAG Series Pulse Gateways, except as it relates to campus networks. In a well-connected campus network, where the connectivity is more LAN-like than WAN-like, the Pulse Secure Gateways can be clustered in separate buildings.

Related Documentation

- Connect Secure Overview

Making the Interface Match a Company's Look-and-Feel

Connect Secure enables you to customize a variety of elements in the end-user interface. Using these customization features, you can update the look-and-feel of the system's end-user console so it will resemble one of your standard company Web pages or applications.

For instance, you can easily customize the headers, background colors, and logos that display in the sign-in page and end-user console to match your company's style. You can also easily customize the order in which bookmarks display and the end user help system.

Or, if you do not want to display the end-user home page to users (either in standard or customized form), you can choose to redirect users to a different page (such as your company Intranet) when users first sign into the system console. If you choose to use this option, you may want to add links to your system's bookmarks on the new page.

If you want to further customize the sign-in page, you can use custom sign-in pages. Unlike the standard customization options that you can configure through the system's admin console, the custom sign-in pages feature does not limit the number of customizations you can make to your pages. Using this feature, you can use an HTML editor to develop a sign-in page that exactly matches your specifications.

Related Documentation

- Connect Secure Overview
- Creating a Seamless Integration Between Connect Secure and the Resources It Intermediates

- Customizable Admin and End-User UIs

Enabling Users on a Variety of Computers and Devices to Use Connect Secure

In addition to allowing users to access Connect Secure from standard workstations and kiosks running Windows, Macintosh, and Linux operating systems, end users can access the system from connected PDAs, handhelds and smart phones such as i-mode and Pocket PC. When a user connects from a PDA or handheld device, the system determines which pages and functionality to display based on settings that you configure.

For more information about specifying which pages display to different devices, see the Pulse Connect Secure supported platforms document available on the Pulse Secure, LLC Customer Support Center web site.

Related Documentation

- Connect Secure Overview

- Handheld Devices and PDAs

Providing Secure Access for International Users

Connect Secure supports English (US), French, German, Spanish, Simplified Chinese, Traditional Chinese, Japanese, and Korean. When your users sign into the device, it automatically detects the correct language to display based on the user's Web browser setting. Or, you can use end-user localization and custom sign-in pages options to manually specify the language that you want to display to your end users.

Related Documentation

- Connect Secure Overview

- About Multi-Language Support for Connect Secure

Configuring Connect Secure

To enable users to start using Connect Secure, you must complete the following basic steps:

1. Plug in the device, connect it to your network, and configure its initial system and network settings (see the hardware guide for your MAG Series or SA Series device).

2. After you connect the device to your network, you need to set the system date and time, upgrade to the latest service package, and install your product licenses. When you first sign into the admin console, an initial configuration task guide is displayed to walk you through this process.

3. After you install your product licenses, use the following procedure to set up your access management framework to enable users to authenticate and access resources. Create a test scenario to familiarize yourself with the process (see "Creating a Test Scenario to Learn Concepts and Best Practices").

   1. Define an authentication server that verifies user names and passwords.

   2. Create user roles that enable access mechanisms, session options, and UI options for user groups.

3. Create a user authentication realm that specifies the conditions that users must meet to sign into the device.

4. Define a sign-in policy that specifies the URL that users must access to sign into the device and the page that they see when they sign in.

5. Create resource profiles that control access to resources, specify which user roles can access them, and include bookmarks that link to the resources.

The admin console includes a task guide to assist you with this process. To view the task guide, click **Guidance** in the upper right corner of the admin console. .

Once you have completed these basic steps, your system is ready for use. You can start using it as is, or configure additional advanced features such as endpoint defense and clustering.

Related Documentation

- Creating a Test Scenario to Learn Concepts and Best Practices

Network and Security Manager and Connect Secure

Network and Security Manager (NSM) is Juniper Networks network management tool that allows distributed administration of network appliances. You can use the NSM application to centralize status monitoring, logging, and reporting, and to administer system configurations.

With NSM you can manage most of the parameters that you can configure through the system's admin console. The configuration screens rendered through NSM are similar to the system's native interface.

NSM incorporates a broad configuration management framework that allows co-management using other methods. You can import and export XML through the system admin console interface.

How Connect Secure and NSM Communicate

The system and the NSM application communicate through the Device Management Interface (DMI). DMI is a collection of schema-driven protocols that run on a common transport (TCP). DMI is designed to work with Juniper Networks platforms to make device management consistent across all administrative realms. The DMI protocols that are supported include NetConf (for inventory management, XML-based configuration, text-based configuration, alarm monitoring, and device specific commands), structured syslog, and threat flow for network profiling. DMI supports third-party network management systems that incorporate the DMI standard, however only one DMI-based agent per device is supported.

The configuration is represented as a hierarchical tree of configuration items. This structure is expressed in XML that can be manipulated with NetConf. NetConf is a network management protocol that uses XML. DMI uses NetConf's generic configuration management capability and applies it to allow remote configuration of the device.

To allow NSM to manage the system using the DMI protocol, NSM must import the schema and meta-data files from the Pulse Secure Schema Repository, a publicly-accessible resource that is updated with each device release. In addition to downloading the system's current schema, NSM may also download upgraded software.

The Schema Repository enables access to XSD and XML files defined for each device, model and software version.

Before attempting to communicate with NSM, you must first complete the initial configuration of the device. Initial configuration includes network interface settings, DNS settings, licensing, and password administration.

If you have several devices that will be configured in a clustering environment, the cluster abstraction must first be created in the NSM Cluster Manager. Then, you can add individual nodes. NSM cannot auto-detect cluster membership.

After you have completed the initial network configuration, you can configure the device to communicate with NSM using the appropriate network information. Once the device has been configured to communicate with NSM, it contacts NSM and establishes a DMI session through an initial TCP handshake.

All communications between the device and NSM occur over SSH to ensure data integrity.

After the device initially contacts NSM and a TCP session is established, interaction between it and NSM is driven from NSM, which issues commands to get hardware, software, and license details. NSM connects to the Schema Repository to download the configuration schema that is particular to the device.

NSM then issues a command to retrieve configuration information from the device. If NSM is contacted by more than one device as a member of a cluster, information from only one of the cluster devices is gathered. NSM attempts to validate the configuration received from the device against the schema from Juniper Networks.

Once the device and NSM are communicating, it delivers syslog and event information to NSM.

After NSM and the device are connected, you can make any configuration changes directly on the device, bypassing NSM. NSM automatically detects these changes and imports the new configuration data. Changes to cluster members will similarly be detected by NSM.

When you make changes to the configuration through NSM you must push the changes to the device by performing an Update Device operation.

When you double-click the Connect Secure icon in the Device Manager and select the Config tab, the configuration tree appears in the main display area in the same orientation as items appear on the system's interface.

Available Services and Configuration Options

The following services and options are provided to NSM by Connect Secure:

- **Inventory management service**— enables management of software, hardware, and licensing details. Adding or deleting licenses is not supported; however upgrading/downgrading software is supported.

- **Status monitoring service**— allows system status to be obtained, including name, domain, OS version, synchronization status, connection details, and current alarms.

- **Logging service**—allows system logs to be obtained in a time-generated order. Logging configuration details that are set on the device apply to NSM.

- **XML-based configuration management service**— enables NSM to manage the configuration of the device. NSM uses the same XML Schema as the device, so you can troubleshoot NSM using XML files downloaded from the device.

The following device configuration items are not supported:

- Editing licensing information, (though licenses can be viewed)

- Creating clusters, joining nodes to clusters, or enabling or disabling cluster nodes

- Packaging log files or debug files for remote analysis

- Rebooting the system

DMI Communication with Connect Secure

To configure the device to communicate with NSM you must coordinate actions between the Connect Secure and NSM administrators. Items such as IP address, password, HMAC key (a one-time password), and the device ID must be shared between administrators of both the device and NSM.

To connect the device and NSM you need to do the following:

- Install and configure the device.

- Add Connect Secure as a device in NSM.

- Configure and activate the DMI agent on Connect Secure.

- Confirm connectivity and import the device configuration into NSM.

Related Documentation

- Configuring Connect Secure for the Initial DMI Connection

- Managing Large Binary Data Files

- Uploading and Linking Large Binary Data Files with NSM

Configuring Connect Secure for the Initial DMI Connection

To permit Connect Secure and NSM to make an initial connection, you must add an NSM administrative user to the Connect Secure configuration. This topic provides a summary of adding the NSM administrator and configuring the DMI agent to allow the device and NSM to communicate. Complete configuration of the device for authenticating users is outside the scope of this topic.

To initiate a DMI session for communication between Connect Secure and NSM:

1. Ensure that basic connection information is configured on the device (network address, DNS, password).

2. Ensure that the proper licenses are installed on the device.

3. From the NSM UI client Device Manager, click the Add icon and select Device to open the Add Device wizard, and enter the applicable information required to add a Connect Secure device to NSM. See the Administration Guide.

> **NOTE:** You must enter a unique NSM admin username and password on the NSM UI client. This username will be used on the device as the username for the administrator account that will be used to manage the Connect Secure device. NSM must have a unique account login to avoid interrupting the communication with device. NSM automatically generates a unique ID which is used for the HMAC key.

4. From the Connect Secure admin console, select **Authentication > Auth. servers** and enter the username and password of the NSM admin using the credentials you entered on NSM in the applicable authentication server. Use the NSM username and password that you entered in the NSM UI Client.

> **NOTE:** Only password-based authentication servers can be used. One-time password authentication is not supported.

5. Select **Administrators > Admin Roles** and create a DMI agent administrator role.

6. Select **Administrators > Admin Realms** and create a new DMI agent admin realm for the DMI agent on the device and use role mapping to associate the DMI agent role and realm.

7. On the NSM interface, select the Domain menu and choose the domain to which the device will be added.

8. In Device Manager, click the Add icon and select Device to open the Add Device wizard, and enter the applicable information required to add a device to NSM. See the Administration Guide.

> **NOTE:**
> - In a clustering environment, each cluster-node must have its own unique DMI agent and its own device-id and HMAC key, as each cluster node maintains its own persistent DMI connection to the management application.
>
> - The HMAC key and the device id are hashed to identify individual devices to the application. Pulse Secure recommends that you use a strong password for the HMAC key value to ensure that the key isn't guessed.

9. After you have added the device to NSM, select **System > Configuration > DMI Agent** on the Connect Secure admin console.

10. Under DMI connection, select the:

- **Inbound Enabled** check box if you are using an SSH secure shell Command Line Interface (CLI) to manage the device. The device can also be managed by integrating an SSH-aware netconf that complies with the DMI specification.

- **Outbound Enabled** check box if you are configuring the device to communicate with NSM.

> **NOTE:** When you enable or disable a connection, it takes a few minutes for the connection state to be updated.

11.   Under DMI settings for inbound connections, enter the TCP port in which the device should accept connections. This TCP port should not be used by any other system processes. We recommend you use the default port or a port number larger than 1024.

12.   Under DMI settings for outbound connection, enter the NSM Primary Server IP address or hostname, Primary Port, Backup Server and Backup Port (if applicable), the Device ID, and the HMAC Key.

13.   Select the Admin realm that you have configured for the DMI agent.

14.   If you do not want DMI logging for both inbound and outbound DMI connections, uncheck the **DMI Logging** check box. By default, DMI logging is enabled.

The device initiates a TCP connection to NSM. After the device is identified to NSM through the HMAC key and device ID hash, the device and NSM negotiate an SSH tunnel, and NSM requests authentication to the device based on the username and password.

If you need to disconnect the device from NSM, you can either disable the DMI agent from the device, or you can delete the device from the NSM interface. If the DMI connection is later reestablished, NSM will automatically retrieve any configuration changes, as well as logs that have accumulated in the interim.

To add a cluster in NSM, you first add the cluster, then you add each member. Adding a member is similar to adding a standalone device. You must have a cluster object and all of the cluster members defined in NSM to allow NSM to access the cluster.

Managing Large Binary Data Files

Large binary data files that form a part of the configuration of the device are handled differently from the remainder of the configuration in NSM. The size of some of these binary files could cause configurations to be so large as to overload resources on the NSM server. Consequently, only the large binary files you specify get imported into NSM, and those files are configured as shared objects, which avoids duplication if applied to multiple devices.

With NSM, large binary data files are not imported with the rest of the configuration during a normal device import operation. Instead, the file is represented in the device configuration tree by a stub containing an MD5 hash and file length designation. If you need to manage such a file in NSM, you upload the file separately, and configure it as a shared object. To include the file as part of the device object in NSM, you must then establish a link between the node in the device configuration tree and the shared binary data object. When you establish the link, a pointer to the shared binary data object replaces the MD5 hash and length.

After you establish the link, an Update Device directive pushes all linked binary data files to the device along with the rest of the device configuration. No binary data is pushed for nodes that still contain the MD5 hash and length designators.

If you do not need to manage a large binary data file from NSM, then you do not need to include it in the device object configuration. For example, suppose you have a hosted Java applet that resides on a device, and you have no intention of updating this applet. In this case, no shared object creation or file upload is necessary. NSM device objects will contain only the MD5 hash stub for these endpoints. Any delta configuration operation between NSM and the

device will indicate identical configurations because the MD5 hash in NSM will match the file on the device. For the same reasons, an Update Device directive will have no effect on the device.

Related Documentation

- Uploading and Linking Large Binary Data Files with NSM

Uploading and Linking Large Binary Data Files with NSM

This topic describes the complete procedure for downloading a large binary data file and linking that file into the configuration tree.

To upload and link a large binary data file:

1. In the Device Manager, right-click the device icon and select Import Device from the list to Import the configuration.

   When the import job is finished, the device object configuration contains the MD5 stubs for each of the large binary data files.

2. Upload each required large binary data file onto the NSM client workstation.

   You'll need to get some files from the device. Other files, such as ESAP configuration files, should be downloaded form the site of origin. Use the device Web UI to upload binary files from Connect Secure.

3. To create a shared object in the NSM Object Manager for the binary file:

   1. In the Configure panel of the NSM navigation tree, select **Object Manager > Binary data** , and then click the Add icon.

   2. In the Binary Data dialog box, enter a name for the object, select a color for the object icon, add a comment if desired, and select the file you uploaded in step 2. Click OK.

3. Link the shared object to the corresponding node in the device configuration tree:

   1. In the Device Manager, double-click the Connect Secure device to open the device editor, and then select the Configuration tab.

   2. Navigate to the node in the configuration where you want to load the binary file.

      For example, to load an ESAP package, expand Authentication and then select Endpoint Security. In the Host Checker tab, select Endpoint Security Assessment Plug-Ins, and then click the Add icon.

   3. Select the shared object.

      To continue the ESAP example, in the New Endpoint Security Assessment Plug-Ins dialog box, enter a version number, select a shared binary data object from the Path to Package list. This list includes all shared binary data objects. Click OK.

      If the object you want is not in the list, you can add it to the shared binary data list by clicking the Add icon. The Binary Data dialog box appears as in step 3.

   4. Click OK to save the newly configured links.

Related Documentation

- Managing Large Binary Data Files

Importing Custom Sign-In Pages with NSM

The customized sign-in pages feature is a licensed feature that enables you to use your own access pages rather than having to modify the sign-in page included with the system.

To create a link from a configuration tree to a shared object containing a custom sign-in access page:

1. In the Device Manager, double-click the Connect Secure device to open the device editor, and then select the Configuration tab.

2. Expand Authentication.

3. Expand Signing-In.

4. Expand Sign-in Pages.

5. Select Users/Administrator Sign-in Pages, and then click the Add icon in the right pane.

6. Enter a name for the access page.

7. Select Custom Sign-in Pages.

8. Select a shared binary data object from the Custom Pages Zip File list.

9. Click OK once to save the link, and again to save the configuration.

To create a link from a configuration tree to a shared object containing a custom sign-in meeting page:

1. In the Device Manager, double-click the Connect Secure device to open the device editor, and then select the Configuration tab.

2. Expand Authentication.

3. Expand Signing-In.

4. Expand Sign-in Pages.

5. Select Meeting Sign-in Pages, and then click the Add icon in the right pane.

6. Enter a name for the sign-in meeting page.

7. Select Custom Sign-in Page.

8. Select a shared binary data object from the Blob list.

9. Click OK once to save the link, and again to save the configuration.

Related Documentation

- Configuring Sign-In Pages

Importing Antivirus LiveUpdate Settings with NSM

Retrieve the latest AV liveupdate file from the Pulse Secure Downloads web site at https://download.juniper.net/software/av/uac/epupdate_hist.xml.

Retrieve the latest patch file from the Pulse Secure Download web site at https://download.juniper.net/software/hc/patchdata/patchupdate.dat.

To create a link from a configuration tree to a shared object containing an antivirus (AV) liveupdate file:

1. In the Device Manager, double-click the Connect Secure device to open the device editor, and then select the Configuration tab.

2. Expand Authentication.

3. Select Endpoint Security.

4. From the Host Checker tab, select Live Update Settings.

5. Select a shared binary data object from the Manually import virus signature list.

6. Click OK to save the configuration.

To create a link from a configuration tree to a shared object containing an AV patch liveupdate file:

1. In the Device Manager, double-click the Connect Secure device to open the device editor, and then select the Configuration tab.

2. Expand Authentication.

3. Select Endpoint Security.

4. From the Host Checker tab, select Live Update Settings.

5. Select a shared binary data object from the Manually import patch management data list.

6. Click OK to save the configuration.

Related Documentation

- Uploading a Third-Party Host Checker Policy with NSM

- Host Checker and Trusted Network Computing

Importing Endpoint Security Assessment Plug-in (ESAP) Packages with NSM

The Endpoint Security Assessment Plug-in (ESAP) on Connect Secure checks third-party applications on endpoints for compliance with the predefined rules you configure in a Host Checker policy.

To download the Endpoint Security Assessment Plug-in from the Pulse Secure Global Support Center (PSGSC) Center to your NSM client computer:

1. Open the following page:

   http://www.pulsesecure.net/support

2. Click the Software tab.

3. Navigate to the ESAP release you want and click the link to download the package file to your computer.

To create a link from a configuration tree to a shared object containing an ESAP package:

1. In the Device Manager, double-click the Connect Secure device to open the device editor, and then select the Configuration tab.

2. Expand Authentication.

3. Select Endpoint Security.

4. From the Host Checker tab, select Endpoint Security Assessment Plug-Ins, and then click the Add icon.

5. In the New Endpoint Security Assessment Plug-Ins dialog box, enter an ESAP version number.

6. Select a shared binary object from the Path to Package list.

7. Click OK once to save the link, and again to save the configuration.

Uploading a Third-Party Host Checker Policy with NSM

For the device to recognize a package definition file, you must:

1. Name the package definition file **MANIFEST.HCIF** and include it in a folder named **META-INF**.

2. Create a Host Checker policy package by creating a zip archive. The archive should include the META-INF folder that contains the MANIFEST.HCIF file along with the interface DLL and any initialization files.

   For example, a Host Checker policy package might contain:

   - META-INF/MANIFEST.HCIF

   - hcif-myPestPatrol.dll

   - hcif-myPestPatrol.ini

3. Upload the Host Checker package (or packages) to the NSM shared object. You can upload multiple policy packages to NSM shared objects, each containing a different MANIFEST.HCIF file.

   > ℹ️ NOTE: After you upload a Host Checker policy package to the NSM shared object, you cannot modify the package contents. Instead, you must modify the package on your local system and then upload the modified version to NSM.

4. Implement the policy at the realm, role, or resource policy levels using the options.

   If you want to verify that the package itself is installed and running on the client computer (as opposed to a specific policy in the package passing or failing) you can use the name you specified when you uploaded the policy package (for example, myPestPatrol). To enforce a particular policy in the package, use the syntax **package-name.policy-name**. For example, to enforce the FileCheck policy in the myPestPatrol package, use myPestPatrol.FileCheck.

Related Documentation

- Host Checker and Trusted Network Computing

Linking to a Third-Party Host Checker Policy Shared Object with NSM

To create a link from a configuration tree to a shared object containing a third-party host checker policy:

1. In the Device Manager, double-click the Connect Secure device to open the device editor, and then select the Configuration tab.

2. Expand Authentication.

3. Select Endpoint Security.

4. From the Host Checker tab, select the Settings tab, and then click the Add icon in the Policies box.

5. From the Policy type list, select 3rd Party Policy.

6. Give the policy a name.

7. Select a shared binary data object from the Package list.

8. Click OK to save the configuration.

Related Documentation

- Host Checker and Trusted Network Computing

Importing Hosted Java Applets with NSM

You can store Java applets of your choice as shared objects in NSM without using a separate Web server to host them. You can then use these applets to intermediate traffic to various types of applications through Connect Secure. For example, you can upload the 3270 applet, 5250 applet, or Citrix Java applet to shared NSM objects. These applets enable users to establish sessions to IBM mainframes, AS/400s, and Citrix MetaFrame servers through terminal emulators. To enable the Citrix Java ICA client through the system session, you must upload multiple Citrix .jar and .cab files or configure a Citrix Terminal Services resource profile to host the Java applets.

You can upload individual .jar and .cab files or .zip, .cab, or .tar archive files to NSM shared objects. Archive files can contain Java applets and files referenced by the applets. Within the .zip, .cab, or .tar file, the Java applet must reside at the top level of the archive.

To ensure compatibility with both Sun and Microsoft Java Virtual Machines (JVMs), you must upload both .jar and .cab files. The Sun JVM uses .jar files, whereas the Microsoft JVM uses .cab files.

> *i* NOTE: When you upload Java applets to NSM, NSM asks you to read a legal agreement before it finishes installing the applets. Please read this agreement carefully. It obligates you to take full responsibility for the legality, operation, and support of the Java applets that you upload.
>
> Uploading Java applets requires signed ActiveX or signed Java applets to be enabled within the browser to download, install, and launch the client applications.

To create a link from a configuration tree to a shared object containing a Java applet:

1. In the Device Manager, double-click the Connect Secure device to open the device editor, and then select the Configuration tab.

2. Expand Users.

3. Expand Resource Profiles.

4. Select Hosted Java Applets, and then click the Add icon in the right pane.

5. Give the applet and file each a name.

6. Select a shared binary data object from the Applet file to be uploaded list.

7. Click OK once to save the link, and then again to save the configuration.

Related Documentation

- About Hosted Java Applet Templates

- Task Summary: Hosting Java Applets

Importing a Custom Citrix Client .cab File with NSM

The custom Citrix client file enables you to provision the Citrix client from Connect Secure instead of requiring that it be pre-installed on end-user machines or downloaded from some other web server.

To create a link from the configuration tree to a shared object containing a Custom Citrix .cab file:

1. In the Device Manager, double-click the Connect Secure device to open the device editor, and then select the Configuration tab.

2. Expand Users.

3. Select User Roles.

4. Select the Global Role Options tab.

5. In the Global Terminal Services Role Options tab, select a shared binary data object from the Citrix Client CAB File list.

6. Click OK to save the configuration.

Related Documentation

- Configuring a Citrix XenDesktop Resource Policy

- Creating Resource Profiles Using Citrix Web Applications

Introducing Pulse Secure Client

Pulse Secure client is an extensible multiservice network client that supports integrated connectivity and secure location-aware network access. Pulse Secure client simplifies the user experience by letting the network administrator configure, deploy, and control the Pulse client software and the Pulse connection configurations that reside on the endpoint.

Pulse Secure comprises client and server software. The client enables secure authenticated network connections to protected resources and services over local and wide area networks. The Pulse Secure client software can connect with Pulse Connect Secure to provide remote access to enterprise and service provider networks. Pulse also delivers secure, identity-enabled network access control (NAC) for LAN-based network and application access when it is

deployed with Pulse Policy Secure. Pulse also integrates with Pulse Collaboration Suite for online meeting services.

Users of mobile devices (smart phones and tablets) can install the Pulse mobile device app from the respective app stores for secure connectivity to Pulse Connect Secure. Mobile device users can also enable an optional security component, the Pulse Mobile Security Suite. Windows 8.1 (Pro and RT) introduced a Pulse Secure VPN client as part of the operating system.

Pulse Secure Client for Windows

The Pulse Secure client for Windows client interface (see Figure 2) lists the deployed Pulse connections. Each connection is a set of properties that enables network access through a specific Pulse server. The user can expand a connection to see more details about the connection.

Figure2: Pulse Secure Client for Windows Client Interface
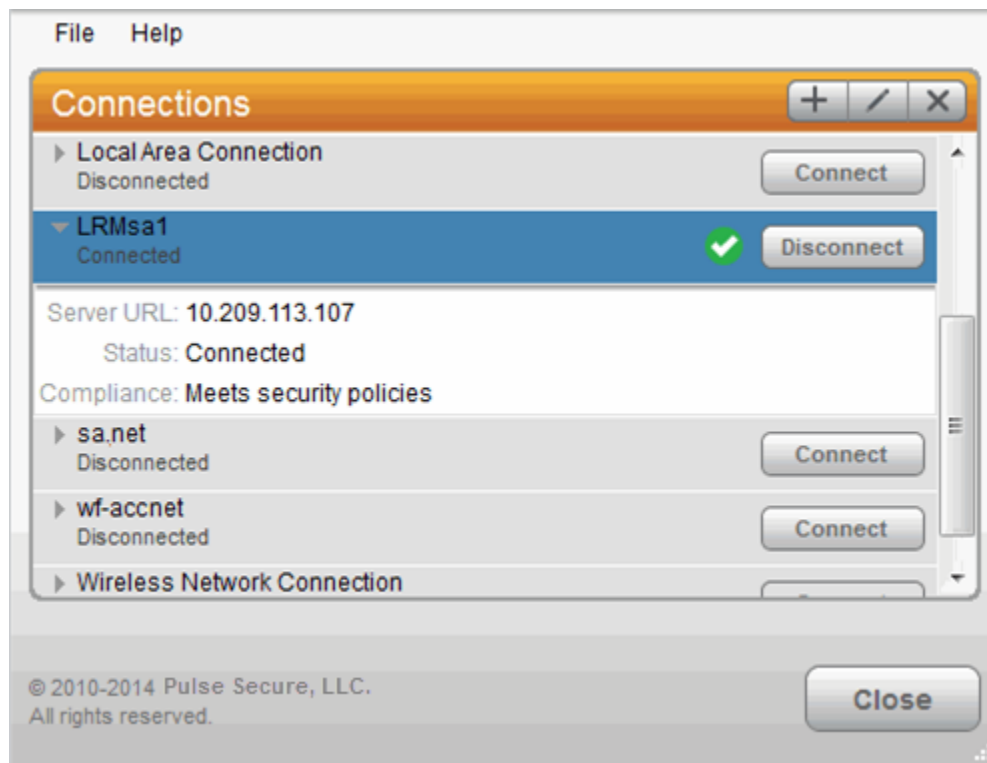
Pulse Secure client also displays a system tray icon that provides connection status, and can allow the user to connect and disconnect and enables quick access to the program interface. One tray icon provides status for all active connections.

Typically, the network administrator defines and deploys the Pulse connections but you can also enable users to define, edit, and remove their own connections.

Table 2 lists the Pulse connection status icons:

Table2: Pulse for Windows Connection Status

✅	Connected.
✅	Connecting.
✅	Connected with limitations
❗	Connection attempt failed.
✅	Connection suspended.
✅	Connected to the local network but no Internet access available.

Pulse Secure client supports the Federal Information Processing Standard (FIPS), which defines secure communication practices for the U.S. government. If FIPS is enabled on the endpoint, **FIPS On** appears near the bottom the Pulse window.

Pulse Secure Client for Mac OS X

Pulse supports Apple computers running Mac OS X. You deploy Pulse to Mac endpoints the same way you deploy the Windows client. Figure 3 shows the Pulse for Mac client interface.

Figure3: Pulse Secure client for Mac Client Interface

Pulse for Mac endpoints supports the following:

- Connections to Pulse Policy Secure

- Connections to Pulse Connect Secure

  Pulse clients connect to the Pulse Connect Secure in SSL fallback mode.

- Connections to Juniper Networks SRX Series gateways.

  - Mac OS X endpoints can connect to SRX Branch series SRX100-SRX650 gateways that are running Junos OS Release 10.x or later, and that have dynamic VPN access enabled and configured. SRX gateways do not support deployment of the Pulse Secure client.

  - Requires Pulse Secure client for Mac 5.0R3 or later and OS X 10.8 or later.

  - Pulse for Mac clients connect to the gateway as an IPsec IKEv1 VPN connection.

  - Pulse Secure Dynamic VPN functionality is compatible with SRX-Branch (SRX100-SRX650) devices only. SRX Data Center (SRX1400-SRX5800 – also called SRX HE or High End) devices do not support Pulse Secure Dynamic VPN from either Windows or Mac clients.

  - On Mac OS X clients, Pulse IPsec connections to SRX are unable to use the DNS IP address supplied by the SRX.

- Host Checker

  Host Checker for Mac OS X supports the following rules and remediation actions:

  - Port

  - Process

  - File

  - Custom IMC

  - Enable Custom Instructions

  - Kill Processes

  - Delete Files

  - Send reason strings

User Experience

From the user perspective, Pulse Secure client presents a clean, uncomplicated interface. The user can enter credentials, select a realm, save settings, and accept or reject the server certificate. When you configure the client, you can specify whether to permit end users to modify settings, such as by adding connections.

i   NOTE: If your Pulse environment uses Security Assertion Markup Language (SAML) for a Single Sign-on (SSO) authentication environment, the Pulse user sees a credential dialog box that is served (HTTPS) from the Pulse server instead of the local Pulse client credential dialog box. The sign-in page is defined as part of the

sign-in policy on the Pulse server and the Pulse client embeds the sign-in page within a Pulse client dialog box.

The client displays the connection status until the connection is made. If a connection fails as a result of the endpoint failing a Host Checker policy, Host Checker reason strings and remediation options appear.

Location Awareness

The location awareness feature enables you to define connections that are activated automatically based on the location of the endpoint. Pulse determines the location of the endpoint by evaluating rules that you define. For example, you can define rules to enable Pulse Secure client to automatically establish a secure tunnel to the corporate network through Pulse Connect Secure when the user is at home, and to establish a Pulse Policy Secure connection when the user is in the office and connected to the corporate network over the LAN. Pulse does not re-establish a VPN tunnel when the endpoint re-enters the trusted/corporate network. Location awareness rules are based on the client's IP address and network interface information. Split tunneling must be enabled for location awareness.

Session Migration

If you configure your access environment to support the Pulse Secure client session migration feature, users can log in once through a Pulse server on the network, and then securely access additional Pulse servers without needing reauthentication. For example, a user can connect from home through Pulse Connect Secure, and then arrive at work and connect through Pulse Policy Secure without having to log in again. Session migration also enables users to access different resources within the network without repeatedly providing credentials. IF-MAP Federation is required to enable session migration for users.

Centralized Pulse Configuration Management

Centralized configuration management is a key feature of Pulse Secure client. Pulse connection sets (the configurations that define how and when a Pulse client connects), are *bound* to a particular Pulse server. The binding server is the one that provides the initial Pulse configuration to the Pulse client. For example, if you create a Pulse connection set on ServerA, and then distribute those connections to endpoints, those clients are bound to ServerA.

A bound client is managed by its particular Pulse server. The Pulse administrator defines the Pulse Secure Connections and software components that are installed on the endpoint. When the Pulse client connects to the Pulse server that is managing it, the server automatically provisions configuration and software component updates. The administrator can permit the user to add, remove, and modify connections. The administrator can also allow dynamic connections (connections that are added by Pulse servers when the user logs into the server using a browser). A dynamic connection enables a bound client to add connections from Pulse servers other than the one the client is bound to. Dynamic connections are created as manual rather than automatic connections, which means that they are run only when the user initiates the connection or the user browses to a Pulse server and launches Pulse from the server's Web interface. Dynamic connections create the connection with the minimum configuration required to make the connection, which means that the URL used to install or launch Pulse from the Pulse server's Web interface is used as the Connection URL and connection name.

Binding Pulse Secure clients to a particular server ensures that the client does not receive different configurations when it accesses other Pulse servers. A bound endpoint receives connection set options and connections from its binding server, but it can have its Pulse client software upgraded from any Pulse server that has the automatic upgrade option enabled. (SRX gateways do not support Pulse software updates.)

> **NOTE:** A Pulse Secure client can be bound to only one Pulse server connection set at a time. The Pulse Secure client can receive updates and changes to that bound connection set from other Pulse servers only if the connection set is exported from the Pulse server and then imported to another Pulse server.

A Pulse client does not need to be bound to a Pulse server. An *unbound client* is managed by its user. If the Pulse Secure client software is installed without any connections, the user must add connections manually. Dynamic connections can be added by visiting the Web portals of Pulse servers. An unbound client does not accept configuration updates from any Pulse server.

Adding a Pulse Configuration to a New Pulse Installation explains the binding process in more detail.

Smart Connections - List of URLs

Each Pulse connection that connects to Pulse Policy Secure or Pulse Connect Secure can be configured with a list of Pulse servers. The Pulse client attempts to connect to each of the servers in the URL list until it succeeds. You can choose different modes to control the behavior of a Pulse connection that is starting from a disconnected state, start at the top of the list, start with the most recently connected URL, or choose randomly. The random option helps distribute the connection load across different Pulse servers. If a Pulse connection that is already established gets disconnected, for example, the wireless connection is interrupted, Pulse always tries to connect to the most recently connected URL. If that connection fails, then Pulse uses the server list. The Pulse user can also choose a connection from the list as shown in Figure 4.

Figure4: Pulse for Windows Client with a List of Connection URLs
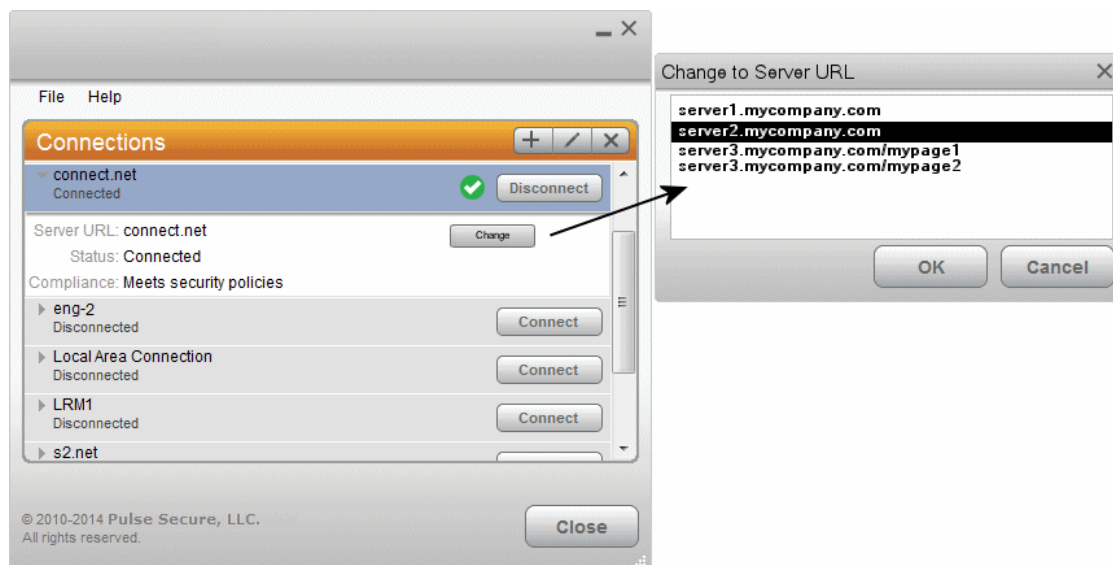
Security Certificates

Users cannot add CA servers or manage the server list. The Pulse client handles certificates in the same way that a browser handles certificates. If the Pulse dynamic certificate trust option is enabled for a connection, the user can accept or reject the certificate that is presented if it is not from a CA that is defined in the endpoint's certificate store.

Compliance and Remediation

Pulse supports the Host Checker application to assess endpoint health and update critical software. Host Checker is a client-side agent that is based on Trusted Network Connect standards. You configure rules in Host Checker policies for Pulse Connect Secure and Pulse Policy Secure to specify the minimum criteria for the security compliance of endpoints that are allowed to enter the network. Endpoints that fail can be connected through a remediation role that provides limited access.

Host Checker can be deployed from a Pulse server to Pulse clients on Windows and Mac OS X endpoints. It will be downloaded and run when a browser is used on a Windows or Mac OS X endpoint to connect to the Pulse server Web portal. You can use Host Checker policies at the realm or role level.

Host Checker for mobile clients (iOS, Android, and Windows Phone) is included as part of the Pulse app. Host Checker runs on the mobile client if Host Checker policies are configured and enabled on the server.

NOTE: Checker is not supported in the use case where the user employs a browser on the mobile device to connect to the Pulse server Web portal.

For Windows and OS X clients, you can use Host Checker to perform the following:

- Virus signature monitoring
  You can configure Host Checker to monitor and verify that the virus signatures, operating systems, and software versions installed on client computers are up to date. You can configure automatic remediation for those endpoints that do not meet the specified criteria.

- Patch management information monitoring and patch deployment
  You can configure Host Checker policies that check for Windows endpoints' operating system service pack, software version, or desktop application patch version compliance.

- Patch verification remediation options

  Pulse and Host Checker support endpoint remediation through Microsoft System Management Server or Microsoft System Center Configuration Manager (SMS/SCCM). With SMS/SCCM, Pulse triggers a preinstalled SMS/SCCM client to get patches from a pre-configured server.

- Endpoint configuration

  You can configure custom rules to allow Host Checker to check for third-party applications, files, process, ports, registry keys, and custom DLLs.

Pulse mobile clients support a set of Host Checker functions that vary from one OS to the next. For complete information on Host Checker for mobile clients, see *Implementing Host Checker Policies for Pulse Secure client for iOS Devices*, *Implementing Host Checker Policies for Pulse Secure for Android Clients*, and *Host Checker for Pulse Secure for Windows Phone*.

Two Factor Authentication

Pulse supports RSA SecurID authentication through soft token, hard token, and smart card authenticators. The SecurID software (RSA client 4.1 and later) must already be installed on the client machine.

Captive Portal Detection

Public WiFi locations often deploy a captive portal that requires the user to enter authentication information or to accept terms of service before network access is granted. Pulse detects the presence of captive portals and does not initiate a connection to a Pulse Connect Secure or Policy Secure server until internet access is granted. Pulse displays appropriate status information to enable the user to establish the portal and network connections.

Captive portal detection notes:

- Captive portal detection is supported on Pulse for Windows and Pulse for Mac. Captive portal detection is not supported on Windows In-Box Pulse Secure client or Pulse Secure client for mobile devices.

- If Pulse connects through a proxy, the detection algorithm is disabled, and Pulse indicates a failed connection attempt to the user.

- SRX connections do not support captive portal detection.

Pulse Client Collaboration Suite Integration

Pulse Collaboration Suite is accessible through the Pulse interface on Windows, Mac OS X, Android, and iOS. (Android clients must be R4.0 or later. iOS clients must be R3.2 or later.) Pulse Collaboration Suite enables users to schedule and attend secure online meetings. In meetings, users can share their desktops and applications with one another over a secure connection. Meeting attendees can collaborate by enabling remote control of their desktops and through text chatting.

Sign In Notifications

The notifications feature on Pulse Connect Secure and Pulse Policy Secure allows the network administrator to display notifications to Pulse client users prior to the user logging in and after the user has already logged in. For example, you could display a legal statement or a message stating who is allowed to connect to the server before you display the Pulse credentials dialog. After the user has connected, you could display a message that notifies the user of scheduled network or server maintenance or of an upcoming company meeting.

Automatic Software Updates

After you deploy Pulse Secure client software to endpoints, software updates occur automatically. If you upgrade the Pulse Secure client configuration on the server, updated

software components are pushed to a client the next time it connects. You can disable this automatic upgrade feature.

> **NOTE:** The automatic update feature is supported on Pulse Connect Secure and Pulse Policy Secure servers only. SRX gateways do not support automatic Pulse software updates.

> **NOTE:** If you configure Pulse Secure client to make 802.1X-based connections, a reboot might be required on Windows endpoints.

Pulse Secure Client Customization and Rebranding

The Pulse Secure client customization tool (**BrandPackager**) enables you to customize the appearance of the Pulse Secure Windows and Apple OS X clients. You can add your own identity graphic to the Pulse splash screen, to the program interface, and to Windows credential provider tiles. Figure 5 shows graphic customizations applied to the Pulse for Windows client. You can also customize error and informational message text, the text that appears in dialog boxes and on buttons, and make limited changes to Pulse online Help. For example, you might want to add your help desk phone number to Pulse error messages and the Pulse online Help.

BrandPackager is available for download from juniper.net.

Figure5: Pulse for Windows Client with a List of Connection URLs
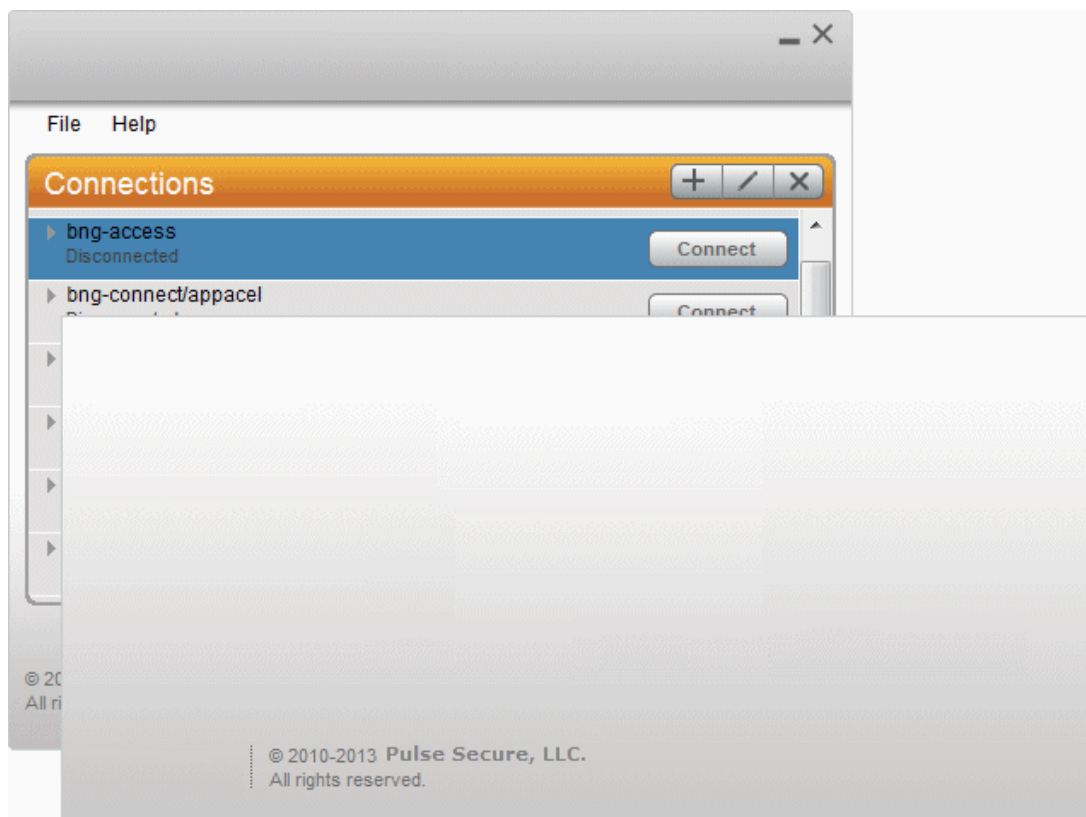
Related Documentation

- *Pulse Secure Client Status Icons*

- *Pulse Secure Client for Mobile Devices Overview*

- *Customizing Pulse Secure Client Overview*

Pulse Secure Client Configuration Overview

You configure the Pulse Secure client settings on the Pulse server so that when users request authentication, they are assigned a role based on the role mappings and optional security profile that you create. Access to specific resources is permitted only for users and devices that provide the proper credentials for the realm, that are associated with the appropriate roles, and whose endpoints meet security restrictions. If a user attempts to connect to the network from an endpoint that does not comply with the security restrictions you have defined, the user cannot access the realm or role.

As you plan your Pulse configuration, be sure you know how you want to deploy Pulse Secure client. You can use one or more of the following Pulse Secure client deployment options:

- Use the defaults or make changes to the Pulse Secure client default component set and default connection set, and then download and distribute Pulse by having users log in to the gateway's user Web portal and be assigned to a role. After the installation is complete, users have all the connections they need to access network resources.

- Create connections that an endpoint needs for connectivity and services, download the Pulse settings file (.jnprpreconfig), download default Pulse .msi installation program, and then run the .msi installation program by using an msiexec command with the settings file as an option. You can use the msiexec command to deploy Pulse using a standard software distribution process, such as SMS/SCCM.

- Distribute Pulse Secure client with no preconfiguration. You can download the default Pulse Secure client installation file (Mac or Win) from the device, and then distribute the file to endpoints using your organization's standard software distribution methods. Because the installer does not contain preconfigured connections, users must define network connections manually. Or you can create dynamic connections on each access gateway. These connections are automatically downloaded to the installed Pulse client when users provide their login credentials to the gateway's user Web portal.

The following tasks summarize how to configure Pulse Secure client on the device:

- Create and assign user roles to control who can access different resources and applications on the network. If you are converting your access environment from agentless or a VPN Tunneling environment, you should create new roles that are specific for Pulse Secure client.

- Define security restrictions for endpoints with Host Checker policies.

- Define user realms to establish authentication domains. If you are converting your access environment from agentless or a NC environment, typically you can use your existing realms.

- Associate the roles with appropriate realms to define your access control hierarchy using role mapping.

- Define Pulse Secure client component sets, connection sets, and connections.

- Deploy Pulse Secure client to endpoints.

Related Documentation

- Introducing Pulse Secure Client

- *Client Connection Set Options for Pulse Policy Secure*

- *Creating a Client Connection Set for Pulse Connect Secure*

- Creating a Client Component Set for Pulse Connect Secure

- Configuring a Role for Pulse Connect Secure

Configuring a Role for Pulse Connect Secure

A user role defines session settings and options, personalization settings (user interface customization and bookmarks), and access features (Web, file, application, Telnet/SSH, Terminal Services, network, meeting, and e-mail access). A user role does not specify resource access control or other resource-based options for an individual request. For example, a user role can define whether a user can perform Web browsing when the user is connected through the Pulse Connect Secure server Web portal. However, the individual Web resources that a user can access are defined by the Web resource policies that you configure separately.

The following procedure describes the role configuration options.

To create a role for Pulse Secure endpoints:

1. Select **Users > User Roles > New User Role** in the admin console.

2. Enter a name for the role and, optionally, a description. This name appears in the list of roles on the Roles page.

3. Under Client Options, select **Pulse Secure**.

   When this option is enabled, the Pulse Secure button appears on the Connect Secure Web portal. When a user clicks it, Pulse is downloaded and installed on the user's endpoint.

   Enabling this option alone does not enable Pulse as the client for the role. This option works in conjunction with the settings you enable in the Access Features section and then configure on the respective role tabs. The combination of settings determines whether you enable Pulse Secure, Pulse Secure Application Manager (SAM), or Network Connect. The following procedures describe how to enable each client option.

   To enable Pulse Secure client:

   1. In the role's **General > Overview > Options** section select **Pulse Secure**.

      This setting applies to both Windows and Apple OS X versions of Pulse Secure.

   2. In the Access Features section select **VPN Tunneling**.

      The VPN Tunnelling tab enables you to specify split tunnel behavior, specify the Pulse component set, and enable 3rd-party software integrations.

   To enable Pulse Secure client for SAM:

1. In the Options section select **Pulse Secure**.

2. In the Access Features section select **Secure Application Manager** and then select **Windows version**.

   The SAM tab enables you to specify applications and servers secured by SAM.

To enable Network Connect:

1. In the Options section make sure **Pulse Secure** is disabled.

2. In the Access Features section select **VPN Tunneling**.

3. Click **Save Changes**. Role configuration tabs appear.

   NOTE: When the Pulse Secure option is enabled and no other access method (VPN Tunneling, WSAM) is enabled, then no client will be delivered.

Configuring General Role Options for Pulse Connect Secure

The General tab includes options for detailed control of how the client interacts with the server and the network. The following describes the options that apply to Pulse Secure client.

General > Restrictions

Source IP—Control from which IP addresses users can access the Web portal sign-in page, be mapped to a role, or access a resource.

Browser—Allow or deny access to the role based on the browser's user agent string.

Certificate—Allow all users or only users with a signed client-side certificate.

Hot Checker—Select configured Host Checker policies to enforce with this role.

General > VLAN/Source IP

VLAN and Select Source IP—To direct traffic to specific sites based on the role, you can define a source IP alias for each role and then use the alias to configure virtual ports you define for the internal interface source IP address. A back-end device can then direct end user traffic based on the alias. This capability enables you to direct various end users to defined sites based on their roles, even though all of the end user traffic has the same internal interface source IP address.

General > Session Options

Idle Timeout—The maximum time a session can remain idle (no traffic) before the server ends the session.

Max. Session Length—The maximum time for a session before the server ends the session.

Reminder Time—When the Enable Session Extension feature is enabled, the Reminder Time specifies the number of minutes prior to a session end when the server sends a notice through Pulse and notifies the user that the session will end soon.

Enable Session Extension—Allows the user to extend the session. The user can choose to extend the session at any time by selecting a menu option in the Pulse client interface. If the Session Timeout Warning is selected. A notice message appears when the Reminder Time is reached and the user can choose to extend the session from within that notice message.

Enable Session Timeout Warning—Enables or disables the session timeout warning, which notifies the user when their Pulse VPN session is close to expiring. The Reminder Time value specifies the point at which the reminder appears.

Roaming Session—Select one of the following options to specify the client's roaming behavior:

- **Enabled**—A roaming session allows a user to retain connectivity when moving a device, such as a laptop with a dynamic IP address, from one subnet to another. Disable this feature to prevent users from accessing a previously established session from a new source IP address. Disabling roaming can help protect against an attack that spoofs a user's session.

- **Limit to Subnet**—Limit the roaming session to the local subnet specified in the endpoint's IP configuration. Users can sign in from one IP address and continue using their sessions with another IP address as long as the new IP address is within the same subnet.

- **Disabled**—Disable roaming user sessions for users mapped to this role.

Browser Session Cookie—Select **Enabled** to remove the Connect Secure session cookie and log users out of their Connect Secure web session after the Pulse client is launched. Removing the browser session cookie enhances Pulse session security.

General > UI Options

UI Options—The settings on this page define the Pulse Connect Secure Web portal page.

SAM > Applications

Add Application—We recommend that you use resource profiles to specify the applications available to users, but you can use role and resource policy settings instead.

SAM > Options

Auto-uninstall Secure Application Manager—This feature is not applicable to the Windows Phone client. Users must download and install Pulse for Windows Phone before the Windows Phone device can connect to the Pulse Connect Secure.

Prompt for username and password for intranet sites—If you enable this option, the Pulse Connect Secure requires users to enter sign-in credentials before connecting to sites on your internal network. This option changes Internet Explorer's intranet zone setting so that Internet Explorer always prompts the user for network sign-in credentials for an intranet site.

Auto-upgrade Secure Application Manager—This feature is not applicable to the Pulse for Windows Phone app.

Resolve only hostnames with domain suffixes in the device DNS domains—If you enable this option, users can only browse to Web sites that are part of their login domain.

Session start script and Session end script—This feature is not applicable to the Pulse for Windows Phone app.

Configuring Role Options for Pulse Connect Secure

All of the options for role configuration tabs are described in *User Access Management Framework Feature Guide*.

To configure role options for Pulse Secure endpoints:

1. From the admin console, select **Users > User Roles**.

2. Click the role you want to configure and then click the VPN Tunneling tab.

3. Under Split Tunneling Options, select your options:

   General VPN Options apply to all Layer 3 VPN clients, Pulse Secure client (Windows, OS X, iOS, and Android), Network Connect, and third-party IKEv2 clients:

   - **Split Tunneling**—Split tunneling options let you define how network traffic flows on the client.

     Enable— Pulse modifies routes on the client so that traffic meant for the corporate intranet uses the virtual adapter created by Pulse (the Pulse tunnel) and all other traffic goes through the local physical adapter.

     Disable—When the Pulse session is established, predefined local subnet and host-to-host routes that might cause split-tunneling behavior are removed, and all network traffic from the client goes through the Pulse tunnel. With split tunneling disabled, users cannot access local LAN resources during an active VPN session. However, when you create a Pulse connection for users, if you allow users to override you can enable an option that allows the user to suspend their VPN connection. While a Pulse connection is in the suspended state, all traffic goes through the local physical adapter.

   Pulse Secure client Options apply only to Pulse Secure and Network Connect:

   - **Route Precedence**—You can define which routing table takes precedence:

     Tunnel Routes—The route table associated with the Pulse virtual adapter take precedence. Pulse overwrites the physical interface routes if there is conflict between the Pulse virtual adapter and the physical adapters. Pulse restores the original routes when the connection is ended.

     Tunnel Routes with local subnet access (Pulse on Windows and Mac OS X only)— Network traffic addressed to the networks defined in the split tunnel resource policies goes through the VPN tunnel. Network traffic that is addressed to the directly-connected (local) subnet goes to the local subnet. The default route is set to the local subnet so all other network traffic is subject to the original endpoint routing table.

     Endpoint Routes—The route table associated with the endpoint's physical adapter take precedence.

   - **Route Monitor**—Pulse can monitor the route tables and take appropriate action.

Yes – VPN tunneling ends the connection only if the route change affects the VPN tunnel traffic. For example, if the route metric is changed higher, it should not disconnect VPN tunneling.

No – Route tables are allowed to change on the client endpoint.

- **Traffic Enforcement**—When Traffic Enforcement is enabled, Pulse creates rules on the endpoint's firewall (Mac and Win) that ensure that all traffic conforms to the split tunneling configuration. For example, a local program might bypass the endpoint's routing tables and bind traffic to the physical interface instead of allowing it to go through the Pulse virtual interface. If you enable traffic enforcement, you ensure that all traffic is bound by the split tunneling configuration.

  IPv4 – When this check box is selected all IPv4 traffic is bound by the split tunneling configuration.

  IPv6 – When this check box is selected all IPv6 traffic is bound by the split tunneling configuration.

  Typically, if you wanted to enable traffic enforcement, you would enable both options. Enabling only one option is useful for nested tunnel (tunnel-in-tunnel) configurations. See the Pulse Secure Supported Platforms Guide for supported nested tunnel configurations.

- **Enable TOS Bits Copy**—Enables you to control the client behavior in networks that employ Quality of Service (QoS) protocols. When you enable this check box, the Pulse Secure client copies IP Type of Service (TOS) bits from the inner IP header to outer the IP Header. Note that enabling this option might require a reboot of the client endpoint when the client software is installed for the first time on Windows endpoints. Pulse Secure clients support TOS bit copy only for IPSec transport and not for SSL transport.

- **Multicast**—Enables the multicast feature on the client when this option is selected.

- **Auto-launch**—Activates the Pulse Secure client software automatically when the endpoint is started when this option is selected.

Options for Pulse Secure client on Windows apply only to Pulse Secure and Network Connect on Windows endpoints:

- **Launch client during Windows Interactive User Logon**—When this option is enabled, the Pulse Secure client starts when the user logs into Windows. Note that this setting is not the same as the Pulse connection settings that control machine authentication and credential provider authentication. Choose one of the following options:

  Require client to start when logging into Windows

  Allow user to decide whether to start client when logging into Windows

- **Windows: Session start script**—Specify a script (.bat, .cmd, or .exe) to run for users assigned to the role after Pulse Secure client connects with Pulse Connect Secure. For example, you can specify a script that maps network drives on an endpoint to shares on protected resources.

- **Windows: Session end script**—Specify a script (.bat, .cmd, or .exe) to run for users assigned to the role after Pulse Secure client disconnects from Pulse Connect Secure. For example, you can specify a script that disconnects mapped network drives. If there is no start script defined, or the start script has not been run, the end script does not run.

- **Skip if Windows Interactive User Logon Enabled**—Select this option to bypass the specified Windows session start script.

  If the client signs in to their Windows Domain via the Credential Provider automatic sign-in function, a script is executed by the Windows client. In this case, the sign-in script might be identical to the specified VPN Tunneling start script. You can use this option, therefore, as a way to avoid executing the same script twice.

  Options for Pulse Secure client on Mac apply only to Pulse Secure and Network Connect on Apple OS X endpoints:

- **Mac: Session start script**—Specify a script (.bat, .cmd, or .exe) to run for users assigned to the role after Pulse Secure client connects with Pulse Connect Secure. For example, you can specify a script that maps network drives on an endpoint to shares on protected resources.

- **Mac: Session end script**—Specify a script (.bat, .cmd, or .exe) to run for users assigned to the role after Pulse Secure client disconnects from Pulse Connect Secure. For example, you can specify a script that disconnects mapped network drives. If there is no start script defined, or the start script has not been run, the end script does not run.

4. In the **Session scripts** area, optionally specify a location for the following:

- **Windows: Session start script**—Specify a script (.bat, .cmd, or .exe) to run for users assigned to the role after Pulse connects with Pulse Policy Secure. For example, you can specify a script that maps network drives on an endpoint to shares on protected resources. The script must be in a location (either local or on the network) that is accessible by the user.

- **Windows: Session end script**—Specify a script (.bat, .cmd, or .exe) to run for users assigned to the role after Pulse Secure client disconnects from Pulse Policy Secure. For example, you can specify a script that disconnects mapped network drives. If there is no start script defined, or the start script has not been run, the end script does not run. The script must be in a location (either local or on the network) that is accessible by the user.

5. Click **Save Changes**.

Configuring Role Options for Host Checker for Pulse Connect Secure

Host Checker options allow you to enable configured Host Checker policies, to choose one or more policies for the role, and to specify whether the endpoint must meet all or just one of the selected Host Checker policies. Before you can assign Host Checker policies for a role, you must have already defined the policies.

To configure Host Checker for a selected role:

1. For a selected role, select **General > Restrictions > Host Checker**.

2. Select the check box **Allow users whose workstations meet the requirements specified by these Host Checker policies**.

3. Click **Add** to move Host Checker policies from the **Available Policies** list to the **Selected Policies** list.

4. Select the check box **Allow access to the role...** to grant access if the endpoint passes any of the selected Host Checker policies.

5. Click **Save Changes**.

Related Documentation

- Creating Global Host Checker Policies

- *Configuring Pulse Secure for Secure Application Manager*

- *Machine Authentication for Pulse Connect Secure Overview*

- Specifying Role-Based Source IP Aliases

Pulse Connection Set Options for Pulse Connect Secure

A Pulse Secure client connection set contains network options and allows you to configure specific connection policies for client access to any Pulse server that supports Pulse Secure. The following sections describe each of the configuration options for a Pulse connection set.

Pulse Secure Connection Set Options

The following items apply to all connections in a connection set.

- **Allow saving logon information**—Controls whether the Save Settings check box is available in login dialog boxes in the Pulse client. If you clear this check box, the Pulse client always requires users to provide credentials. If you select this check box, users have the option of saving their credentials.

  The Pulse Secure client can retain *learned user settings*. These settings are retained securely on the endpoint, evolving as the user connects through different Pulse servers. The Pulse Secure client can save the following settings:

  - Certificate acceptance

  - Certificate selection

  - Realm

  - Username and password

  - Proxy username and password

  - Secondary username and password

  - Role

    i   NOTE: If the authentication server is an ACE server or a RADIUS server and authentication is set to Users authenticate using tokens or one-time passwords, Pulse ignores the Allow saving logon information option. If the user sees a username and token prompt and the Save settings check box is disabled. Pulse supports soft token, hard token, and smart card authentication.

  When a user opts to save settings, that information is used for each subsequent connection without prompting. If a setting changes (for example, if a user changes a password), the

saved setting is invalid and connection attempts fail. In this case, the user must use the client's Forget Saved Settings feature, which clears all user-saved settings.

- **Allow user connections**—Controls whether connections can be added by the user.

- **Display splash screen**—Clear this check box to hide the Pulse splash screen that normally appears when the Pulse client starts.

- **Dynamic certificate trust**—Determines whether users can opt to trust unknown certificates. If you select this check box, a user can ignore warnings about invalid certificates and connect to the target Pulse server.

- **Dynamic connections**—Allows connections within this connection set to be automatically updated or added to a Pulse Secure client when the user connects to the Pulse server through the user Web portal, and then starts Pulse through the Web portal interface. Dynamic connections are created as manual rather than automatic connections, which means that they are run only when the user initiates the connection or the user browses to a Pulse server and launches Pulse from the server's Web interface.

  If dynamic connections are disabled, and the user logs in through the Web portal of a Pulse server that is not already included in the Pulse client's connection set, then starting Pulse from the Web portal does not add a new Pulse connection for that Pulse server. If you choose to disable dynamic connections, you can still allow users to manually create connections by enabling **Allow User Connections**.

- **FIPS mode enabled**—Enable FIPS mode communications for all Pulse connections in the connection set. The Federal Information Processing Standard (FIPS) defines secure communications for the U.S. government. When a Pulse connection is operating in FIPS mode, **FIPS On** appears in the lower corner of the Pulse client interface. If the Pulse server hardware does not support FIPS mode operations, FIPS mode configuration options are not present in the admin console interface. FIPS mode operations are supported on MAG Series Pulse Secure Gateways and some SA series appliances. The device must be running Pulse Policy Secure R5.0 or later or Pulse Connect Secure R8.0 or later.

  > NOTE: Users cannot enable FIPS mode from within the Pulse client. You must create FIPS-enabled connections on the server and deploy them.

- **Wireless suppression**—Disables wireless access when a wired connection is available. If the wired connection is removed, Pulse enables the wireless connections with the following properties:

  - Connect even if the network is not broadcasting.

  - Authenticate as computer when computer information is available.

  - Connect when this network is in range.

    > NOTE: Wireless suppression occurs only when the wired connection is connected and authorized. If you enable wireless suppression, be sure to also configure a connection that enables the client to connect through a wired connection.

Policy Secure 802.1X Connection Type Options

Use this connection type to define authenticated connectivity to 802.1X devices, wired or wireless. Users cannot create 802.1X connections from the Pulse client interface. Users see 802.1X connections in the Pulse interface only when the connection has been deployed from the server and the specified network is available.

- **Adapter type**—Specifies the type of adapter to use for authentication: wired or wireless.

- **Outer username**—Enables a user to appear to log in anonymously while passing the actual login name (called the inner identity) through an encrypted tunnel. As a result, the user's credentials are secure from eavesdropping, and the user's inner identity is protected. In general, enter **anonymous**, which is the default value. In some cases, you might need to add additional text. For example, if the outer identity is used to route the user's authentication to the proper server, you might be required to use a format such as anonymous@acme.com.

  > **NOTE:** If you leave the box blank, the client passes the user's or the machine's Windows login name as the outer identity.

- **Scan list**—If you selected wireless as the adapter type, the scan list box is available to specify the SSIDs, including non-broadcast SSIDs, to connect to in priority order. If you leave the list empty, the user can connect to any available wireless network.

- **Support Non-broadcast SSID**—Allows a user to connect to a non-broadcast wireless network from within the Pulse interface. Selecting this field enables the following options:

  - **Wireless Security Cipher**—Specify the type of encryption used by the non-broadcast network:

    - TKIP

    - AES

If the non-broadcast SSID options are configured, the Pulse connection configuration includes the values and they are used to configure the wireless profile on the endpoint.

Trusted Server List (for Policy Secure 802.1X Connection)

For a Layer 2 connection that uses certificate authentication, make sure that the connection has at least one entry in the Trusted Server list.

- **Server certificate DN**—The server certificate distinguished name (DN) and its signing certificate authority (CA). An empty DN field allows a client to accept any server certificate signed by the selected CA.

Connect Secure or Policy Secure (L3) Connection Type Options

Use a Connect Secure or Policy Secure (L3) connection for a Layer 3 connection to Pulse Connect Secure or Pulse Policy Secure.

- **Allow user to override connection policy**—Allows a user to override the connection policy by manually connecting or disconnecting. Typically, you leave this option selected to make sure that a user can establish a connection under all conditions. If you disable this check box, the

user cannot change the endpoint's connection status, suspend/resume a connection to Pulse Connect Secure or shut down Pulse Secure client.

- **Support Remote Access (SSL VPN) or LAN Access (UAC) on this connection**—This option must be selected if this connection is for Pulse Policy Secure. If the connection is for Pulse Connect Secure, you can disable this check box and use the connection for accessing Pulse Collaboration meetings only by also selecting **Enable Pulse Collaboration integration on this connection**.

- **Enable Pulse Collaboration integration on this connection**—This option must be disabled if this connection is for Pulse Policy Secure. If the connection is for Pulse Connect Secure, you can enable this check box and use the connection for accessing Pulse Collaboration meetings.

- **Connect to URL of this server only**—Specifies whether the endpoint connects to this Pulse server exclusively or if it can connect to the any of the servers listed in the list of connection URLs. Disable this check box to enable the List of Connection URLs.

- **List of Connection URLs**—Allows you to specify a list of Pulse servers (Pulse Policy Secure or Pulse Connect Secure) for this connection. The Pulse client attempts to reach each server in the list, in the order listed, until it succeeds. You can specify up to 8 Pulse servers.

    NOTE: IF-MAP federation must be configured to ensure that a suspended session can be resumed to a different URL.

    NOTE: When this feature is used with Pulse Policy Secure, all of the Pulse servers in the list must be configured for failover, so that any one of them can provision the firewall enforcer.

Figure 6 shows how the Pulse user can select a server from the connection's list of URLs.

Figure6: Pulse for Windows Client with a List of Connection URLs
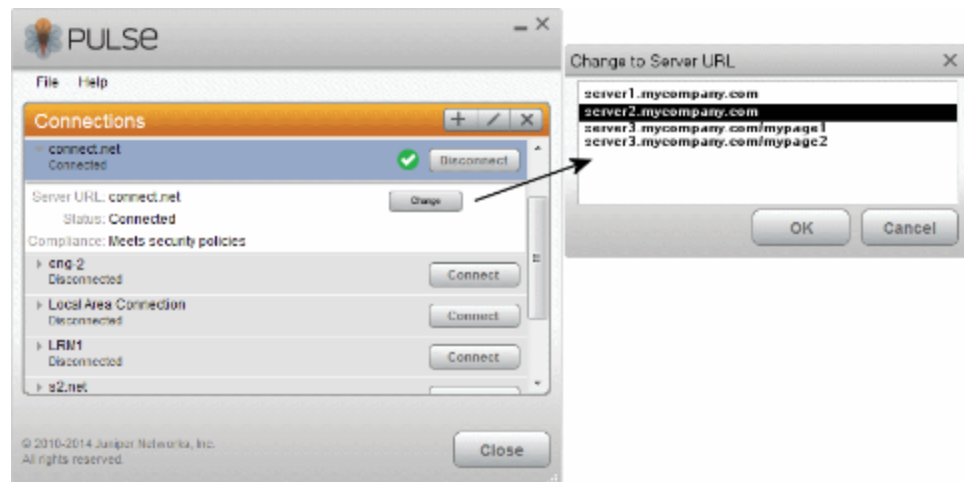
- **Attempt most recently connected URL first**—If you have specified a list of connection URLs, you can select this check box to have the Pulse client always attempt the most recent

successful connection. If that connection is not successful, Pulse then starts at the top of the list. The most recently connected URL is saved across reboots.

- **Randomize URL list order**—If you have specified a list of connection URLs, select this check box to have the Pulse client ignore the order in which the servers are listed. You can select this option to spread the connection load across multiple Pulse servers.

If you enabled Attempt most recently connected URL first, then Pulse attempts that connection first. If the connection attempt fails, Pulse chooses randomly from the list for the next connection attempt. During a credential provider connection attempt, Pulse chooses the URL automatically. It does not display a window to let the user choose a URL. Connections that use machine authentication ignore this option and always use the ordered list of connection URLs. Any preferred roles and realms you specify must be applicable to all of those servers. In the case of an interrupted connection, such as temporarily losing the WiFi link, Pulse always tries to reconnect to the most recently connected URL.

The Pulse servers should be configured for IF-MAP federation to ensure that a session can be resumed to a different URL. If the Pulse servers are not federated, then Pulse might prompt for credentials.

SRX (for Dynamic VPN) Connection Type Options

Use an SRX connection for a dynamic VPN connection to an SRX Series Services Gateway.

- **Address**—Specifies the IP address of the SRX Series device.

- **Allow user to override connection policy**—Allows users to override the connection policy by manually connecting or disconnecting. Typically, you leave this option selected to make sure that a user can establish a connection under all conditions. If you disable this check box, the user cannot change the endpoint's connection status or shut down Pulse Secure client.

Pulse Connection is Established Options

For all connection types, specify how the connection is established. The options vary according to the type of connection. Automatic connections include machine authentication and credential provider connections. Connections can be established using the following options.

> ℹ️ NOTE: All connections that are configured to start automatically attempt to connect to their target networks at startup time. To avoid multiple connection attempts, be sure that only one connection is configured to start automatically, or configure location awareness rules.

- Modes:

  - **User**—Enables user authentication.

  - **Machine** —Enables machine authentication, which requires that Active Directory is used as the authentication server and that machine credentials are configured in Active Directory. A machine connection is, by default, an automatic connection.

  - **Machine or user**—Enables machine authentication for the initial connection. After user authentication, the machine authentication is dropped. When the user logs out, the machine authentication connection is restored.

- Options:

    - **Connect automatically**—Connections are attempted when the conditions specified in the location awareness rules are true, and disconnected when the conditions are no longer true.

    - **Enable pre-desktop login (Credential provider)**—Enables Pulse client interaction with the credential provider software on the endpoint. The user credentials are used to establish the authenticated Pulse connection to the network, login to the endpoint, and login to the domain server.

Pulse Connection is Established Examples

The following configurations show how to select the Connection is established options of a Pulse connection set for specific user login behavior:

Figure7: Connect Manually

Figure8: Connect Automatically After User Signs in to the Desktop

Figure9: Connect Automatically When the Machine Starts; Machine Credentials are Used for Authentication

NOTE: When you use machine credentials for authentication and no user credentials, Pulse cannot perform user-based tasks. The following tasks can be run only when the user is logged in:

- Run session scripts

- Detect or modify proxy settings

- Run automatic Pulse client upgrade

- Install or upgrade Pulse components

Figure10: Connect Automatically When the Machine Starts; the Connection is Authenticated Again When the User Signs in to the Desktop

The configuration in Figure 10 enables machine authentication for the initial connection. After the user connects with user credentials, the machine authentication is dropped. When the user logs out, the machine authentication connection is restored. In one typical use case, the machine credentials provide access to one VLAN, and the user credentials provide access to a different VLAN.

NOTE: If the machine and user have different roles, each role should map to the same connection set. Otherwise, after the user connects, the existing connection set might be replaced.

Figure11: Connect Automatically at User Login

The configuration in Figure 11 enables Pulse client interaction with the credential provider software on the endpoint. The user credentials are used to establish the authenticated Pulse connection to the network, to log in to the endpoint, and to log in to the domain server.

Figure12: Connect Automatically When the Machine Starts; Connection is Authenticated Again at User Login
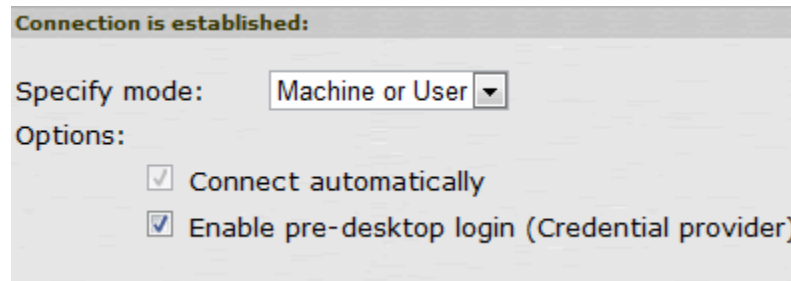
The configuration in Figure 12 enables Pulse client interaction with the credential provider software on the endpoint. Machine credentials are used to establish the authenticated Pulse connection to the network. When the user provides user credentials, the connection is authenticated again. In one typical use case, the machine credentials provide access to one VLAN, and the user credentials provide access to a different VLAN.

Location Awareness Rules

For Connect Secure or Policy Secure (L3) and SRX connections, you can define location awareness rules that enable an endpoint to connect conditionally. If you do not have location awareness rules defined, Pulse attempts to connect with each connection that is defined as an automatic connection until it connects successfully. Location awareness rules allow you to define an intelligent connection scheme. For example, the endpoint connects to Pulse Policy Secure if it is connected to the company intranet, or it connects to Pulse Connect Secure if it is in a remote location.

> **NOTE:** Split tunneling must be enabled to allow resolution of location awareness rules.

A Pulse connection uses the IP address of a specified interface on the endpoint to determine its network location. Each location awareness rule includes the following settings:

- **Name**—A descriptive name, for example, "corporate-DNS." A name can include letters, numbers, hyphens, and underscores.

- **Action**—The method the connection uses to discover the IP address. Choose one of the following values:

  - **DNS Server**—Allows the endpoint to connect if the endpoint's DNS server on the specified interface is set to one of the specified values. Use the Condition box to specify IP addresses or address ranges.

  - **Resolve Address**—Allows the endpoint to connect if the hostname specified in the DNS Name box can be resolved by the DNS server for the specified interface. If one or more address ranges are specified in the Address Range box, the address must resolve to one of the ranges to satisfy the expression.

  - **Endpoint Address**—Allows the endpoint to connect if the IP address of the specified interface is within a range specified in the IP Address Range box.

NOTE: To create a negative location awareness rule, you first create the positive state and then use rule requirement logic to use the rule as a negative condition.

Machine Connection Preferences

The Machine Connection Preferences appear if you have selected one of the machine authentication options for how the connection is established. Normally Pulse presents a selection dialog box if more than one realm or role is available to the user. However, a connection that is established through machine authentication fails if a dialog box is presented during the connection process. To suppress the selection dialogs, either ensure that only one role and realm is available to users, or specify a preferred realm and role for this connection.

- **Preferred Machine Realm**—Specify the realm that this connection uses when establishing the machine connection. The connection ignores any other realm available for the specified login credentials

- **Preferred Machine Role Set**—Specify the role or the name of rule for the role set that this connection uses when establishing the machine connection. The role or rule name used must be a member of the preferred machine realm.

User Connection Preferences

The User Connection Preferences options enable you to specify a realm and role for automatic connections that would otherwise present a selection dialog box to the user. To suppress the selection dialogs, either ensure that only one role and realm is available to users, or specify a preferred realm and role for this connection.

- **Preferred User Realm**—Specify the realm for this connection that is used when a user logs onto the endpoint. The connection ignores any other realm available for the user's login credentials

  If one of the credential provider connection options is enabled, the following options are available:

  - **Preferred Smartcard Logon Realm**—Preferred realm to be used when user logs in with a smart card.

  - **Preferred Password Logon Realm**—Preferred realm to be used when user logs in with a password.

    NOTE: Be sure that the authentication realms you specify exist, and that they support the appropriate login credential option.

- **Preferred User Role Set**—Specify the preferred role or the name of the rule for the role set to be used for user authentication. The role or rule name used must be a member of the preferred user realm.

- **Select client certificate from machine certificate store**—Enables you to specify the location of the client certificate on a Windows endpoint as part of a Pulse connection that verifies the identity of both the machine and the user before establishing a connection. When this check box is selected, the Pulse connection looks at client certificates located in the Local Computer personal certificate store. When this check box is not selected, the connection

accesses the user certificate store an a Windows endpoint. For more information, see *Machine and User Authentication Through a Pulse Connection for Pulse Connect Secure*.

Related Documentation

- *Pulse Connect Secure Overview*

- *Creating a Client Connection Set for Pulse Connect Secure*

- FIPS Supported Platforms

- *Machine and User Authentication Through a Pulse Connection for Pulse Connect Secure*

Creating a Client Connection Set

To create a client configuration:

1. From the admin console, select **Users > Pulse Secure > Connections**.

2. Click the **New** button.

3. Enter a name and optional description for this connection set.

    i   NOTE: You must enter a name, otherwise you cannot create a connection set.

4. Click **Save Changes**.

5. From the main Pulse Secure Connections page, select the connection set.

6. Under Options, select or clear the check box for each of the following items.

    - **Allow saving logon information**

    - **Allow user connections**

    - **Dynamic certificate trust**

    - **Dynamic connections**

    - **Wireless suppression**

7. Under connections, click **New** to define a new connection.

8. Enter a name and an optional description for this connection.

9. Select a Type for the connection. The Type identifies the device type for the connections and can be any of the following:

    - **Policy Secure (802.1X)**

    - **Connect Secure or Policy Secure (L3)**

    - **SRX**

10.   If you select **Policy Secure 802.1X** from the type list enter a value or select or clear the following check boxes:

    - **Adapter type**—Select wired or wireless.

    - **Outer username**—Enter the outer username.

    - If you select the wieless adapter type, you can specify the following:

- **Scan list**—Enter the SSIDs to connect to in your order of priority.

- **Support Non-broadcast SSID**—Select this option to allow users to connect to non-broadcast networks.

- **Wireless Security Algorithm** —Select the wireless security type (WPA or WPA2).

- **Wireless Security Cipher** —Select the wireless security type (AES or TKIP).

11. Click **Save Changes**.

12. If you select **Connect Secure or Policy Secure (L3)**, after Options select or clear the check box for each of the following items:

- **Allow user to override connection policy**

- **Support Remote Access (SSL VPN) or LAN Access (UAC) on this connection**—This option must be selected if this connection is for Policy Secure. If the connection is for Connect Secure, you can disable this check box and use the connection for accessing Pulse Collaboration meetings only by also selecting **Enable Pulse Collaboration integration on this connection**.

- **Enable Pulse Collaboration integration on this connection**—This option must be disabled if this connection is for Policy Secure. If the connection is for Connect Secure, you can enable this check box and use the connection for accessing Pulse Collaboration meetings.

- **Connect to URL of this server only**—This connection will use the URL of the server where you are creating the connection.

- **URL**—If you did not enable the check box for this server only, you must specify the URL of the server for the connection.

13. If you select **SRX**, enter an IP address in the **Address** box.

14. From the Options list, select or clear the following check boxes:

- **Allow user to override connection policy**

- **Connect automatically**

- **URL**—enter the network address for the firewall device.

15. (Optional) You can enable location awareness by creating location awareness rules. Location awareness can force a connection to a particular interface.

16. After you have created the client connection set, create a client component set profile, and select this connection set.

Related Documentation

- *Pulse Secure Component Set Options*

- Creating a Client Component Set for Pulse Connect Secure

- Configuring Connection Rules for Location Awareness

Configuring Connection Rules for Location Awareness

The location awareness feature enables a Pulse client to recognize its location and then make the correct connection. For example, a Pulse client that is started in a remote location automatically connects to the device. But that same client automatically connects to a device when it is started in the corporate office.

> NOTE: Location awareness and session migration are similar because they both simplify connectivity for the user, but they do so under different conditions. With location awareness, the Pulse client makes a decision on where to connect when a user logs in to the computer. Session migration occurs when the user puts the computer into a standby or hibernate mode without first logging off, and then opens the computer in a different network environment. Location awareness enables the Pulse client to intelligently start a new session. Session migration enables Pulse servers to intelligently migrate an existing session.

Location awareness relies on rules you define for each connection. If the conditions specified in the rules are true, Pulse attempts to make the connection. To set up the location awareness rules that select among many connections, you must define location awareness rules for each connection. Each location awareness rule is based on the endpoint's ability to reach an IP address or resolve a DNS name over a specified network interface.

> NOTE: Location awareness behavior is affected by split tunneling configuration. For example, if a location awareness rule relies on a address resolution made on the physical adapter, and split tunneling is disabled, the rule always resolves to FALSE after Pulse establishes the connection.

> NOTE: Connections can be set to manual, automatic, or controlled by location awareness rules. When the user logs in, the Pulse client attempts every connection in its connections list that is set to automatic or controlled by location awareness rules.

To configure location awareness rules:

1. If you have not already done so, create a connection or open an existing connection.

2. In the Connection is established area, select **According to location awareness rules**, and then click **New**.

3. Enter a name for this rule.

4. In the Action list, select one of the following:

   - **DNS server**—Connect if the DNS server associated with the endpoint's network properties is (or is not) set to a certain value or set of values. Specify the DNS server IP address in the IP address box. Also specify a network interface on which the condition must be satisfied:

     - **Physical**—The condition must be satisfied on the physical interfaces on the endpoint.

     - **Pulse Secure**—The condition must be satisfied on the virtual interface that Pulse Secure creates when it establishes a connection.

- **Any**—Use any interface.

- **Resolve address**—Connect if the configured hostname or set of hostnames is (or is not) resolvable by the endpoint to a particular IP address. Specify the hostname in the DNS name box and the IP address or addresses in the IP address box. Also specify a network interface on which the condition must be satisfied.

  > **ⓘ** NOTE: The Pulse client software evaluates IP and DNS policies on network interface changes. DNS lookups occur on DNS configuration changes or when the time-to-live setting (10 minutes) expires for a particular host record. If Pulse cannot resolve the host for any reason, it polls the configured DNS server list every 30 seconds. If the host had been resolved successfully previously and the time-to-live timer has not expired, the polling continues until the timer expires. If the host had not been resolved successfully previously, the resolution attempt fails immediately.

- **Endpoint Address**—Connect if a network adapter on the endpoint has an IP address that falls within or outside of a range or a set of ranges. Specify the IP address or addresses in the IP address box. Also specify a network interface on which the condition must be satisfied.

5. Click **Save Changes**.

After you create the rule or rules, you must enable each rule you want to use for the connection. To enable a negative form of a rule, use a custom version of the rule. To enable location awareness rules:

1. In the list of connection awareness rules for a connection, select the check box next to each rule you want to enable.

2. To specify how to enforce the selected location awareness rules, select one of the following options:

   - **All of the above rules**—The condition is TRUE and the connection is attempted only when all selected location awareness rules are satisfied.

   - **Any of the above rules**—The condition is TRUE and the connection is attempted when any select location awareness rule is satisfied.

   - **Custom**—The condition is TRUE and the connection is attempted only when all selected location awareness rules are satisfied according to the Boolean logic you specify in the Custom box. Use the Boolean condition to specify a negative location rule. For example, connect to the device when Rule–1 is false and Rule–2 is true. The boolean logic in the custom box would be: NOT Rule-1 AND Rule-2. The accepted Boolean operators are AND, OR, NOT, and the use of ( ).

3. Click **Save Changes**.

Related Documentation

- Creating a Client Connection Set

Creating a Client Component Set for Pulse Connect Secure

A Pulse Secure client component set includes specific software components that provide Pulse Secure client connectivity and services.

i NOTE: Client component set options affect Web-based installations only.

To create a Pulse client component set:

1. From the admin console, select **Users > Pulse Secure > Components**.

2. Click **New** to create a new component set.

3. If you have not yet created a client connection set, select **Users > Pulse Secure > Connections** and create a new connection set. Or you can use the default client configuration, which permits dynamic connections, supports the outer username anonymous, and allows the client to automatically connect to Pulse Policy Secure or Pulse Connect Secure.

4. Specify a name for the client component set.

5. (Optional) Enter a description for this client component set.

6. Select a connection set that you have created, or use the default connection set.

7. For Pulse Secure client components, select one of the following options:

   - **All components**—Supports all Pulse connection types.

   - **No components**—Updates existing Pulse client configurations, for example, to add a new connection. Do not use this setting for a new installation.

8. Click **Save Changes**.

9. After you create a component set, distribute the client to users through a role. When users access the role, the installer automatically downloads to the endpoint. The installer components and connections are applied to the endpoint client.

   If client connections associated with the component set for a role are changed even though the list of components has not, the existing configuration on the endpoint is replaced immediately if the endpoint is currently connected, or the next time the endpoint connects.

   If a user is assigned to multiple roles and the roles include different component sets, the first role in an endpoint's list of roles is the one that determines which client (component set) is deployed.

Related Documentation

- *Component Set Options for Pulse Connect Secure*

- *Creating a Client Connection Set for Pulse Connect Secure*

Pulse Secure Client Installation Overview

This section describes how to deploy Pulse for Windows and Pulse for Mac OS X client software from Pulse Policy Secure and Pulse Connect Secure platforms.
Pulse Policy Secure and Pulse Connect Secure include a default connection set and a default component set. These defaults enable you to deploy the Pulse client to users without creating

new connection sets or component sets. The default settings for the client permit dynamic connections, install only the components required for the connection, and permit an automatic connection to Pulse Connect Secure or Pulse Policy Secure to which the endpoint connects.

In all deployment scenarios, you must have already configured authentication settings, realms, and roles.

You can deploy the Pulse Secure Client to endpoints from Pulse Connect Secure and Pulse Policy Secure in the following ways:

- *Web install*—With a Web install (also called a server-based installation), users log in to the Pulse server's Web portal and are assigned to a role that supports a Pulse installation. When a user clicks the link to run Pulse Secure client, the default installation program adds Pulse to the endpoint and adds the default component set and the default connection set. If you do not make any changes to the defaults, the endpoint receives a Pulse installation in which a connection to the Pulse server is set to connect automatically. You can edit the default connection set to add connections of other Pulse servers and change the default options.

  NOTE: A Web install requires that the user have Java installed and enabled for an installation through the Firefox or Chrome browsers or ActiveX enabled for an installation through Internet Explorer. If the browser does not meet this requirement, the user receives a descriptive message at the beginning of the installation process.

  NOTE: A Web install is not compatible with the Pulse rebranding tool, BrandPackager.

- *Preconfigured installer*—Create the connections that an endpoint needs for connectivity and services, download the settings file (.jnprpreconfig), and download default Pulse installation program. For Windows endpoints you run the Pulse installation program by using an msiexec command with the settings file as an option. For OS X endpoints, you run the default installer and then import the .jnprpreconfig file using a separate command.

- *Default installer*—You can download the default Pulse installation program and distribute it to endpoints using your local organization's standard software distribution method (such as Microsoft SMS/SCCM). The Pulse client software is installed with all components and no connections. After users install a default Pulse installation, they can add new connections manually through the Pulse client user interface or by using a browser to access a Pulse server's Web portal. For the latter, the Pulse server's dynamic connection is downloaded automatically and the new connection is added to the Pulse client's connections list when the user starts Pulse by using the Pulse server's Web portal interface. Dynamic connections are created as manual rather than automatic connections, which means that they are run only when the user initiates the connection or the user browses to a Pulse Server and launches Pulse from the server's Web interface.

If the Windows endpoints in your environment do not have admin privileges, you can use the Juniper Installer Service program, which is available on the admin console System Maintenance Installers page. The Juniper Installer Service allows users to download, install, upgrade, and run client applications without administrator privileges. In order to perform tasks that require administrator privileges, the Juniper Installer Service runs under the client's Local System account (a powerful account with full access to the system) and registers itself with

Windows' Service Control Manager (SCM). An Active-X control or a Java applet running inside the user's Web browser communicates the details of the installation processes to be performed through a secure channel between the Pulse server and the client system.

When installing the Juniper Installer Service on client systems, note the following:

- When installing a Pulse Secure client application on a user's Windows system, the Juniper Installer Service deploys two files on the client machine:

  - JuniperSetupSP1Control.ocx

  - AccessServiceComponent.exe

    The device auto-starts this service when installing, and, then stops and removes it when uninstalling.

- Installing the Pulse Secure Installer MSI package requires administrator rights to install onto your client systems. If you plan to use the EXE version, administrator rights is not needed as long as a previous version of the access service component (deployed through, for example, JIS, Pulse, and so forth) is already present. If policies are defined for your client with the group policy "Run only Allowed Windows Application", the following files must be allowed to run in the group policy. If not, client applications might not install.

  - dsmmf.exe

  - JuniperCompMgrInstaller.exe

  - JuniperSetupClient.exe

  - JuniperSetupClientOCX.exe

  - JuniperSetupXP.exe

  - uninstall.exe

  - x86_Microsoft.*.exe

- You should ensure that the Microsoft Windows Installer exists on the client system prior to installing the Juniper Installer Service.

- Your end-users' client systems must contain either a valid and enabled Java Runtime Engine (JRE) or a current Pulse Connect Secure ActiveX control. If the client systems do not contain either of these software components, the users will be unable to connect to the gateway. If there is no JRE on your end-users' client systems, you should download an appropriate installer package from Maintenance > System > Installers. The service appears in the Windows Services (Local) list as Neoteris Setup Service. The service starts automatically on install and during client system start up.

Related Documentation

- *Adding a Pulse Configuration to a New Pulse Installation*

- Installing Pulse Secure Client from the Web

- Installing the Pulse Secure Client on Windows Endpoints Using a Preconfiguration File

- Installing the Pulse Secure Client on OS X Endpoints Using a Preconfiguration File

Installing Pulse Secure Client from the Web

For a Web install, you direct users to the Web interface of the Pulse server. After a successful login, a user is assigned to a role that includes an automatic download and installation of the Pulse client software.

> *i* NOTE: A Web install requires that the user have Java installed and enabled for an installation through the Firefox browser or ActiveX enabled for an installation through Internet Explorer. If the browser does not meet this requirement, the user receives a descriptive message at the beginning of the installation process.

The default Pulse Secure client installation settings includes minimal components, which includes the Host Checker component, and a connection to the Pulse server. If you want a Web install that has customized settings, you can do any of the following:

- Edit the default connection set and add new connections. The default installer uses the default component set which includes the default connection set.

- Create a new connection set and edit the default component set to include the new connection set.

- Edit the role to specify a component set that includes the connections you want for the default installation.

> *i* NOTE: A Pulse installation causes a restart of active network connections on a Windows endpoint. When a user initiates a Pulse installation through a WAN connection to the Web interface of a Pulse server, the user might need to log in to their service provider again to reestablish network connectivity. Users need to be aware of this issue before they begin the installation.

Related Documentation

- Pulse Secure Client Installation Overview
- *Launching Pulse Secure Client from the Pulse Server Web Portal*

Installing the Pulse Secure Client on Windows Endpoints Using a Preconfiguration File

The following procedures apply to Windows installations only.
After you create client connection sets and specify the connections to include within a client component set, you can create a preconfiguration file with all of the connections you want to distribute with the Pulse client. You specify the preconfiguration file as an option when you run the Pulse MSI installer program using an msiexec (**windows\system32\msiexec.exe**) command. To create a preconfigured Pulse installer for distribution to Windows endpoints:

1. Select **Users > Pulse Secure > Connections** and create a connection set with the connections that you want to distribute.

2. Select **Users > Pulse Secure > Components**.

3. If necessary, create a new component set with the connection sets you want to distribute.

It does not matter which component option you select, All components or No components. The Pulse installer installs all components.

4. Select the check boxes next to the component sets that you want to distribute.

5. Click **Download Installer Configuration**.

You are prompted to save the preconfiguration. You can also specify the name of the target Pulse server for the connections, which enables you to create configuration files that are the same except for the target server.

The default filename of the **.jnprpreconfig** file is the name of the selected component set. Make note of the filename and location where you put the file. The preconfiguration file must be available to the clients either through a network share or distributed along with the Pulse Secure client installation file.

6. Select **Maintenance > System > Installers**.

If necessary for your environment, download and install the Juniper Installer Service. To install Pulse, users must have appropriate privileges. The Juniper Installer Service allows you to bypass privilege restrictions and allow users with limited privileges to install Pulse. See "Downloading Client Installer Files" for more information about Juniper Installer Service.

7. Download the appropriate Pulse Secure client installer for your Windows environment:

- Pulse Secure client Installer (32-bit)

- Pulse Secure client Installer (64-bit)

> **NOTE:** For a Windows installation (.msi) that uses an automated distribution mechanism and where the users do not have administrator privileges, you should ensure that the installation is run in the proper context, typically the USER context. To install in USER context, first advertise the .msi while in the SYSTEM context. For example, to advertise the 64-bit Windows installation to all users, use the following msiexec command:
> **msiexec /jm \Pulsesecure.x64.msi**
> The advertisement allows the installation to be run in USER context even if the user is a restricted (non-admin) user. The location where the advertisement is run and where the actual installation is run *must* be the same. If the installation is an upgrade, you must advertise the upgrade version before running it. (Note that it is much easier to upgrade the Pulse client by not disabling the automatic upgrade feature on the Pulse server.) After the installation is run by the user, the Pulse client will use the correct user certificate and context.

Installing the Pulse Client Using Advanced Command-Line Options

The Pulse Secure client installer includes the Pulse client and all the software components for all the Pulse services. The preconfiguration file contains the definitions of the Pulse connections that provide client access to specific Pulse servers and services.

Usage Notes:

- The preconfigured installer installs all Pulse components unless you specify the specific components you want using ADDLOCAL options. If you use one or more ADDLOCAL

options, the preconfigured installer installs only the components specified by ADDLOCAL. A preconfigured installer ignores the component set you select when you create the preconfiguration file.

- When you run **msiexec**, you should append **/qn** or **/qb** (**msiexec** options) to the command line to suppress the installation program user interface. For example, the user interface lets the user choose a complete installation or a custom installation, which can override the components you specify with ADDLOCAL options. If the user selects Complete, the **msiexec** program ignores the ADDLOCAL options in the command line and installs all components. The **/qn** option specifies a silent install, so no user interface appears. The **/qb** option also hides the user interface but it displays a progress bar.

- The procedures in this topic are valid with Windows installations only. For information about installing Pulse on OS X endpoints, see "Installing the Pulse Secure Client on OS X Endpoints Using a Preconfiguration File".

You run the Pulse preconfigured installer program with **msiexec** (the command line for launching .msi programs on Windows platforms) and specify the following options.

> *i*   NOTE: Command line options (CONFIGFILE and ADDLOCAL) are case sensitive and must be all caps.

> *i*   NOTE: If the path to the **.jnprpreconfig** file includes spaces, be sure to use quotes around the path.

- **CONFIGFILE**—This property specifies a configuration file to be imported into Pulse during installation. The property must include the full path to the configuration file. For example:

  msiexec /i Pulsesecure.msi CONFIGFILE="c:\temp\my configuration.jnprpreconfig"

- **ADDLOCAL**—This optional property specifies which features and feature options (subfeatures) to install when you want to install optional Pulse components. If you do not specify ADDLOCAL options, all Pulse components are installed. When you use ADDLOCAL, you should append **msiexec** options **/qn** or **/qb** to the command line to suppress the installation program user interface.

  Feature and subfeature names are case sensitive. To specify multiple features in a single command, separate each feature with a comma.

  The ADDLOCAL property has the following optional subfeatures:

  - **SAEndpointDefense**—Available with Connect Secure.

  - **UACEndpointDefense**—Available with Policy Secure.

    > *i*   NOTE:

    The Enhanced Endpoint Security product has been discontinued.

Examples

When you use ADDLOCAL, you should append **msiexec** options **/qn** or **/qb** to the command line to suppress the installation program user interface. These examples use **/qb**.

To install Pulse on a Windows 32-bit endpoint using a configuration file:

    msiexec /i Pulsesecure.x86.msi CONFIGFILE=c:\pulse\Pulse-Connection-no.jnprpreconfig
       ADDLOCAL=UACEndpointDefense /qb

To install Pulse on a 32-bit Windows endpoint using a configuration file:

    msiexec /i Pulsesecure.x86.msi CONFIGFILE=c:\temp\myconfiguration.jnprpreconfig /qb

To install Pulse on a 64-bit Windows endpoint using a configuration file:

    msiexec /i Pulsesecure.x64.msi CONFIGFILE=c:\temp\myconfiguration.jnprpreconfig
       ADDLOCAL=SAEndpointDefense /qb

To install Pulse on a 64-bit Windows endpoint using a configuration file:

    msiexec /i Pulsesecure.x64.msi CONFIGFILE=c:\temp\myconfiguration.jnprpreconfig /qb

Repairing a Pulse Installation on a Windows Endpoint

Pulse Secure client uses an MSI installer, which supports a repair function. If problems with Pulse on a Windows endpoint indicate missing or damaged files or registry settings, the user can easily run the installation repair program. The repair program performs a reinstallation and replaces any missing files. The repair program does not install any files that were not part of the original installation. For example, if the file that holds Pulse connection configurations is damaged, the file installed by the repair program does not replace any Pulse connections that were created by the user or deployed to the endpoint after the original Pulse installation.

To repair a Pulse installation on a Windows endpoint:

1. On the Windows endpoint where Pulse is installed, click **Start > Programs > Pulse Secure > Repair Pulse Secure**.

2. Follow the prompts for the installation wizard.

   When the program is finished, you might be prompted to reboot the system.

Related Documentation

- Installing the Pulse Secure Client on OS X Endpoints Using a Preconfiguration File

Installing the Pulse Secure Client on OS X Endpoints Using a Preconfiguration File

The following procedures apply to OS X installations only.
After you create client connection sets and specify the connections to include within a client component set, you can create a preconfiguration file with all the connections you want to distribute with the Pulse client. After you run the Pulse installer on the endpoint, you run a special command that imports the settings from the preconfiguration file into Pulse.
To create a preconfigured Pulse installer for distribution to OS X endpoints:

1. Select **Users > Pulse Secure > Connections** and create a connection set with the connections that you want to distribute.

2. Select **Users > Pulse Secure > Components**.

3. If necessary, create a new component set with the connection sets you want to distribute.

The All components or No components options apply to Web-based installations only. The Pulse installation program for OS X always installs all components.

4. Select the check boxes next to the component sets that you want to distribute.

5. Click **Download Installer Configuration**.

   You are prompted to save the preconfiguration. You can also specify the name of the target Pulse server for the connections, which enables you to quickly create multiple configuration files that are the same except for the target server.

   The default filename of the **.jnprpreconfig** file is the name of the selected component set. Make note of the filename and location where you put the file. The preconfiguration file must be available to the clients either through a network share or distributed along with the Pulse Secure client installer file.

6. Select **Maintenance > System > Installers**.

7. Download the Pulse Secure installer, **Pulse Secure Installer (Macintosh)**.

Installing the Pulse Client on OS X Endpoints Using Command-Line Options

The Pulse Secure installer includes the Pulse client and all of the software components for all of the Pulse services. The preconfiguration (**.jnprpreconfig**) file contains the definitions of the Pulse connections that provide client access to specific Pulse servers and services. After you distribute the Pulse installation package, you must first run the installer, and then run a separate program called jamCommand, which imports the settings from the **.jnprpreconfig** file. The jamCommand program is part of the Pulse installation.

The Pulse file you download from the Pulse server is in compressed (.dmg) format. You must unpack the file before you run the Pulse installation program.

The following steps include sample commands to install Pulse on an OS X endpoint and then import Pulse connections from a **.jnprpreconfig** file.

1. Run the Pulse installation program:

   sudo /usr/sbin/installer -pkg <full-path-to-the-pulse-install-package> -target /

2. Import the settings from the **.jnprpreconfig** file:

   /Applications/Pulse Secure\ Pulse.app/Contents/Plugins/JamUI/jamCommand –ImportFIle <full-path-to-the-jnprpreconfig-file>

Related Documentation

- Installing the Pulse Secure Client on Windows Endpoints Using a Preconfiguration File

- *jamCommand Reference*

Pulse Secure Command-line Launcher

The Pulse Launcher (pulselauncher.exe) is a standalone client-side command-line program that allows you to launch Pulse and connect to or disconnect from a Pulse server (Pulse Connect Secure or Pulse Policy Secure) without displaying the Pulse graphical user interface.

Pulse Launcher Usage Notes:

- Pulse Launcher runs on Windows 32-bit and 64-bit endpoints.

- The Pulse Launcher program, pulselauncher.exe, is installed as part of a Pulse client installation in **Program Files\Common Files\Juniper Networks\Integration** or **Program Files (x86)\Common Files\Juniper Networks\Integration**.

- Pulse Launcher works only for the Connect Secure or Policy Secure (L3) connection type. Pulse Launcher does not support the SRX Series or Policy Secure (802.1X) connection types.

- The Pulse Launcher program does not support the role mapping option that prompts a user to select from a list of assigned roles. If you use the Pulse Launcher and more than one role can be assigned to a user, you must configure the role mapping settings for the realm to merge settings for all assigned roles. If the realm settings require the user to select a role, the Pulse Launcher command fails and exits with return code 2.

- Pulse Launcher does not support secondary authentication.

To use the Pulse Launcher program:

1. Write a script, batch file, or application.

2. Include a call to the Pulse Launcher executable, pulselauncher.exe.

3. Include logic in your script, batch file, or application to handle the possible return codes.

    Table 3 lists the Pulse Launcher arguments.

The following command shows the complete pulselauncher.exe command syntax:

```
pulselauncher.exe [-version|-help|-stop] [-url <url> -u <user> -p <password> -r
<realm>] [-d <DSID> -url <url>] [-cert <client certificate> -url <url> -r
<realm>] [-signout|-suspend|-resume -url <url>] [-t timeout)]]
```

Table3: Pulse Launcher Arguments

Argument	Action
-version	Display the Pulse Launcher version information, then exit.
-help	Display available arguments information.
-stop	Stop Pulse and disconnect all active connections.
-url <url>	Specify the Pulse server URL.
-u <user>	Specify the username.
-p <password>	Specify the password for authentication.
-r <realm>	Specify the realm on the Pulse server.
-d <DSID>	Passes a cookie to Pulse Launcher for a specified Pulse server from another authentication mechanism when Pulse Launcher starts. When you use the -d argument, you must also specify the -url argument to specify the Pulse server.

-cert <client certificate>	Specify the certificate to use for user authentication. For <client certificate> use the string specified in the Issued To field of the certificate. When using the -cert argument, you must also specify the -url and -r <realm> arguments.

To use certificate authentication with the Pulse Launcher program, you must first configure the Pulse server to allow the user to sign in via user certificate authentication. You must also configure a trusted client CA on the Pulse server and install the corresponding client-side certificate in the Web browsers of end-users before running the Pulse Launcher.

If the certificate is invalid, the Pulse Launcher displays an error message on the command line and logs a message in the log file.

NOTE: If Pulse is launched through a browser, the browser handles certificate verification. If Pulse is launched through an application on Windows, the application handles certificate verification. If Pulse is launched through the Pulse Launcher on Windows, Pulse Launcher handles the expired or revoked client certificates. |
| -signout <url>

-suspend <url>

-resume <url> | Signout disconnects and signs out from a specific server. Suspend puts an active connection in the suspend state without removing the session information from the server. Resume restores a suspended connection. Pulse can have multiple simultaneous connections so the -url argument is required when you use -signout, -suspend, or -resume. |
| -t <timeout in seconds> | The amount of time allowed for the connection to take place before the attempt fails. Min = 45 (default), Max = 600. |

Table 4 lists the possible return codes pulselauncher.exe returns when it exits.

Table4: Pulse Launcher Return Codes

Code	Description
-1	Pulse is not running.
0	Success.
1	A parameter is invalid.
2	Connection has failed or Pulse is unable to connect to the specified gateway.
3	Connection established with error.
4	Connection does not exist. Example: the command attempts to sign out from a server that has not been added on the Pulse UI.
5	User cancelled connection.
6	Client certificate error.
7	Timeout error.
8	No user connection allowed from Pulse UI.
9	No policy override from Pulse UI.
25	Invalid action for current connection state. This error code would occur if you tried to suspend or resume a connection that was disconnected.
100	General error.

Examples

The following command is a simple login application that captures the credentials the user enters, and passes the credentials as arguments to pulselauncher.exe:

pulselauncher.exe -u JDoe -p my$Pass84 -url https://int-company.portal.com/usr -r Users
pulselauncher return code: 0

The following Pulse Launcher example shows a certificate authentication:

pulselauncher.exe -url https://int-company.portal.com/usr -cert MyCert -url https://int-company.portal.com/usr -r Users
pulselauncher return code: 0

The following example shows a command to use Pulse Launcher to specify a cookie (-d) for a specific Pulse server (-url):

pulselauncher.exe –d 12adf234nasu234 -url https://int-company.portal.com/usr
pulselauncher return code: 0

Related Documentation

- Using jamCommand to Import Pulse Secure Connections

- *jamCommand Reference*

Using jamCommand to Import Pulse Secure Connections

The jamCommand.exe program is a command line program that imports a .jnprpreconfig file into the Pulse client. The jamCommand program is available for Windows (Vista, Windows 7, and Windows 8) and Mac OS X.

> NOTE: What is "jam?" The working name for Pulse Secure client during the early development phase was Juniper Access Manager. Many Pulse filenames, directories, and other Pulse Secure client elements include the acronym "jam."

A .jnprpreconfig file includes Pulse connection parameters. You can create a .jnprpreconfig file on the Pulse server, and then use it as part of a Pulse installation to ensure that Pulse users have one or more Pulse connections when they start Pulse for the first time.

> NOTE: One typical use case for jamCommand on a Windows endpoint is to first run jamCommand to import one or more Pulse connections from a .jnprpreconfig file, and then run pulselauncher.exe to start Pulse.

To install Pulse connections using jamCommand:

1. Create a .jnprpreconfig file on the Pulse server.

   In the Pulse server admin console, click **Users > Pulse Secure > Components**.

2. Select the component sets you want, and then click **Download Installer Configuration**.

3. Distribute the .jnprpreconfig file to the Pulse endpoints.

4. Run jamCommand with the .jnprpreconfig file as an option. For example:

   On Windows:

   \Program Files\Common Files\Juniper Networks\JamUI\jamCommand -importfile myfile.jnprpreconfig

   On Mac OSX:

/Applications/Pulse Secure/Contents/Plugins/JamUI/jamCommand -ImportFile myfile.jnprpreconfig

If the Pulse client is running when you run jamCommand, the new Pulse connection or connections appear immediately. The connection name appears as it was defined when you created the connection in the Pulse server admin console.

Related Documentation

- Pulse Secure Command-line Launcher

- Installing the Pulse Secure Client on Windows Endpoints Using a Preconfiguration File

- Installing the Pulse Secure Client on OS X Endpoints Using a Preconfiguration File

- *jamCommand Reference*

PART 2 Access Management Framework

CHAPTER 3 General Access Management

Access Management Overview

The system enables you to secure your company resources using authentication realms, user roles, and resource policies. These three levels of accessibility allow you to control access from a very broad level (controlling who may sign into the device) down to a very granular level (controlling which authenticated users may access a particular URL or file). You can specify security requirements that users must meet to sign in to the device, to gain access to features, and even to access specific URLs, files, and other server resources. The system enforces the policies, rules and restrictions, and conditions that you configure to prevent users from connecting to or downloading unauthorized resources and content.

To permit endpoints that are not directly connected to a Juniper Networks security device to access resources behind the firewall, you can configure a Policy Secure device to provision shared user sessions from any number of different Connect Secure devices and Infranet Controllers. IF-MAP Federation allows users to access resources protected by any number of Juniper Networks firewalls (Infranet Enforcers) without requiring additional authentication.

The access management framework is available on all Connect Secure products. The access management features, including realms, roles, resource policies, and servers, are the base of the platform on which all Connect Secure products are built.

Related Documentation

- User Roles Overview

- Understanding Authentication Realms

Policies, Rules & Restrictions, and Conditions Overview

The system enables you to secure your company resources using authentication realms, user roles, and resource policies. These three levels of accessibility allow you to control access from a very broad level (controlling who may sign into the device) down to a very granular level (controlling which authenticated users may access a particular URL or file).

Accessing Authentication Realms

Resource accessibility begins with the authentication realm. An authentication realm is a grouping of authentication resources, including:

- **An authentication server**—verifies that the user is who he claims to be. The system forwards credentials that a user submits on a sign-in page to an authentication server.

- **An authentication policy**—specifies realm security requirements that need to be met before the system submits a user's credentials to an authentication server for verification.

- **A directory server**—specifies an LDAP server that provides user and group information to the system that it uses to map users to one or more user roles.

- **Role mapping rules**—specifies the conditions a user must meet for the system to map the user to one or more user roles. These conditions are based on either user information returned by the realm's directory server or the user's username.

You can associate one or more authentication realms with the sign-in page. When more than one realm exists for a sign-in page, a user must specify a realm before submitting her credentials. When the user submits her credentials, the system checks the authentication policy defined for the chosen realm. The user must meet the security requirements you define for a realm's authentication policy or else the system does not forward the user's credentials to the authentication server.

At the realm level, you can specify security requirements based on various elements such as the user's source IP address or the possession of a client-side certificate. If the user meets the requirements specified by the realm's authentication policy, the system forwards the user's credentials to the appropriate authentication server. If this server successfully authenticates the user, then the system evaluates the role mapping rules defined for the realm to determine which roles to assign to the user.

Accessing User Roles

A role is a defined entity that specifies session properties for users who are mapped to the role. These session properties include information such as session time-outs and enabled access features. A role's configuration serves as the second level of resource access control. Not only does a role specify the access mechanisms available to a user, but you can also specify restrictions with which users must comply before they are mapped to a role.

At the role level, you can specify security requirements based on elements such as the user's source IP address and possession of a client-side certificate. If the user meets the requirements specified either by a role mapping rule or a role's restrictions, then the system maps the user to the role. When a user makes a request to the backend resources available to the role, the system evaluates the corresponding access feature resource policies.

Note that you may specify security requirements for a role in two places—in the role mapping rules of an authentication realm (using custom expressions) or by defining restrictions in the role definition. The system evaluates the requirements specified in both areas to make sure the user complies before it maps the user to a role.

Accessing Resource Policies

A resource policy is a set of resource names (such as URLs, hostnames, and IP address/netmask combinations) to which you grant or deny access or other resource-specific actions, such as rewriting and caching. A resource policy serves as the third level of resource access control. While a role may grant access to certain types of access features and resources (such as bookmarks and applications), whether or not a user can access a specific resource is controlled by resource policies. These policies may even specify conditions that, if met, either deny or grant user access to a server share or file. These conditions may be based on security requirements that you specify. The user must meet these security requirements or else the system does not process the user's request.

At the resource level, you can specify security requirements based elements such as the user's source IP address or possession of a client-side certificate. If the user meets the requirements specified by a resource policy's conditions, then the system either denies or grants access to the requested resource. You may enable Web access at the role level, for example, and a user mapped to the role may make a Web request. You may also configure a Web resource policy to deny requests to a particular URL or path when Host Checker finds an unacceptable file on the user's machine. In this scenario, the system checks to see if Host Checker is running and indicates that the user's machine complies with the required Host Checker policy. If the user's machine complies, meaning the unacceptable file is not found, then the system grants the user access to the requested Web resource.

Note that you can create separate resource policies or you can create automatic resource policies (called autopolicies) during resource profile configuration (recommended).

Related Documentation

- Policies, Rules & Restrictions, and Conditions Evaluation

- Dynamic Policy Evaluation

Policies, Rules & Restrictions, and Conditions Evaluation

Figure 13 illustrates the access management security checks that the system performs when a user tries to access resources through the device. A detailed description of each step follows after the diagram.

Figure13: Security Checks Performed During a User Session
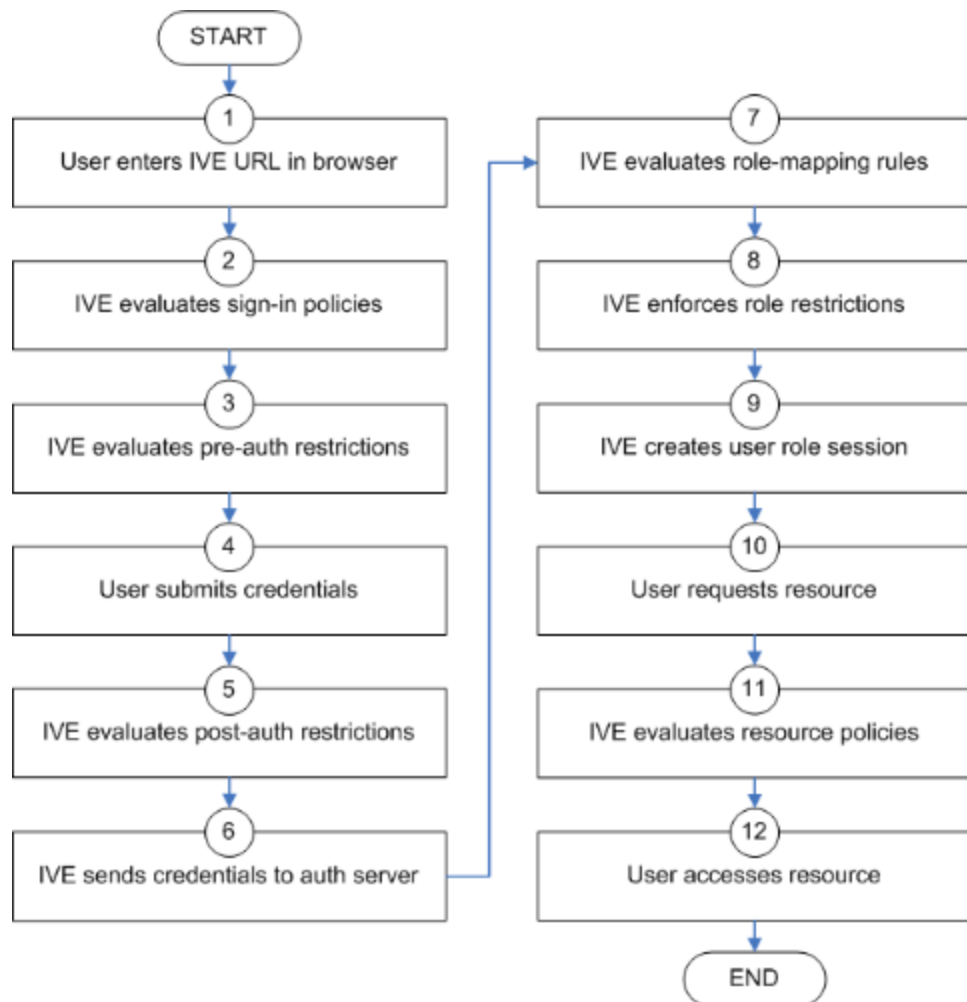
1. The user enters the URL of the device end-user console (such as http://employees.yourcompany.com/marketing) in a Web browser.

2. The system evaluates its sign-in policies (starting with the administrator URLs and continuing to user URLs) until it matches the hostname entered by the user.

3. The system evaluates pre-authentication restrictions and determines if the user's system passes host checks and other requirements. If the pre-authentication checks fail, the system denies the user access. If the checks pass, the system prompts the user to enter the username and password for the realms whose preauthentication checks succeeded. (If required by the realm, the system prompts the user to enter two sets of credentials.) If more than one realm exists, the user must enter a realm or choose one from a list.

4. The system evaluates the post-authentication restrictions and determines whether the user's password conforms to specified limits and requirements. If the postauthentication checks fail, the system denies the user access. If the checks pass, the system passes the user's credentials to the realm's authentication server.

5. The system forwards the user's username and password to the authentication server, which returns success or failure. (A RADIUS or SiteMinder authentication server also returns attributes for the system to use in role mapping.) If the authentication server returns failure, the system denies the user access. If the server returns success, the system stores the user's credentials. If the realm has a separate LDAP authorization server, the system also queries the LDAP server for attribute and group information and saves the information returned by LDAP. If the realm includes a secondary authentication server, the system repeats this process with the secondary server.

6. The system evaluates the realm's role mapping rules and determines the roles for which the user is eligible. The system determines eligibility using information from the LDAP or RADIUS server or the user's username.

7. The system evaluates the restrictions of the eligible roles, enabling the user to access those roles whose restrictions the user's computer meets. Restrictions may include source IP, browser type, client-side certificate, Host Checker, and Cache Cleaner.

8. The system creates a "session role," determining the user's session permissions. If you enable permissive merging, the system determines session permissions by merging all valid roles and granting the allowed resources from each valid role. If you disable merging, the system assigns the user to the first role to which he is mapped.

9. When the user requests a resource, the system checks whether the corresponding access feature is enabled for the session user role. If not, the system denies the user access. If the access feature is enabled, the evaluates resource policies.

10. The system evaluates resource profiles and policies related to the user's request, sequentially processing each until it finds the profile or policy whose resource list and designated roles match the user's request. The system denies user access to the resource if specified by the profile or policy. Otherwise, the system intermediates the user request if the profile or policy enables access.

11. The system intermediates the user request, forwarding the user's request and credentials (if necessary) to the appropriate server. Then, the system forwards the server's response to the user.

12. The user accesses the requested resource or application server. The user session ends when the user signs out or his session times out due to time limits or inactivity. The system may also force the user out if the session if you enable dynamic policy evaluation and the user fails a policy.

> NOTE: If you enable dynamic policy evaluation, the system performs additional checks beyond the ones mentioned here.

Related Documentation

- Dynamic Policy Evaluation

Dynamic Policy Evaluation

Dynamic policy evaluation allows you to automatically or manually refresh the assigned roles of users by evaluating a realm's authentication policy, role mappings, role restrictions, and resource policies. When the system performs a dynamic evaluation, it checks whether the client's status has changed. (For instance, the client's Host Checker status may have changed. Or, if the user is roaming, the computer's IP address may have changed.) If the status has changed, the system enables or denies the user access to the dependent realms, roles, or resource policies accordingly.

The system does not check for changes in user attributes from a RADIUS, LDAP, or SiteMinder server when performing dynamic policy evaluation. Instead, the system re-evaluates rules and policies based on the original user attributes that it obtained when the user signed into the device.

Understanding Dynamic Policy Evaluation

Please note the following about Dynamic Policy Evaluation:

- Clients that use Network Communications Protocol (NCP) do not honor policy changes. This includes NC, WSAM, Pulse Collaboration, and WTS/CTS.

- NC has a persistent NCP control channel. NC does not get this new policy until NC is disconnects and attempts to reconnect by the user. When NC reconnects via the UI level (NC service reconnect times out), it will obtain the new policy.

- WSAM establishes a new NCP tunnel when the protected application opens a new connection, so WSAM establishes new NCP connections frequently. This means WSAM gets the new policy frequently.

- Pulse Collaboration has a persistent NCP data channel, so Pulse Collaboration will not get the new policy. The down side of Pulse Collaboration not getting the new policy is not significant because Pulse Collaboration only tunnels its own data traffic.

- WTS has a persistent NCP tunnel so it does not get policy changes until the user disconnects and then reconnects.

  NOTE: Because the system evaluates Web and Files resource policies whenever the user makes a request for a resource, dynamic policy evaluation is unnecessary for Web and Files. The system does not use dynamic policy evaluation for Meetings resource policies and Email Client resource policies.

If the system determines after a dynamic policy evaluation that a user no longer meets the security requirements of a policy or role, the system terminates the connection immediately with the user. The user may see the closing of a TCP or application connection, or the termination of a user session for VPN Tunneling, Secure Application Manager, Terminal or Telnet/SSH. The user must take the necessary steps to meet the security requirements of the policy or role, and then sign into the system again.

The system logs information about policy evaluation and changes in roles or access in the Event log.

Understanding Standard Policy Evaluation

If you do not use dynamic policy evaluation, the system evaluates policies and roles only when the following events occur:

- When the user first tries to access the system sign-in page, the system evaluates the Host Checker policies (if any) for a realm.

- Immediately after the user's initial authentication, the system evaluates the user's realm restrictions in the authentication policy, role mapping rules, and role restrictions.

- When the user makes a request for a resource, the system evaluates resource access policies to determine if the associated role is allowed to access the resource.

- When the Host Checker status of the user's machine changes, the system evaluates the Host Checker policies (if any) for the role.

If you do not use dynamic policy evaluation and you make changes to an authentication policy, role mapping rules, role restrictions, or resource policies, the system enforces those changes only when the events described above occur.

If you use dynamic policy evaluation, the system enforces changes when the events described above occur, and it also enforces changes at the times you specify.

Enabling Dynamic Policy Evaluation

You can use dynamic policy evaluation in the following ways:

- **Evaluate all signed-in users in a realm**—You can automatically or manually refresh the roles of all currently signed-in users of a realm by using the General tab of the Administrators > Admin Realms > Select Realm or Users > User Realms > Select Realm page. You can trigger the system to perform a dynamic policy evaluation at the realm level based on:

  - **An automatic timer**—You can specify a refresh interval that determines how often the system performs an automatic policy evaluation of all currently signed-in realm users, such as every 30 minutes. When using the refresh interval, you can also fine-tune the system performance by specifying whether or not you want to refresh roles and resource policies as well as the authentication policy, role mapping rules, and role restrictions.

  - **On-demand**—At any time, you can manually evaluate the authentication policy, role mapping rules, role restrictions, and resource policies of all currently signed-in realm users. This technique is especially useful if you make changes to an authentication policy, role mapping rules, role restrictions, or resource policies and you want to immediately refresh the roles of a realm's users.

- **Evaluate all signed-in users in all realms**—At any time, you can manually refresh the roles of all currently signed-in users in all realms by using settings in the System > Status >Active Users page.

- **Evaluate individual users**—You can automatically refresh the roles of individual users by enabling dynamic policy evaluation for Host Checker on the Authentication > Endpoint Security > Host Checker page. Host Checker can trigger the system to evaluate resource policies whenever a user's Host Checker status changes. (If you do not enable dynamic policy evaluation for Host Checker, the system does not evaluate resource policies but it

does evaluate the authentication policy, role mapping rules, and role restrictions whenever a user's Host Checker status changes.)

Related Documentation

- Displaying Active Users

Specifying Source IP Access Restrictions

This topic describes options to enforce source IP restrictions for access to the corporate network or intranet resources. It includes the following information:

- About Source IP Restrictions
- Specifying Source IP Restrictions at the Realm Level
- Specifying Source IP Restrictions at the Role Level
- Specifying Source IP Restrictions in Resource Policies

About Source IP Restrictions

You can enforce access rules based on the source IP address of the request. You can configure rules related to sign-in, role-mapping, and resource access.

At the realm level, you can add source IP rules to the realms associated with sign-in pages. The user must sign in from a host with an IP address that is allowed by the source IP requirements for the authentication realm. If the source IP policy does not allow the host to access the realm, the system does not forward the user's credentials to the authentication server, and the user is denied access. You can set up multiple rules. For example, you can deny access to all users on a wireless network (10.64.4.100), and allow access to all other network users (0.0.0.0).

At the user role level, you can add source IP rules to the criteria that determine user role membership. If the source IP rule disqualifies a user from a role, subsequent role mapping rules are consulted.

In resource policies, you can add allow/deny rules based on source IP.

Specifying Source IP Restrictions at the Realm Level

To specify source IP restrictions:

1. Navigate to the administrator or user realm you want to configure:

   - **Administrators > Admin Realms >** *Realm*

   - **Users > User Realms >** *Realm*

2. Select **Authentication Policy > Source IP** to display the Source IP policy configuration page.

3. Choose one of the following options:

   - **Allow users to sign in from any IP address**—Essentially, this option turns off source IP restrictions.

   - **Allow or deny users from the following IP addresses**—Specifies source IP restrictions. If you select this option, use the policy table controls to create source IP rules.

4. Add a rule to the table:

   1. Use the text boxes to specify source IP address match criteria:

      - For IPv4 clients, enter IPv4 address and netmask pairs.

      - For IPv6 clients, enter IPv6 address and prefix length pairs.

   2. Use the Allow and Deny option buttons to specify the action when a rule matches.

   3. Click **Add** to add the rule to the table.

   4. Use the selection box and arrow buttons to order the list. Move the highest priority restrictions to the top of the list. For example, to deny access to all users on a wireless network (10.64.4.100) and allow access to all other network users (0.0.0.0), move the wireless network address (10.64.4.100) to the top of the list and move the (0.0.0.0) network below the wireless network.

5. For administrator realms, select the ports where the administrator can log in (internal, external, and management). On virtual appliances that use traffic segregation, administrators can log in on the management port on the Default Network or Administrative Network (see "Using the Traffic Segregation Feature"). If necessary, click **External Port** or **Management Port** to enable the port.

6. Save the configuration.

Specifying Source IP Restrictions at the Role Level

To specify source IP restrictions:

1. Navigate to the administrator or user role you want to configure:

   - **Administrators > Admin Roles >** *Role*

   - **Users > User Roles >** *Role*

2. Select **General > Restrictions > Source IP** to display the Source IP policy configuration page.

3. Choose one of the following options:

   - **Allow users to sign in from any IP address**—Essentially, this option turns off source IP restrictions.

   - **Allow or deny users from the following IP addresses**—Specifies source IP restrictions. If you select this option, use the policy table controls to create source IP rules.

4. Add a rule to the table:

   1. Use the text boxes to specify source IP address match criteria:

      - For IPv4 clients, enter IPv4 address and netmask pairs.

      - For IPv6 clients, enter IPv6 address and prefix length pairs.

   2. Use the Allow and Deny option buttons to specify the action when a rule matches.

   3. Click **Add** to add the rule to the table.

4. Use the selection box and arrow buttons to order the list. Move the highest priority restrictions to the top of the list. For example, to deny access to all users on a wireless network (10.64.4.100) and allow access to all other network users (0.0.0.0), move the wireless network address (10.64.4.100) to the top of the list and move the (0.0.0.0) network below the wireless network.

5. Save the configuration.

Specifying Source IP Restrictions in Resource Policies

A third way to use source IP restrictions is by creating custom rules in resource policies. The action for custom rules is either allow or deny. The match criteria include resources and conditions. One of the conditions you can set is source IP, so you can enforce source IP restrictions through resource policies. For example:

1. Navigate to **Users > Resource Policies**.

2. Select a policy. Click **Web Access Policies**, for example, to display its policies list.

3. Click the **Initial Policy for Local Resources** policy to edit it.

4. Click the **Detailed Rules** tab.

5. Under Conditions, expand the Prebuilt Conditions list, expand the **SourceIPStr** selections, select one of the example expressions, such as **SourceIPStr = "192.168.10.0/24"** or **SourceIPStr = "2001:DB8::15"**, and click **Insert Expression** to add the string to the Conditions box.

6. Modify the IP address. In other words, replace **192.168.10.0/24** with an IPv4 address / netmask pair; replace **2001:DB8::15** with an IPv6 address.

7. Specify the other match condition (resource) and specify the action (allow or deny).

8. Save the configuration.

Related Documentation

- Role Restrictions

- Creating an Authentication Realm

- Specifying Browser Access Restrictions

- Specifying Certificate Access Restrictions

- Specifying Password Access Restrictions

- Configuring Host Checker Restrictions

Specifying Browser Access Restrictions

Use a browser restriction to control from which Web browsers users can access a system sign-in page or be mapped to a role. If a user tries to sign in to the device using an unsupported browser, the sign-in attempt fails and a message displays stating that an unsupported browser is being used. This feature also enables you to ensure that users sign in to the device from browsers that are compatible with corporate applications or are approved by corporate security policies.

You can restrict system and resource access by browser-type:

- **When administrators or users try to sign in to Connect Secure**—The user must sign in from a browser whose user-agent string meets the specified user-agent string pattern requirements for the selected authentication realm. If the realm "allows" the browser's user-agent string, then the system submits the user's credentials to the authentication server. If the realm "denies" the browser's user-agent string, then the system does not submit the user's credentials to the authentication server.

- **When administrators or users are mapped to a role**—The authenticated user must be signed in from a browser whose user-agent string meets the specified user-agent string pattern requirements for each role to which the system may map the user. If the user-agent string does not meet the "allowed" or "denied" requirements for a role, then the system does not map the user to that role.

- **When users request a resource**—The authenticated, authorized user must make a resource request from a browser whose user-agent string meets the specified "allowed" or "denied" requirements for the resource policy corresponding to the user's request. If the user-agent string does not meet the "allowed" or "denied" requirements for a resource, then the system does not allow the user to access the resource.

The browser restrictions feature is not intended as a strict access control since browser user-agent strings can be changed by a technical user. It serves as an advisory access control for normal usage scenarios.

To specify browser restrictions:

1. Select the level at which you want to implement browser restrictions:

    - **Realm level**—Navigate to:

        - **Administrators > Admin Realms >** *Select Realm* **> Authentication Policy > Browser**

        - **Users > User Realms >** *Select Realm* **> Authentication Policy > Browser**

    - **Role level**—Navigate to:

        - **Administrators > Admin Realms >** *Select Realm* > *Role Mapping* **> Select|Create Rule > Custom Expressions**

        - **Administrators > Admin Roles >** *Select Role* **> General > Restrictions > Browser**

        - **Users > User Realms >** *Select Realm* **> Role Mapping > Select|Create Rule > Custom Expression**

        - **Users > User Roles >** *Select Role* **> General > Restrictions > Browser**

2. Choose one of the following options:

    - **Allow all users matching any user-agent string sent by the browser**— Allows users to access the system or resources using any of the supported Web browsers.

    - **Only allow users matching the following User-agent policy**—Allows you to define browser access control rules. To create a rule:

    1. For the User-agent string pattern, enter a string in the format

        browser_string

where start (*) is an optional character used to match any character and *browser_string* is a case-sensitive pattern that must match a substring in the user-agent header sent by the browser. Note that you cannot include escape characters (\) in browser restrictions.

For example, the following is a browser sent user-agent header:

```
Mozilla/5.0 (Windows NT 6.1; WOW64) AppleWebKit/537.22 (KHTML, like Gecko)
```

where:

- *Mozilla/5.0* indicates compatibility with the Mozilla rendering engine.

- (*Windows NT 6.1; WOW64*) are the details of the system in which the browser is running.

- *AppleWebKit/537.22* is the platform the browser users.

- (*KHTML, like Gecko*) is the browser platform details.

Using the above example, enter *Windows NT* as a string pattern for specifying the Windows NT system. For more details on user-agent strings, see your specific browser's documentation.

2. Select either:

- **Allow** to allow users to use a browser that has a user-agent header containing the <browser_string> substring.

- **Deny** to prevent users from using a browser that has a user-agent header containing the <browser_string> substring.

3. iii. Click **Add**.

4. Click **Save Changes** to save your settings.

Rules are applied in order, so the first matched rule applies.

Literal characters in rules are case sensitive, and spaces are allowed as literal characters.

For example, the string *Netscape* matches any user-agent string that contains the substring Netscape.

The following rule set grants resource access only when users are signed in using Internet Explorer 5.5x or Internet Explorer 6.x. This example takes into account some major non-IE browsers that send the 'MSIE' substring in their user-agent headers:

*Opera*Deny

*AOL*Deny

*MSIE 5.5*Allow

*MSIE 6.*Allow

* Deny

Specifying Certificate Access Restrictions

When you install a client-side certificate on the device through the System > Configuration > Certificates > Trusted Client CAs page of the admin console, you can restrict system and resource access by requiring client-side certificates:

- **When administrators or users try to sign in to Connect Secure**—The user must sign in from a machine that possesses the specified client-side certificate (from the proper certificate authority (CA) and possessing any optionally specified field/value pair requirements). If the user's machine does not possess the certificate information required by the realm, the user can access the sign-in page, but once the system determines that the user's browser does not possess the certificate, the system does not submit the user's credentials to the authentication server and the user cannot access features on the device.

  To implement certificate restrictions at the realm level, navigate to:

  - **Administrators > Admin Realms >** *SelectRealm* **> Authentication Policy > Certificate**

  - **Users > User Realms >** *SelectRealm* **> Authentication Policy > Certificate**

- **When administrators or users are mapped to a role**—The authenticated user must be signed in from a machine that meets the specified client-side certificate requirements (proper certificate authority (CA) and optionally specified field/value pair requirements) for each role to which the system may map the user. If the user's machine does not possess the certificate information required by a role, then the system does not map the user to that role.

  - **Administrators > Admin Roles >** *SelectRole* **> General > Restrictions > Certificate**

  - **Users > User Realms >** *Select Realm* **Role Mapping >** *Select|CreateRule* **>** *CustomExpression*

  - **Users > User Roles >** *SelectRole* **> General > Restrictions > Certificate**

- **When users request a resource**—The authenticated, authorized user must make a resource request from a machine that meets the specified client-side certificate requirements (proper certificate authority (CA) and optionally specified field/value pair requirements) for the resource policy corresponding to the user's request. If the user's machine does not possess the certificate information required by a resource, then the system does not allow the user to access the resource.

  - **Users > Resource Policies>** *SelectResource* **>** *SelectPolicy* **> Detailed Rules***Select|CreateRule* **>** *ConditionField*

Related Documentation

- Dynamic Policy Evaluation

Specifying Password Access Restrictions

You can restrict system and resource access by password-length when administrators or users try to sign in to the device. The user must enter a password whose length meets the minimum password-length requirement specified for the realm. Note that local user and administrator records are stored in the system authentication server. This server requires that passwords are a minimum length of 6 characters, regardless of the value you specify for the realm's authentication policy.

To specify password restrictions:

1. Select an administrator or user realm for which you want to implement password restrictions.

   Navigate to:

   - **Administrators > Admin Realms >** *Select Realm* **> Authentication Policy > Password**

   - **Users > User Realms >** *Select Realm* **> Authentication Policy > Password**

2. Choose one of the following options:

   - **Allow all users (passwords of any length)** — Does not apply password length restrictions to users signing in to the device.

   - **Only allow users that have passwords of a minimum length** — Requires the user to enter a password with a minimum length of the number specified.

     > NOTE: This option is not applicable for IKEv2 users and therefore is not enforced for IKEv2 users.

3. Select **Enable Password Management** if you want to enable password management. You must also configure password management on the authentication server configuration page (local authentication server) or through an LDAP server. For more information about password management,

4. If you have enabled a secondary authentication server, specify password length restrictions using the restrictions above as a guideline.

5. Click **Save Changes** to save your settings.

By default, the system requires that user passwords entered on the sign-in page be a minimum of four characters. The authentication server used to validate a user's credentials may require a different minimum length. The local authentication database, for example, requires user passwords to be a minimum length of six characters.

Related Documentation

- Dynamic Policy Evaluation

Specifying Session Limits

In addition to the access management options you may specify for an authentication policy, you may also specify a limit for concurrent users. A user who enters a URL to one of this realm's sign-in pages must meet any access management and concurrent user requirements specified for the authentication policy before the system presents the sign-in page to the user.

Setting the minimum or maximum setting limit amount allows you to configure realms that are more likely to be available when the device is nearing the amount of licensed users.

Valid numbers for the minimum amount of sessions are between 0 and the license limit. A default of 0 means there are no limits. All of the realms minimum limits can add up to the license limit but cannot exceed it. You cannot modify an existing realm's minimum limit or add a new realm's minimum limit that exceeds the license limit. The maximum limit can be equal to or

greater than the minimum limit for a particular realm. Value 0 for maximum limit means no user can log in to the realm.

You can also limit the number of concurrent users per session; a user can have multiple sessions. For example, if a user logs in from two machines in the same realm, an additional session is created. Each session counts towards the user license.

Users who enter through a realm with this feature enabled must have no more than the specified number of sessions open. If the user attempts to open a new session that exceeds the limit, a message appears that denies access or allows the user to continue or cancel.

When considering concurrent users per session, make note of the following:

- All session-related SSO attributes are saved in their respective session in the cache. These attributes are not shared with other sessions.

- All form-related SSO attributes are saved in their respective session in the cache. These attributes are not shared with other sessions.

- All Form-SSO related attributes will be saved in its respective session in the cache. The Form SSO state will not be shared with other sessions. The admin configured Form SSO values will be shared across all sessions.

- End-user's home page changes are reflected across all sessions. Any changes to the following will appear in the other concurrent sessions:

  - Bookmarks

  - Panel sorting (Preferences > User Home)

  - E-mail information, Daylight Saving Time, Pulse Collaboration (Preferences > General)

  - Autostart Client Application Session session (Preferences > Applications)

  - Cached E-mail Info settings (Preferences > Advanced)

  - Delete Cookies (Preferences > Advanced) now has options to let you remove cookies from the current session only or to remove cookies from all sessions.

  - Delete Password (Preferences > Advanced) now has options to let you remove passwords from the current session only or to remove passwords stored by all sessions.

- Cache Cleaner and Host Checker information is saved in each session. They are not shared across concurrent sessions

- Log messages will contain session identifiers (concatenated at the end of the log message) to differentiate which session the message refers to.

- Only one session can host a scheduled meeting. users cannot launch multiple scheduled meetings from concurrent sessions.

- Users can attend meetings from any sessions. However, since only one meeting client can be run per system, if a user wishes to attend more than one meeting, they must attend the other meetings from a different end-user system.

- Meeting host passes from one session to the other when you log out of a session. For example, suppose you are the meeting host, you join the meeting in user session A and then join the meeting again with user session B. User session A retains the meeting host.

However, if you are the meeting host from user session A, exit the meeting from user session A and then join the meeting in user session B then user session B assumes the meeting host role.

- Each user session maintains its own VPN Tunneling information. This information is not shared between concurrent sessions. However, administrator network connect sessions are shared between concurrent sessions.

- If you log in to the device as administrator, the first session is edit mode. If you log in as an administrator in a concurrent session, that administrator is logged in as read-only mode.

- VPN Tunneling bandwidth allocation is enforced on a per-session basis. For example, if a user is allocated a 1M bandwidth then each user session has a 1M bandwidth. The total bandwidth for this user is the number of sessions of this user times 1M.

- Users can launch terminal services, JSAM or WSAM from any session. Session information is saved per each session, they are not shared across concurrent sessions. Multiple instances of terminal services, JSAM and WSAM cannot be started on the same client.

- If a user has concurrent sessions and starts the e-mail client from multiple sessions, the e-mail client from the last session is the only one that can access the backend e-mail server through the system. For example, if a user has two concurrent sessions and starts the e-mail client from both sessions, only the second session can access the e-mail server.

> *i* NOTE: If you enable the multiple sessions per user feature, IKEv2 clients and VPN Tunneling clients may not be assigned the same IP address. For example, an IKEv2/VPN Tunneling client may be assigned a different VPN Tunneling VIP address each time they connect to the device when the system is obtaining the DHCP addresses from a DHCP server.

Use limits restrictions to set minimum and maximum concurrent users on the realm.

To specify the number of concurrent users limit restrictions:

1. Select an administrator or user realm for which you want to implement limits restrictions.

   - **Administrators > Admin Realms >** *SelectRealm***> Authentication Policy > Limits**

   - **Users > User Realms >** *SelectRealm* **> Authentication Policy > Limits**

2. To limit the number of concurrent users on the realm, select **Limit the number of concurrent users** and then specify limit values for these options:

   - **Guaranteed minimum**—You can specify any number of users between zero (0) and the maximum number of concurrent users defined for the realm, or you can set the number up to the maximum allowed by your license if there is no realm maximum.

   - **Maximum** (optional)—You can specify any number of concurrent users from the minimum number you specified up to the maximum number of licensed users. If you enter a zero (0) into the Maximum field, no users are allowed to login to the realm.

3. Click **Save Changes**.

To specify the number of concurrent users per session limit restriction:

1. Select **Authentication > Signing In > Sign-in Policies**.

2. Select the **Enable multiple user sessions** check box to allow users to have multiple concurrent sessions, and specify whether the user can log in when the maximum number of sessions is reached:

   - **Deny any more session from the user**—Displays a message saying the login is denied because it would exceed the maximum number of concurrent sessions.

   - **Allow the user to login**—Allows the user to log in. If the **Display open user session[s] warning notification** option is enabled, the user can select which session to close; otherwise the session that has been idle the longest is closed automatically.

3. Select the **Display open user session[s] warning notification** check box to allow users who have met the maximum session count to close one of their existing sessions before continuing with the current log in. If this option is disabled, the system terminates the session that has been idle the longest. This option applies only if **Enable multiple user sessions** is enabled along with **Allow the user to login**. Specify when the user is warned about concurrent sessions:

   - Select **Always** to notify users each time they log in when they already have another active session

   - Select **If the maximum session has been exceeded** to display the warning message only when the user's maximum session count has been met.

4. To specify the maximum number of concurrent sessions:

   1. Select **Users > User Realms >** *RealmName* **> Authentication Policy > Limits**.

   2. Specify the number of sessions permitted for users in the **Maximum number of sessions per user** text box.

   3. Click **Save Changes**.

      NOTE: If you do not select the **Enable multiple user sessions** check box, only one session per user is allowed regardless of the value you specify in the **Maximum number of sessions per user** text box.

Related Documentation

- Dynamic Policy Evaluation

IF-MAP Federation Overview

You can configure a Policy Secure device to store user session information for other Infranet Controllers and Connect Secure devices. Federation allows users to authenticate to a single Connect Secure or Infranet Controller, and then access resources that are protected by any number of Juniper Networks firewall devices known as Infranet Enforcers that are controlled by different Infranet Controllers. Federation enhances network performance. If a user is required to login to multiple Connect Secure or Infranet Controller devices during the course of a day to access different resources, each device must perform authentication and Host-Checking, often with periodic Host Checker updates throughout the day. The overhead can lead to decreased performance not only on the devices, but also on the network and the endpoint. Imported IF-MAP sessions eliminate redundant logins and Host Checks.

Federation on the device uses the standard IF-MAP (Interface for Metadata Access Point) protocol to share session information and other data between connected devices over distributed networks. IF-MAP is a protocol defined by the Trusted Network Connect Working Group (TNC-WG) as a standard interface between different network elements and devices. Federation is accomplished using an IF-MAP server and IF-MAP clients.

It is important as an administrator to understand the fundamental underlying communication method for data transmission in a Federation network over IF-MAP. Policies that you configure on the device permit this communication.

In a federated network, the IF-MAP server functions as the repository, or data store for IF-MAP clients to use for publishing information regarding activity on the network. For example, IF-MAP clients can publish information about sessions on the network, and Juniper Networks IDP devices can communicate information about potential threats to the IF-MAP client for publishing. IF-MAP clients can search for information about sessions or threats, and an IF-MAP client can establish a subscription so the IF-MAP server notifies the client when other clients publish new or changed information. In addition, IF-MAP clients can purge data that is no longer valid. All transactions are initiated by the IF-MAP client.

IF-MAP Federation is not supported for non-root IVSes on the device.

IF-MAP Federation is available on all Connect Secure devices with version 6.4 or greater. No licensing is required.

1. The endpoint authenticates through the IF-MAP client (Connect Secure). The IF-MAP client publishes session information to the IF-MAP server.

2. The endpoint attempts to access protected resources that are behind the Infranet Enforcer.

3. The Infranet Enforcer notifies the Infranet Controller (also an IF-MAP client). The IF-MAP client searches for session information on the IF-MAP server.

4. The Infranet Controller subscribes to session information about the endpoint's IP address.

5. The Infranet Controller notifies the Infranet Enforcer that session information exists for the IP address attempting to access resources, and the Infranet Enforcer provisions an auth table entry.

6. Access is granted to the protected resources. If any session information about the user changes, the authenticating IF-MAP client publishes the new information. Having subscribed to the user's session information, the Infranet Controller will be aware of any changes and provision access in accordance with the changed session information.

For details about configuring the system to work in an IF-MAP Federated network with the Infranet Controller, see *IF-MAP Feature Guide*.

IF-MAP Federation Workflow

Configuring an IF-MAP Federated network requires coordination between administrators of the different devices that will be in the federated network.

This document describes IF-MAP deployments that include only Pulse Secure devices: Infranet Controllers, Connect Secure devices, Infranet Enforcer firewalls, and Juniper Networks IDP. For

implementations that incorporate third-party components, contact Pulse Secure Technical Support.

The mix of devices in the federated network is important when planning the network. Will your network consist of only Infranet Controllers, or will you incorporate Connect Secure? Do the devices in your network have similar role mapping policies, or is each device different?

Determine and understand your goals for the federated network. The big picture guides your implementation as it becomes more complex. Pulse Secure recommends that you begin slowly. For example, start with a single role on each device, and then build the network incrementally.

In the simplest model, you can use the default policies. Using this model, you can quickly establish a federated network, and session information will automatically be shared among distributed devices in the network. This simple model should be adequate for most implementations in which the devices in the federated network have identical or very similar role mapping policies.

If your configuration requires more complex policies, you will need to perform a number of tasks to achieve your intended results. The following guidelines will help you plan your workflow:

- Ensure that communications between IF-MAP servers and IF-MAP clients is established

- Determine the resources that will be shared among the different devices

- Define who can access specific resources

- Distribute resources and users into roles

- Establish a naming convention that is shared and implemented between all administrators and devices

- Create Session-Export and Session-Import policies that reflect the role designations that you have configured on the devices

Related Documentation

- IF-MAP Federation Details

- Task Summary: Configuring IF-MAP Federation

IF-MAP Federation Details

You can configure the system as an IF-MAP client for an IF-MAP server. You configure an Infranet Controller as an IF-MAP server. Any endpoint sessions with an IP address created on an IF-MAP server are automatically published to that IF-MAP server.

You can create source IP policies for endpoints that authenticate to the device to permit access to resources behind Infranet Enforcers (ScreenOS Enforcers and Pulse Policy Secure s). Session-Export policies that you configure on the IF-MAP clients allow the clients to publish endpoint user data to the IF-MAP server. Devices that are IF-MAP clients can subscribe to the information on an IF-MAP server.

When a user accesses the device that is configured as an IF-MAP client, the client publishes basic session information, including the IP address, username and roles, to the IF-MAP server. The server stores the information as metadata. Other IF-MAP clients in the network can poll the

server for metadata when session information is needed as a result of an endpoint attempting to access protected resources behind an Infranet Enforcer.

When an authenticated user from the device that is configured as an IF-MAP client attempts to access resources that are protected by an Infranet Enforcer for an Infranet Controller that is also configured as an IF-MAP client, the Infranet Controller automatically provisions an auth table entry for the user on the Infranet Enforcer to allow access without requiring the user to authenticate to the Infranet Controller.

The Infranet Enforcer as an IF-MAP client subscribes to session information and other data for the endpoint based on the originating IP address. The authenticating device (the original IF-MAP client) publishes any changes in session parameters to the IF-MAP server. Since the Infranet Controller that is protecting the accessed resources subscribes to the metadata on the Federation server, session information is always current.

The Infranet Enforcer allows or denies traffic based on the resource access policies that are configured on the Infranet Controller to which it is connected.

You configure server settings on the Infranet Controller that will be the IF-MAP server. You configure client settings on each of the Connect Secure and Infranet Controller devices and that will be connected in the network.

In addition to the server and client settings, you configure Session-Export policies on Connect Secure and Infranet Controllers that are IF-MAP clients You configure and Session-Import policies on Infranet Controller IF-MAP clients that are connected to Infranet Enforcers.

IF-MAP allows servers and clients to publish, search, poll, and subscribe to data within a network of IF-MAP servers and clients. All of the data from Connect Secure in the network that is published to the IF-MAP server uses the IF-MAP protocol. Session-Export and Session-Import polices that you configure on Connect Secure and Infranet Controller allow the devices to utilize the IF-MAP protocol to share session information.

Session-Export policies specify how to translate an endpoint's session on Connect Secure or Connect Secure into IF-MAP data. To translate session information into IF-MAP data, you enter detailed user information. Connect Secure evaluates the Export policies to determine a session's IF-MAP roles, capabilities, identities, and device attributes and publishes the data to the IF-MAP server.

The Session-Import policies that you configure on Connect Secure specify how the device should derive a username and a set of roles based on IF-MAP data that it receives from the IF-MAP server from other Connect Secure devices. Import policies are similar to Role Mapping policies on a realm. You must be precise when configuring Export and Import policies, otherwise roles cannot be assigned properly.

The following figure depicts a scenario in which there are two Infranet Controllers configured as IF-MAP clients, one Connect Secure device configured as a IF-MAP client, and another Infranet Controller configured as the IF-MAP server. Endpoints that authenticate through any of the IF-MAP clients can access protected resources behind the enforcement point attached to Infranet Controller 1.

Figure14: IF-MAP Federated Network Topology
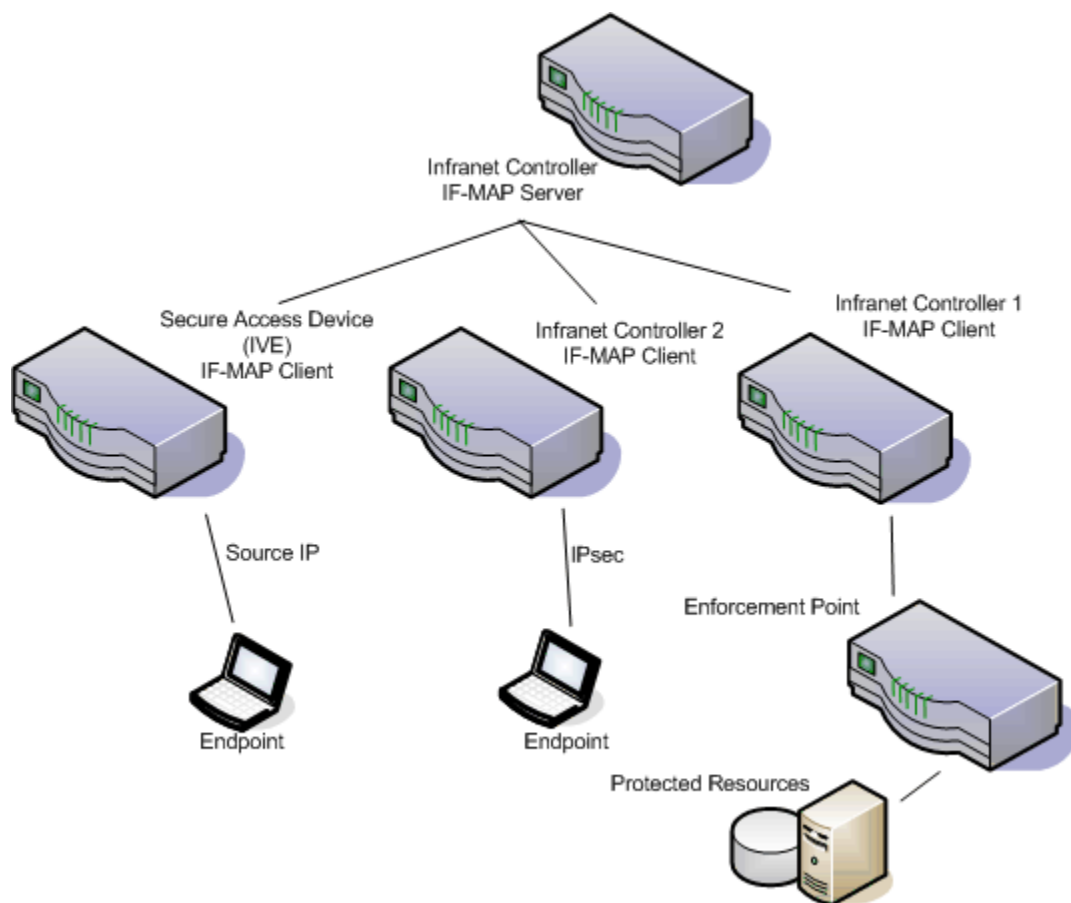
The interaction between the endpoints, the clients and the server is as follows:

- An endpoint authenticates through Connect Secure depicted in the figure and starts VPN Tunneling or Pulse Secure client.

- Connect Secure provisions an IP address for the endpoint to use on the internal network. Once the endpoint's IP address on the internal network is known, Connect Secure derives IF-MAP data from the endpoint's session.

- The Connect Secure IF-MAP client publishes the session information as IF-MAP data to the IF-MAP server using Session-Export policies.

- When the user attempts to access resources behind the enforcement point, access is blocked since the Infranet Enforcer has no information about the endpoint. The Infranet Enforcer sends out a dynamic discovery message that includes the endpoint's source IP address.

- Infranet Controller 1 uses the IP address to retrieve session data from the IF-MAP server.

- The Infranet Controller uses Session-Import policies to retrieve session data from the IF-MAP server.

The endpoint authenticating to Connect Secure must be running VPN Tunneling.

Imported user sessions do not count against the maximum user count for either platform, as each user is counted on the Connect Secure device from which they authenticated.

For details on configuring an IF-MAP Federation network, see *IF-MAP Feature Guide*.

IF-MAP Logging

IF-MAP related events are logged on both the IF-MAP server and the IF-MAP client.

Task Summary: Configuring IF-MAP Federation

The tasks listed in this topic are performed by an Policy Secure administrator, in conjunction with an administrator for Connect Secure. On Connect Secure, you configure Session-Export policies and you configure IF-MAP client settings. For details on configuring an IF-MAP Federated network, see *IF-MAP Feature Guide*.

To use IF-MAP Federation, perform the following tasks on Policy Secure and Connect Secure:

1. Enable dynamic auth table provisioning on any connected Infranet Enforcers that you want to use with Federation.

2. On Policy Secure, configure IF-MAP server settings to permit the server to communicate with IF-MAP clients.

3. Configure IF-MAP client settings to permit clients to communicate with the IF-MAP server.

4. On Policy Secure and Connect Secure, coordinate Session-Import policies, Session-Export policies, roles, and resource access policies between all of the clients in the Federated network.

5. Configure Session-Export policies on Connect Secure to define how sessions are translated into IF-MAP data.

6. Configure Session-Import policies on Connect Secure that correspond with Export policies to translate IF-MAP data into roles.

7. On Policy Secure, configure Source IP policies for Connect Secure users who will use Source IP to access the network.

Related Documentation

- Configuring IF-MAP Server Settings
- Configuring the IF-MAP Federation Client
- Session-Export and Session-Import Policies
- Troubleshooting the IF-MAP Federated Network

Configuring IF-MAP Server Settings

You must add all IF-MAP clients to the IF-MAP server. To add clients, you must specify the IP address and the security mechanism and credentials for the client.

For details on configuring an IF-MAP server, see *IF-MAP Feature Guide*.

Related Documentation

- Troubleshooting the IF-MAP Federated Network

Configuring the IF-MAP Federation Client

You must identify the IF-MAP server to each IF-MAP client. To add the server, you specify the IF-MAP URL of the server and how to authenticate to the server. Match the URL and security settings to equal those on the IF-MAP server(s) to which the IF-MAP client will connect.

To configure IF-MAP client settings on the devices that will be IF-MAP clients:

1. From the admin console select **System > IF-MAP Federation > Overview**.

2. Select the Enable IF-MAP Client option button.

3. Type the Server URL for IF-MAP Web service on the IF-MAP server. Append the server URL with **/dana-ws/soap/dsifmap** for all Pulse Secure IF-MAP servers.

4. Select the Client Authentication Method: Basic or Certificate.

    1. If you select Basic, enter a Username and Password. This is the same as the information that was entered on the IF-MAP server.

    2. If you select Certificate, select the Device Certificate to use.

    3. Ensure that the certificate of the CA that signed the IF-MAP server certificate is added from the System > Configuration > Certificates > Trusted Server CA page.

       The IF-MAP client validates the IF-MAP server certificate: if validation fails, the connection fails. Ensure that the hostname in the IF-MAP URL on the client machine matches the hostname of the server certificate on the IFMAP server, and that the CA that signed the server certificate is configured as a trusted server CA on the IF-MAP client.

5. Click **Save Changes**.

Related Documentation

- Configuring IF-MAP Server Settings

IF-MAP Federated Network Timing Considerations

It is important that the time on all IF-MAP servers is correct, as timeout issues are critical to ensure that IF-MAP provides complete and timely information. The IF-MAP Federation is designed to fail secure. If any component in the network does not receive timely information, the IF-MAP metadata will be purged from the data stores.

The components are designed to fail-secure. If complete and timely information cannot be provided, a user's session will be deleted. For example, if the chain of connections between an IF-MAP client that publishes a session and a client that grants access to a resource breaks, the client that granted access will remove the session. The fail-secure time limit is three minutes.

The timeout limit for IF-MAP is three minutes and applies to the following events:

- An IF-MAP server (or cluster) loses contact with one of its IF-MAP clients

- An IF-MAP server (or cluster) loses contact with one of its IF-MAP clients

- An IF-MAP server (cluster) loses contact with one of the other IF-MAP server (clusters) in the IF-MAP federation

- A Pulse Secure IF-MAP client loses contact with its IF-MAP server (cluster) for too long

Related Documentation

- Configuring IF-MAP Server Settings

- Configuring the IF-MAP Federation Client

- Session-Export and Session-Import Policies

- Troubleshooting the IF-MAP Federated Network

Session-Export and Session-Import Policies

You configure Session-Export policies on all of the Connect Secure and Infranet Controller devices in the Federated network that are IF-MAP clients. These policies allow IF-MAP clients to translate outgoing session information into IF-MAP data and incoming IF-MAP data into session information. These translations enable sessions to be shared between Connect Secure and Infranet Controller devices even if the devices sharing sessions have different role configurations.

To accurately configure Session-Export and Session-Import policies you need a minimal understanding of IF-MAP identifiers and IF-MAP metadata. An identifier is a unique value required for all metadata operations. Each instance of metadata is associated with an identifier. Examples of identifiers include access-request, identity, IP address, and MAC address. Examples of metadata include capability, role, and device-attribute.

IF-MAP recognizes two metadata types that are similar to roles on Connect Secure: IF-MAP roles and IF-MAP capabilities. An IF-MAP role is an attribute assigned to a user in the organization. When IF-MAP metadata is published to the IF-MAP server, this information could be one way to identify individuals on the network. This is somewhat different from the concept of roles. An IF-MAP capability is closer to the concept of a role. An IF-MAP capability is a collection of privileges assigned as a result of an access request. This is more analogous to a role since they are derived through role mapping in an authentication realm.

The data that is published to the IF-MAP server about a user session is derived by applying the Session-Export policies to the user session. The Session-Import policies are applied to the data from the IF-MAP server to assign a set of roles to the user.

When an endpoint attempts to access protected resources associated with Connect Secure, the device queries the IF-MAP server for data. The Infranet Controller uses Session- Import policies to derive roles and a username from the IF-MAP data. For example, you could configure a Session-Import policy that looks for a specific Host Checker policy (you specify the Host Checker policy in the Session-Import policy). If the Infranet Controller finds a match (in this case the Host Checker device attribute), the user can be assigned a role specified in the Session-Import policy.

97

All of the administrators who are configuring devices in the IF-MAP Federated network must agree on a set of capabilities, roles and device attributes and their meanings to be used with IF-MAP. Then, each administrator configures their device to map between local sessions and IF-MAP data. Figure 15 illustrates a coordinated IF-MAP Federated network configuration with policies that permit an example user to access protected resources.

Figure15: Session-Import and Session-Export Policies
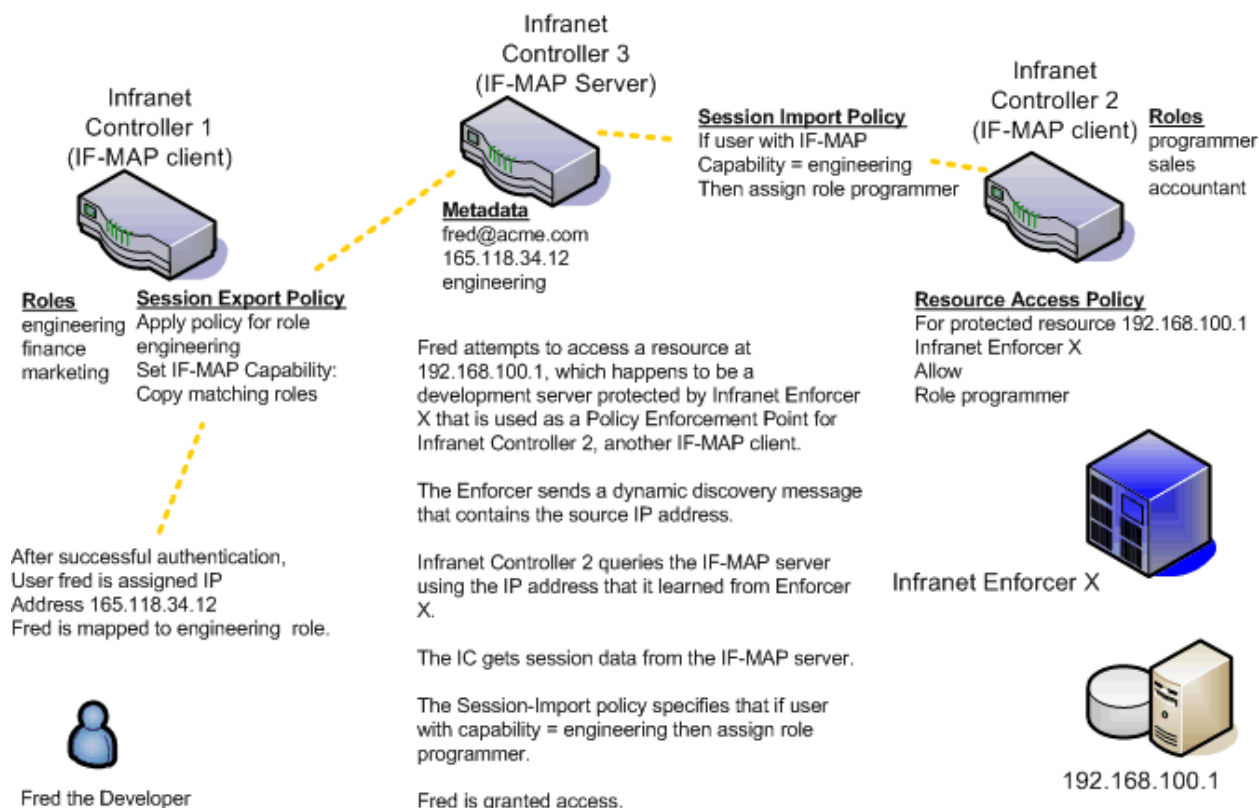
To further your understanding of Session-Import and Session-Export policies, please note the following Pulse Secure IF-MAP conventions:

- the system maps to the identical IF-MAP username.

- A role is paired with an IF-MAP capability.

- Capabilities can have the same name as the roles they are paired with, or a different name.

- When different IF-MAP clients have different but equivalent role names (e.g. Legal and Law, both referring to members of the corporate legal department) a single IF-MAP capability must be chosen.

- Not every role needs to be paired with an IF-MAP capability: roles can be local to Connect Secure.

- After you decide on pairings between IF-MAP capabilities and the roles, you create a session export policy for each pairing. On an Infranet Controller that controls Infranet Enforcers, you create a session import policy.

- The only parameters for the policies are the roles and the IF-MAP capability; everything else is fixed.

Default Session-Export and Session-Import Policy Action

By default, Session-Import and Session-Export IF-MAP policies are configured to allow IF-MAP capabilities (the equivalent of roles) to be published to the IF-MAP server and retrieved from the IF-MAP server, provided there are matching roles on each IF-MAP client. You can open new Session-Import and Session-Export policies on each device, and then name and close the policies. Any matching roles that the IF-MAP clients in the federated network have can be used to access resources.

Advanced Session-Export and Session-Import Policies

By default, advanced policy actions are not visible unless you click the advanced options links on the Session-Export and Session-Import policy pages. In default mode, you configure Session-Export and Session-Import policies using IF-MAP capabilities and roles.

Device attributes, IF-MAP roles and identities can be accessed through the advanced options links. IF-MAP capabilities and Connect Secure roles should provide the functionality that most IF-MAP Federation requires.

Related Documentation

- Configuring Session-Export Policies
- Session-Import Policies

Configuring Session-Export Policies

Session-Export policies determine how users are identified on the IF-MAP server when their session is published via IF-MAP: the policy sets the IF-MAP identifiers. You define attributes for users that will be used to determine role matching on different Infranet Controllers. For example, you might configure a Session-Export policy to specify that any users that belong to the "engineering" role should be identified with the "engineering" IF-MAP capability on the IF-MAP server. That identity will be included in the session information to which other IF-MAP clients subscribe. You configure corresponding Session-Import Policies on Infranet Controllers to identify which roles the user should be assigned.

You configure Session-Export policies based on Infranet Controller or Connect Secure roles, and users belonging to those roles can access resources on an Infranet Enforcer only if the role can be successfully matched with a role on the target Infranet Controller. You configure Session-Export policies on all Infranet Controller and Connect Secure devices for which you have users that will be allowed to access resources behind an Infranet Enforcer in the network.

When a user for whom Session-Export policies has been configured successfully authenticates to the network, the Session-Export policies are used to translate the user session into IF-MAP data which is then sent to the IF-MAP server. When the user attempts to access a resource that is protected by an Infranet Enforcer, the target Infranet Controller then attempts to translate the

IF-MAP data for the user into a username and roles using the Session-Import policies that are configured on the second Infranet Controller device.

Administrative Domains In Session-Export Policies

In a Layer 2 environment, session information on the IF-MAP server includes a MAC address. If an export policy specifies an Administrative Domain, the domain is associated with the MAC address published to the IF-MAP server (the administrative domain is also associated with the identity published to the IF-MAP server).

A DHCP server assigns an IP address to the endpoint after authentication. An IF-MAP enabled DHCP server publishes an ip-mac link to IF-MAP, associating the endpoint's IP address with its IF-MAP session information.

Including administrative domains in MAC addresses allows the ip-mac link to be created based on the administrative domain.

If your IF-MAP Federated network spans different administrative domains, you should configure separate Session-Export policies for each domain to prevent MAC address spoofing. Each administrative domain should have an associated DHCP server and unique Session-Export policies.

Other aspects of the Session-Export policies within the IF-MAP Federated network can overlap.

To configure a Session-Export policy:

1. From the admin console select **System > IF-MAP > Session-Export Policies**.

2. Click **New** to create a new policy.

3. Type a **Policy Name**, and optionally a **Description**.

4. Optionally, add **Available Roles** to the **Selected Roles** column to determine the roles for which this policy should apply. If you do not add any roles, the policy applies to all sessions. However, if you have non-interactive devices such as printers that do not need access, you may want to manually add roles and exclude those roles with non-interactive devices.

5. Under Policy Actions, Select **Set IF-MAP Capabilities** and choose the applicable roles.

   • **Copy Matching Roles**—Selecting this action copies all of the user roles that match the roles specified in the Roles section of this policy into the IF-MAP capabilities data.

   • **Copy all Roles**—Selecting this action copies all of the roles from the user session to the IF-MAP capabilities data.

   • **Set capabilities specified below**—Enter capabilities, one per line.

6. Select **Stop processing policies when this policy matches** to specify that when this policy is matched, no more Session-Export policies should be applied.

7. Select **Save Changes**, or continue to configure Advanced Actions.
To configure advanced options (generally not required for Infranet Controller and Connect Secure IF-MAP Federation):

1. Select the **View Advanced Actions** link. Additional options appear on the page.

2. **Set IF-MAP Identity**—If this action is chosen, enter the Identity and select an Identity Type from the menu. Identity is normally specified as <NAME>, which assigns the user's login name. Any combination of literal text and context variables may be specified. If you choose other for Identity Type, enter a unique Identity Type in the text box.

3. Optionally type the **Administrative Domain** for the Session-Export policy. This optional field is applied to identity and MAC address data. One example for using this field is in a large network environment with several domains in which a username could be duplicated. By entering the domain, you ensure that the correct user is identified.

4. **Set IF-MAP Roles**—If this action is selected, select the applicable roles.

   - **Copy Matching Roles**—Selecting this action copies all of the user roles that match the roles specified in the Roles section of this policy into the IF-MAP capabilities data.

   - **Copy all Roles**—Selecting this action copies all of the roles from the user session to the IF-MAP capabilities data.

   - **Set capabilities specified below**—Enter capabilities, one per line.

5. **Set IF-MAP Device Attributes**—Device attributes represent a passed Host Checker policy on the Infranet Controller or Connect Secure.

   - **Copy Host Checker policy names**—The name of each Host Checker policy that passed for the session is copied to a device attribute.

   - **Set Device Attributes**—Type Device Attributes, one per line, into the text box.

6. Select **Save Changes** to save this advanced Session-Export policy.

You must create corresponding Session-Import policies that allow IF-MAP client Infranet Controllers that are connected to an Infranet Enforcer in front of protected resources to collect IF-MAP data from the IF-MAP server.

Related Documentation

- Session-Export and Session-Import Policies

- Troubleshooting the IF-MAP Federated Network

Session-Import Policies

The Session-Export policies that you create allow IF-MAP data that represents a session to be stored on the IF-MAP server. Session-Import policies specify how the Infranet Controller derives a set of roles and a username from the IF-MAP data in the IF-MAP server. Session-Import policies establish rules for importing user sessions from Connect Secure. Import policies allow you to match authenticated users with corresponding roles on the target device. For example, you might configure an Import policy to specify that when IF-MAP data for a session includes the "Contractor" capability, the imported session should have the "limited" role. Session-Import policies allow the Infranet Controller to properly assign roles based on information that the IF-MAP server provides.

You configure Session-Import policies on IF-MAP client IVEs that are connected to an Infranet Enforcer in front of protected resources. For information about configuring Session-Import policies, see *IF-MAP Feature Guide*.

Related Documentation

- Troubleshooting the IF-MAP Federated Network

Troubleshooting the IF-MAP Federated Network

Diagnostic tools on the Infranet Controller and Connect Secure can assist you with troubleshooting a federated network.

IF-MAP Client User Messages—On the IF-MAP client, logs information that is published and removed from the IF-MAP server.

- Enable IF-MAP Client User Messages from **Log/Monitoring > User Access > Settings** on the Infranet Controller and Connect Secure IF-MAP client.

IF-MAP Server Trace—On the IF-MAP server, logs the XML for all IF-MAP requests and responses.

- Enable the IF-MAP Server Trace from **Log/Monitoring > Events > Settings** on the IF-MAP server.

  IF-MAP Server Trace should only be enabled for troubleshooting purposes, as running this diagnostic incurs a large performance impact.

Related Documentation

- Viewing Active Users on the IF-MAP Client

Viewing Active Users on the IF-MAP Client

On an IF-MAP client, you can view all of the sessions from other Infranet Controllers or Connect Securedevices that currently access the client (the imported sessions). Session information that can be viewed includes the username, roles, the user's endpoint IP address, and the IP address of the Infranet Controller or Connect Securedevice that authenticated the user. You can select and remove sessions either temporarily or permanently. A temporarily removed session can be restored in response to a request for continued access. A permanently removed session cannot be restored.

To view, de-activate, or activate current sessions on an IF-MAP client:

1. Select **System > IF-MAP > Active Users** from the IF-MAP client admin console.

2. Select **Imported** or **Exported**.

3. Select **Activate** or **De-activate**.

Related Documentation

- Troubleshooting the IF-MAP Federated Network

Trusted Server List

The system uses two mechanisms to install and launch client software from a web browser:

- ActiveX controls (available only for Windows/IE)

- Java applets

With both mechanisms, the user is prompted to trust ActiveX controls and Java applets they have not run before. Inherent problems with these types of mechanisms are:

- When the user trusts an ActiveX control that control is trusted forever.

- When trusting a Java applet, users are trusting all code that is signed by the exact same code signing certificate.

To address the above, administrators can create a text file (called a whitelist) that contains a list of trusted devices, fully qualified domain names or IP addresses, one per line. Administrators can configure two types of whitelists:

- Admin whitelist—The admin whitelist file can be modified only by the endpoint administrator. The administrator must use SMS or other mechanism to copy the admin whitelist file to the end-user's system. Admin whitelist files are located in:

  %ProgramFiles%\Pulse Secure\Whitelist.txt (Windows)

  /usr/local/pulse secure/whitelist.txt (Macintosh and Linux)

- User whitelist—Users can themselves make the decision to trust a device. When the user makes a decision to trust device, the device gets added to the user whitelist. User whitelist files are located in:

  %AppData%\Pulse Secure\Whitelist.txt (Windows)

  /~/Library/Application Support/Pulse Secure/whitelist.txt (Macintosh)

  /~/.pulse_secure/whitelist.txt (Linux)

  NOTE: The trusted server list feature is for applications launched from a browser window. It does not apply to applications launched from the command-line or other means.

Administrator and User Configuration

The following is a snippet of a whitelist file:

qa.juniper.netdev1.juniper.net66.129.224.48

NOTE: Whitelist files are not deleted when the software is removed.

There are two modes of enforcement:

- Allow Admin List Only—When software launches from the device that is not in the administrator whitelist, the launch fails and the user receives the error message "You are not allowed to launch software downloaded from <server>. Contact your system administrator for assistance." If the device is in the administrator whitelist, the launch proceeds as requested.

- Prompt—When software launches from the device that is not in the administrator whitelist or the user whitelist, the user is prompted if they want to launch the software with the message "Do you want to download, install and/or execute software from the following server". If the user declines, the launch fails. If the user accepts, the launch proceeds. The user also has the option to automatically add the device to the user whitelist file by selecting one of the following options from the message window:

- Always —Add the server to the user whitelist file and download, install or launch the software

- Yes—Download, install or launch the software but don't add the server to the user whitelist file

- No—Don't download, install or launch software and don't add the server to the user whitelist file

If the first line of the whitelist file contains "AllowAdminListOnly" (case insensitive) then Allow Admin List Only enforcement mode is used. Otherwise, prompt mode enforcement is used.

A snippet of a whitelist file using Allow Admin List Only enforcement is shown here:

AllowAdminListOnly qa.juniper.net dev1.juniper.net 66.129.224.48

NOTE: Prompt enforcement is the default mode when you upgrade your software to the latest revision.

To add clusters to the whitelist file:

- For Active/Passive clusters enter the VIP in the whitelist.

- For Active/Active clusters enter the load balancer hostname in the whitelist.

White List Flow Chart

The following steps outline the process for determining whether to launch the software

1. If the URL of the page initiating the launch does not begin with https, abort the launch and notify the user.

2. Else if the admin whitelist exists,

   - If the origin site is listed in the whitelist, proceed with the launch.

   - If the origin site is not in the whitelist and the whitelist starts with "AllowAdminListOnly", abort the launch and notify the user.

3. Else if the user whitelist exists,

   - If the origin site is in the user whitelist, proceed with the launch.

4. Prompt the user if they trust the origin site.

5. If the user agrees to trust the origin:

   - If they select Always then add the server to user whitelist file.

   - Proceed with the launch.

6. Abort the launch.

Related Documentation

- Uploading Java Applets to Connect Secure

CHAPTER 4 User Roles

User Roles Overview

A user role is an entity that defines user session parameters (session settings and options), personalization settings (user interface customization and bookmarks), and enabled access features (Web, file, application, Telnet/SSH, Terminal Services, network, meeting, and e-mail access). A user role does not specify resource access control or other resource-based options for an individual request. For example, a user role may define whether or not a user can perform Web browsing. However, the individual Web resources that a user may access are defined by the Web resource policies that you configure separately.

The access management framework supports two types of user roles:

- **Administrators**—An administrator role specifies management functions and session properties for administrators who map to the role. You can customize an administrator role by selecting the feature sets and user roles that members of the administrator role are allowed to view and manage. You can create and configure administrator roles through the Delegated Admin Roles page. Click Administrators > Admin Roles in the admin console.

- **Users**—A user role is an entity that defines user session parameters, personalization settings, and enabled access features. You can customize a user role by enabling specific access features, defining Web, application, and session bookmarks, and configuring session settings for the enabled access features. You can create and configure user roles through the Roles page. Click Users > User Roles in the admin console.

User roles are an integral part of the access management framework, and therefore are available on all Connect Secure products. However, you can only access features through a user role if you are licensed for the feature.

User Role Evaluation

The role mapping engine determines a user's session role, or combined permissions valid for a user session, as illustrated in Figure 16. A detailed description of each step follows the diagram.

Figure16: Security Checks Performed by Connect Secure to Create a Session Role
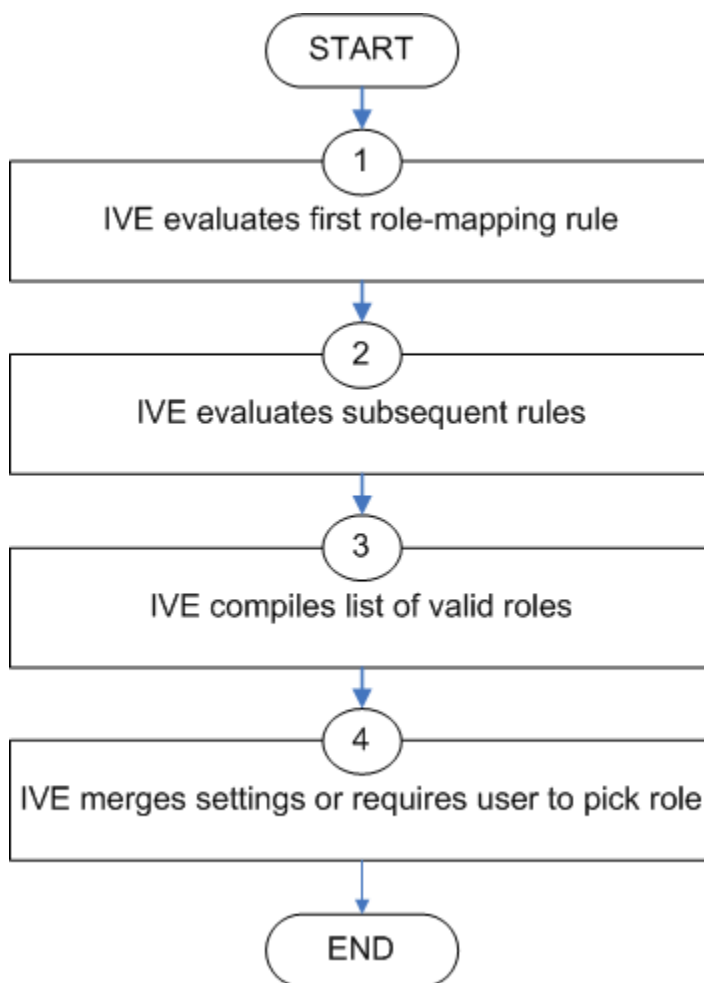
The system performs the following security checks to create a session role:

1. The system begins rule evaluation with the first rule on the Role Mapping tab of the authentication realm to which the user successfully signs in. During the evaluation, the system determines if the user meets the rule conditions. If so, then:

   - The system adds the corresponding roles to a list of "eligible roles" available to the user.

   - The system considers whether or not the "stop on match" feature is configured. If so, then the engine jumps to step 5.

2. The system evaluates the next rule on the authentication realm's Role Mapping tab according to the process in Step 1 and repeats this process for each subsequent rule. When

the system evaluates all role mapping rules, it compiles a comprehensive list of eligible roles.

3. The system evaluates the definition for each role in the eligibility list to determine if the user complies with any role restrictions. The system then uses this information to compile a list of valid roles, whose requirements the user also meets.

   If the list of valid roles contains only one role, then the system assigns the user to that role. Otherwise, the system continues the evaluation process.

4. The system evaluates the setting specified on the Role Mapping tab for users who are assigned to more than one role:

   - **Merge settings for all assigned roles**—If you choose this option, then the system performs a permissive merge of all the valid user roles to determine the overall (net) session role for a user session.

   - **User must select from among assigned roles**—If you choose this option, then the system presents a list of eligible roles to an authenticated user. The user must select a role from the list, and the assigns the user to that role for the duration of the user session.

   - **User must select the sets of merged roles assigned by each rule**—If you choose this option, the system presents a list of eligible rules to an authenticated user (that is, rules whose conditions the user has met). The user must select a rule from the list, and the system performs a permissive merge of all the roles that map to that rule.

   i NOTE: If you use automatic (time-based) dynamic policy evaluation or you perform a manual policy evaluation, the system repeats the role evaluation process described in this section.

Permissive Merge Guidelines

A permissive merge is a merge of two or more roles that combines enabled features and settings following these guidelines:

- Any enabled access feature in one role takes precedence over the same feature set disabled in another role. For example, if a user maps to two roles, one of which disables Pulse Collaboration while the other role enables Pulse Collaboration, the system allows the user to use Pulse Collaboration for that session.

- In the case of Secure Application Manager, the system enables the version corresponding to the first role that enables this feature. Furthermore, the system merges the settings from all the roles that correspond to the selected version.

  i NOTE: If you are using Pulse Secure client, then Pulse Secure is always enabled as the default client.

- In the case of user interface options, the system applies the settings that correspond to the user's first role.

- In the case of session timeouts, the system applies the greatest value from all of the roles to the user's session.

- If more than one role enables the Roaming Session feature, the system merges the netmasks to formulate a greater netmask for the session.

- When merging two roles that a user is mapped to—one in which bookmarks open in a new window and one in which bookmarks open in the same window—the merged role opens bookmarks in the same window.

- When merging two roles in which the first role disables the browsing toolbar and the second role enables either the framed or standard toolbar, the merged role uses the settings from the second role and displays the specified browsing toolbar.

- The merged role uses the highest value listed for each HTTP Connection Timeout. Click Users > User Roles > Select Role > Web > Options then click View advanced options.

- Merging of conflicting VPN Route Precedence Options is discouraged. But if it is done, the order of precedence is Allow Local Subnet Access, then Tunnel Routes and then Endpoint Routes.

Configuration of User Roles

To create a user role:

1. In the admin console, choose **Users > User Roles**.

2. Click **New Role** and then enter a name and optionally a description. This name appears in the list of Roles on the Roles page.

Once you have created a role, you can click the role's name to begin configuring it using the instructions in the following sections.

> *i* NOTE: When you delete a role, the personal bookmarks, SAM settings, and other settings may not be removed. Therefore, if you add a new role with the same name, any users added to that new role may acquire the old bookmarks and settings. In general, the system enforces referential integrity rules and does not allow you to delete any objects if they are referenced elsewhere. For example, if a role is used in any of the realm's role mapping rules, then the system rejects the deletion of the role unless you modify or delete the mapping rules.

> When you create individual user accounts, you must add the users through the appropriate authentication server (not the role). Or for instructions on how to create users on third-party servers, see the documentation that comes with that product.

Related Documentation

- Policies, Rules & Restrictions, and Conditions Overview

- Configuring General Role Options

Configuring General Role Options

Click Overview at the top of the General tab to edit a role's name and description, toggle session and user interface options on and off, and enable access features. When you enable an access feature, make sure to create corresponding resource policies.

To manage general role settings and options:

1. In the admin console, click **Users > User Roles >** *Role Name* **> General > Overview**.

2. Revise the name and description and then click **Save Changes** (optional).

3. Under Options, check the role-specific options that you want to enable for the role.

   The system uses default settings for newly created roles or when you do not select role-specific options.

   Role-specific options include:

   - **VLAN/Source IP**—Select this option to apply the role settings configured on the General > VLAN/Source IP page.

   - **Session Options**—Select this option to apply the role settings in the General > Session Options page to the role.

   - **UI Options**—Select this option to apply the role settings in the General > UI Options page to the role.

   - **Pulse Secure**—Select this option to download the desktop Pulse Secure client to Windows and MAC OS X users.

4. Under Access features, check the features you want to enable for the role. Options include:

   - **Web**—intermediate Web URLs through the Content Intermediation Engine.

   - **Files (Windows or UNIX/NFS version)**—resource profile that controls access to resources on Windows server shares or UNIX servers.

   - **Secure Application Manager (Windows version or Java version)**—provides secure, application-level remote access to enterprise servers from client applications.

   - **Telnet/SSH**—connect to internal server hosts in the clear using Telnet protocols or to communicate over an encrypted Secure Shell (SSH) session through a Web-based terminal session emulation.

   - **Terminal Services**—enable terminal emulation sessions on a Windows terminal server, Citrix NFuse server, or Citrix Metaframe server.

   - **Meetings**—securely schedule and hold online meetings between both system users and non-system users.

   - **Email Client**—enables users to use standards-based e-mail clients to access corporate e-mail securely from remote locations without the need for any additional software, such as a VPN client.

   - **VPN Tunneling**—provides secure, SSL-based network-level remote access to all enterprise application resources using the system.

   - **Secure Mail**—enables automatic synchronization with an Exchange server (ActiveSync) and e-mail encryption for iOS devices that have the Pulse Secure client.

   - **Enterprise Onboarding**—allows users to securely access enterprise network resources with almost any device. Wi-Fi, VPN, certificate, and Secure Mail profiles can be defined for enterprise resources and downloaded to a device during onboarding, depending on the device type.

5. Click **Save Changes** to apply the settings to the role.

Related Documentation

- Role Restrictions

- Specifying Role Session Options

Role Restrictions

Click Restrictions at the top of the General tab to specify access management options for the role. The system considers these restrictions when determining whether or not to map a user to the role. The system does not map users to this role unless they meet the specified restrictions.

You may configure any number of access management options for the role. If a user does not conform to all of the restrictions, the system does not map the user to the role.

To specify access management options for the role:

1. In the admin console, click **Users > User Roles >** *Role Name* **> General > Restrictions**.

2. Click the tab corresponding to the option you want to configure for the role, and then configure it.

Related Documentation

- Specifying Source IP Access Restrictions

- Specifying Browser Access Restrictions

- Specifying Certificate Access Restrictions

Specifying Role-Based Source IP Aliases

Click VLAN/Source IP at the top of the General to define role-based source IP aliases. If you want to direct traffic to specific sites based on roles, you can define a source IP alias for each role. You use these aliases to configure virtual ports you define for the internal interface source IP address. A back-end device can then direct end user traffic based on these aliases, as long as you configure the back-end device, such as a firewall, to expect the aliases in place of the internal interface source IP address. This capability enables you to direct various end users to defined sites based on their roles, even though all of the end user traffic has the same internal interface source IP address.

NOTE: You must define virtual ports to take advantage of the role-based source IP aliases.

To specify a source IP alias for the role:

1. In the admin console, click **Users > User Roles >** *Role Name* **> General General > VLAN/Source IP**.

2. Select the VLAN you want to use from the VLAN list, if you have defined VLAN ports on your system.

   If you have not defined VLAN ports, the option defaults to the Internal Port IP address. If you have provisioned IVS systems, and you have defined VLAN ports and you want any of those VLAN ports to appear in the VLAN list, then you must include the VLAN ports in the Selected VLANs text box on the Root IVS configuration page.

3. Select a source IP address from the list.

4. Click **Save Changes** to apply the settings to the role.

> *i*   NOTE: If an end user is mapped to multiple roles and the system merges roles, the system associates the source IP address configured for the first role in the list with the merged role.
>
> You can specify the same source IP address for multiple roles. You cannot specify multiple source IP addresses for one role.

Related Documentation

- Using Virtual Ports

Specifying Role Session Options

Use the Session tab to specify session time limits, roaming capabilities, session and password persistency, request follow-through options, and idle timeout application activity. Select the **Session Options** check box on the Overview tab to enable these settings for the role.

To configure general session options:

1. In the admin GUI, click **User > User Roles >** *RoleName* **> General > Session Options**.

2. Configure session options, as described in Table 5.

3. Save the configuration.

Table5: Session Options

Option	Guidelines
Session lifetime	• For Idle Timeout, specify the number of minutes a nonadministrative user session may remain idle before ending. The minimum is five minutes. The default idle session limit is 10 minutes, which means that if a user's session is inactive for 10 minutes, the system ends the user session and logs the event in the system log (unless you enable session timeout warnings described later).
	• For Max. Session Length, specify the number of minutes an active nonadministrative user session may remain open before ending. The minimum is six minutes. The default time limit for a user session is 60 minutes, after which the system ends the user session and logs the event in the system log. During an end user session, prior to the expiration of the maximum session length, the system prompts the user to reenter authentication credentials, which avoids the problem of terminating the user session without warning.
	• For Reminder Time, specify when the system should prompt non-administrative users, warning them of an impending session or idle timeout. Specify the number of minutes before the timeout is reached.
	• Optionally, select **Use Session/Idle timeout values sent by the primary Radius authentication Server** to override the idle timeout and session length specified above. If the received values are below the minimums (5 minutes for the idle timeout and 6 minutes for the session length), the minimum values are used.
	• Optionally, select **Enable Session Extension** to allow users to extend the session beyond the maximum session length. If this feature is enabled, users will be reauthenticated and extend their current session without interruption.
	NOTE: We recommend the difference between Idle Timeout and Reminder Time be greater than two minutes. This ensures that the reminder pop-up window appears at the correct time.

Enable session timeout warning	Enable to notify non-administrative users when they are about to reach a session or idle timeout limit. These warnings prompt users to take the appropriate action when they are close to exceeding their session limits or idle timeouts, helping them save any in-progress form data that would otherwise be lost. Users approaching the idle timeout limit are prompted to reactivate their session. Users approaching the session time limit are prompted to save data. For example, a user may unknowingly reach the idle timeout set for his role while using an e-mail client configured to work with the system, because the system does not receive data while the user composes e-mail. If the session timeout warning is enabled, however, the system prompts the user to reactivate his session before the session times out and forces the user's session to end. This warning gives the user the opportunity to save his partially composed e-mail. Optionally, select **Display sign-in page on max session time out** to display a new browser sign-in page to the end user when their session times out. This option only appears when you choose to enable the session timeout warning. NOTE: If you do not select the **Enable session timeout warning** option, the system only displays expiration messages to users—it does not give them the option to extend their sessions. Instead, users need to access the sign-in page and authenticate into a new session. The Enable session timeout warning option only applies to expiration messages displayed by the end user's browser, not by other clients such as WSAM or VPN Tunneling.
Roaming session	Select one of the following options: • **Enabled**—Enables unlimited roaming sessions. An unlimited roaming session allows mobile users (laptop users) with dynamic IP addresses to sign in to the device from one location and continue working from another. Disable this feature to prevent users from accessing a previously established session from a new source IP address. This helps protect against an attack spoofing a user's session, provided the hacker was able to obtain a valid user's session cookie. If you enable unlimited session roaming, a session is maintained within an IPv4 network, within an IPv6 network, or from IPv4 to IPv6 and vice versa. • **Limit to subnet**—Limits the roaming session to the local subnet, specified by netmask for IPv4 subnets and prefix length for IPv6 subnets. Users may sign in from one IP address and continue using their sessions with another IP address as long as the new IP address is within the same subnet. If you configure limited session roaming, you can specify IPv4 or IPv6 subnets within which the session is maintained. However, with limited session roaming, you cannot allow sessions to roam from IPv4 to IPv6 networks, or vice versa. • **Disabled**—Disables roaming user sessions for users mapped to this role. Users who sign in from one IP address may not continue an active session from another IP address; user sessions are tied to the initial source IP address.

Persistent session	By default, the session cookie is flushed from the browser's memory when the browser is closed. The session length is determined by both the idle timeout value and maximum session length value that you specify for the role. The session does not terminate when a user closes the browser; a session only terminates when a user signs out of the device. Enable the persistent session option to write the session cookie to the client hard disk so that the user's credentials are saved for the duration of the session. Assume persistent session is enabled and a user starts a VPN Tunneling session from a browser. Later, the user quits the browser application. The next time the user opens a new browser window and logs in to the same device, the user is not prompted to enter his or her credentials again. NOTE: (Macintosh only) Persistent session applies only for browser login as stated earlier. If you start VPN Tunneling from the standalone launcher (by opening NetworkConnect.dmg) and later open a new browser and log in to that same device, you are prompted to reenter your credentials. NOTE: If you enable the Persistent session option and a user closes the browser window without signing out, any user can open another instance of the same browser to access the device without submitting valid credentials, posing a potential security risk. We recommend that you enable this feature only for roles whose members need access to applications that require system credentials and that you make sure these users understand the importance of signing out of the device when they are finished.
Remove Browser Session Cookie	Enable to remove the session cookie and logs users out of their Web session once the client component launches, enhancing security for your VPN Tunneling, WSAM and Pulse session. Disable to retain the session cookie and keep users logged in to their Web session once the client component starts. Because browser cookies are plain text files, they are susceptible to malicious attacks. The Remove Browser Session Cookie option removes the session cookie, making your VPN Tunneling, WSAM and Pulse Secure client sessions more secure. When enabled, users are logged out of their Web session once the client component (for example, Pulse Secure) launches. Users are logged out of their Web session regardless of whether the client component launches successfully or not. If the client component does not successfully launch, users can restart their Web session and try launching their client component again. This option also prevents any client component from launching a browser through the client. NOTE: The Remove Browser Session Cookie removes only the session cookie. It does not remove non-system cookies or other any other cookie.
Persistent password caching	Enable to allow cached passwords to persist across sessions for a role. The system supports Windows NT LAN Manager (NTLM) authentication protocol and HTTP Basic Authentication and supports servers that are set up to accept both NTLM and anonymous sign-in. The system caches NTLM and HTTP Basic Authentication passwords provided by users so that the users are not repeatedly prompted to enter the same credentials used to sign in to the server or another resource in the NT domain. By default, the system flushes cached passwords when a user signs out. A user can delete cached passwords through the Advanced Preferences page. After the end user logs in to the device, click **Preferences** and then click the **Advanced** tab.
Browser request follow-through	Enable to allow the system to complete a user request made after an expired user session after the user reauthenticates.
Idle timeout application activity	Enable to ignore activities initiated by Web applications (such as polling for e-mails) when determining whether a session is active. If you disable this option, periodic pinging or other application activity may prevent an idle timeout.
Upload Logs	Enable to allow the user to transmit (upload) client logs to the system. NOTE: Use the System > Log/Monitoring > Client Logs > Settings page to completely enable client-side logs for the user.

Related Documentation

- Pulse Collaboration Overview

Customizing the Welcome Page

Click **UI Options** at the top of the General tab to specify customized settings for the welcome page and the browsing toolbar for users mapped to this role. The welcome page (or home page) is the Web interface presented to authenticated users. Click **Overview** at the top of the General tab, and then select the **UI Options** check box to enable custom settings for the role; otherwise, the system uses the default settings.

Personalization settings include the sign-in page, page header, page footer, and whether or not to display the browsing toolbar. If the user maps to more than one role, then the system displays the user interface corresponding to the first role to which a user is mapped.

To customize the welcome page for role users:

1. Click **Users > User Roles >** *RoleName* **> General > UI Options**.

2. Under Header, specify a custom logo and alternate background color for the header area of the welcome page (optional):

   - Click **Browse** and locate your custom image file. The new logo appears in the Current appearance box only after you save your changes.

     > ℹ️ NOTE: You can only specify a JPEG or GIF file for a custom logo image. Other graphics formats are not displayed properly in the JSAM status window on some OS platforms.

   - Type the hexadecimal number for the background color or click the Color Palette icon and pick the desired color. The Current appearance box updates immediately.

3. Under Sub-headers, select new background and text colors (optional):

   - Type the hexadecimal number for the Background color or click the **Color Palette** icon and pick the desired color. The Current appearance box updates immediately.

   - Type the hexadecimal number for the Text color or click the **Color Palette** icon and pick the desired color. The Current appearance box updates immediately.

4. Under Start page, specify the start page that you want users to see after they sign in and when they click the Home icon on the toolbar:

   - **Bookmarks page**—Select this option to display the standard Bookmarks page.

   - **Meetings page**—Select this option to display the standard meetings page.

   - **Custom page**—Select this option to display a custom start page and then specify the URL to the page. The system rewrites the URL and creates an access control rule to allow users access to the URL. (Note that users can also enter the custom URL in the Browse field on the toolbar.) The system evaluates the access control rule after all other policies, which means another policy could deny access to the URL.

   - **Also allow access to directories below this url**—Select this option to allow users access to subdirectories of the custom-page URL. For example, if you specify http://www.domain.com/, users can also access http://www.domain.com/dept/.

5. Under Bookmarks Panel Arrangement, arrange the panels as you want to display them on the user's bookmarks page:

   1. To select the name of a panel, click in the Left Column or Right Column list.

   2. To position a panel above or below the other panels, click Move Up or Move Down.

   3. To move a panel to the other side of the user's bookmarks page, click Move > or < Move.

   > **NOTE:** The system displays all panels under Bookmarks Panel Arrangement for all licensed features regardless of whether or not you enable the corresponding feature for the role.
   >
   > The maximum number of combined bookmarks a role can have is approximately 500. If a role has more than 500 bookmarks, some operations (for example, delete role, duplicate role) may not work correctly. The workaround is to split a role with a large number of bookmarks into multiple roles.

4. Under Help Page, select options to control the Help page that appears when users click the Help button on the toolbar:

   - **Disable help link**—Select this option to prevent users from displaying Help by removing the Help button from the toolbar.

   - **Standard help page**—Select this option to display the standard end-user Help.

   - **Custom help page**—Select this option to display a custom Help page. Specify the URL to the custom help page, and then provide an optional width and height for the help page's window. The system rewrites the URL and creates an access control rule to allow users access to the URL. (Note that users can also enter the custom URL in the Browse field on the toolbar.) The system evaluates the access control rule after all other policies, which means another policy could deny access to the URL. (Note that when you choose this option, the system disables the Tips link next to the Browse field.)

   - **Also allow access to directories below this url**—Select this option to allow users access to subdirectories of the custom help page URL. For example, if you specify http://www.domain.com/help, users can also access http://www.domain.com/help/pdf/.

5. Under User Toolbar, select options for the toolbar on the Bookmarks page and other secure gateway pages:

   - **Home**—Select this option to display the Home icon on the Bookmarks page and other secure gateway pages.

   - **Preferences**—Select this option to display the Preferences button.

   - **Session Counter**—Select this option to display a time value on the user toolbar that indicates the maximum remaining time allowed in the user's current session. Note that a period of user inactivity could also end the current session before this maximum time expires.

   - **Client Application Sessions**—Select this option to display the Client Apps button on the user toolbar. Users can click this button to display the Client Application Sessions page where they can start client applications such as VPN Tunneling or Secure Application Manager. If you do not select this option, the system displays the Client Application Sessions panel on the Bookmarks page.

6. Under Browsing toolbar, select options for the toolbar that users see when browsing pages not located on the system, such as external web sites:

- Show the browsing toolbar—Select this option to display the browsing toolbar.

- Toolbar type—Select the type of browsing toolbar you want to display:

  - Standard—This toolbar can be moved to the top left or top right side of the browser window. Users can also collapse and expand the toolbar. When collapsed, the toolbar displays the custom logo only. The toolbar's default state is expanded and on the top right side of the browser window.

  - Framed—This toolbar remains fixed in a framed header section at the top of the page.

    NOTE: We recommend that you do not use the top variable when working with a frame set because after the system intermediates the page, top might reference a different frame than you intend. This change might make the framed toolbar disappear or could cause your intermediated application to work erratically or incorrectly. See *Content Intermediation Engine Developer Guide*,

- Toolbar logo and Toolbar logo (mobile)—Specify a custom logo (such as your company's logo) that you want to display on the standard and framed toolbars by browsing to the image file (optional). When the user clicks the logo, the page you specify for the Logo links to option appears. The current logo for the browsing toolbar appears next to these options.

- Logo links to— Select an option to link the browsing toolbar logo to a page that appears when users click the logo:

  - Bookmarks page—Links the logo to the Bookmarks page.

  - "Start Page" settings—Links the logo to the custom start page you specified under the Start Page section.

  - Custom URL—Links the logo to the URL you enter in the associated text box (optional). This resource must be accessible to the system. The system rewrites the URL and creates an access control rule to allow users access to the URL. (Note that users can also enter the custom URL in the Browse field on the toolbar.) The system evaluates the access control rule after all other policies, which means another policy could deny access to the URL.

  - Also allow access to directories below this url—Select this option to allow users access to subdirectories of the custom URL.

- Specify the items you want to display in the browsing toolbar:

  - Enable "Home" link—Select this option to display the Home Page button, which is linked to the Bookmarks page.

  - Enable "Add Bookmark" link—Select this option to display the Bookmark this Page button.

  - Enable "Bookmark Favorites" link—Select this option to display the Bookmark Favorites button. When the user clicks this button, the system displays a list of the bookmarks

that the user specified as favorites on the Add Web Bookmark page of the secure gateway.

- Display Session Counter— Select this option to display a time value on the browsing toolbar that indicates the maximum remaining time allowed in the user's current session. Note that a period of user inactivity could also end the current session before this maximum time expires.

- Enable "Help" link—Select this option to display the Help button, which is linked to the Help page you specify for under Help page.

> **NOTE:** If you click **Users > User Roles >** *Role Name***> Web > Options** and clear the **User can add bookmarks** check box, then the system does not display the Bookmark this Page and Bookmark Favorites buttons on the browsing toolbar even if you select the **Enable "Add Bookmark"** link and **Enable "Bookmark Favorites"** link options.

- Use Iframe in Toolbar—Select this option if you are having problems with using iframes with JavaScript rewriting and with the Firefox web browser. This option resolves interoperability problems with the above.

7. Under Personalized greeting, specify a greeting and notification message on the Bookmarks page (optional):

- Enabled—Select this option to display the personalized greeting. The system displays the username if the full name is not configured.

- Show notification message—Select this option and enter a message in the associated text box (optional). The message appears at the top of the Bookmarks page after you save changes and the user refreshes that page. You may format text and add links using the following HTML tags: <i>, <b>, <br>, <font> and <a href>. However, the system does not rewrite links on the sign-in page (because the user has not yet authenticated), so you should only point to external sites. Links to sites behind a firewall will fail. You may also use Connect Secure system variables and attributes in this field.

> **NOTE:** The length of the personalized greeting cannot exceed 12K, or 12288 characters.
>
> If you use unsupported HTML tags in your custom message, the system may display the end user's home page incorrectly.

8. Under Other, specify whether or not you want the copyright notice and label shown in the footer (optional). This setting applies only to those users whose license permits disabling the copyright notice. For more information about this feature, call Pulse Secure Support.

9. Click **Save Changes**. The changes take effect immediately, but current user browser sessions may need to be refreshed to see the changes.

10. Click **Restore Factory Defaults** to reset all user-interface options back to factory defaults (optional).

Related Documentation

- Specifying Role Session Options

Optimized Interface for the Apple iPad

The system is optimized for a number of platforms, including the Apple iPad. This optimization includes:

- Login pages – includes the login and logout pages as well as intermediate pages that appear after the user enters their credentials on the sign-in page and before the Home page appears. The following is a list of these customized login pages:

    - Cancel.thtml

    - Defender.thtml

    - ExceededConcurrent.thtml

    - GeneratePin.thtml

    - GraceLoginUsed.thtml

    - LoginPage.thtml

    - Logout.thtml

    - NewPin.thtml

    - NextToken.thtml

    - PasswordChange.thtml

    - PasswordExpiration.thtml

    - SelectRole.thtml

    - ShowSystemPin.thtml

    - SigninNotifPostAuth.thtml

    - SigninNotifPreAuth.thtml

    - SM-NewPinSelect.thtml

    - SM-NewPinSystem.thtml

    - SM-NewUserPin.thtml

    - SM-NextToken.thtml

    - SSL.thtml

    - confirmation.thtml

    - confirmation_opensessions.thtml

    - user_unknown.thtml

- Home page – This home page displays the welcome panel and any applicable notification messages as well as the Web Bookmark panel, the File Bookmark (or Files) panel, the VPN and Preferences button.

- Web Bookmark pages – Located on the home page, the Web Bookmark panel lists each individual bookmarks and allows user to tap and browse the bookmark destination page. To to edit bookmarks, tap the **Edit** button on the panel header and the Edit Bookmark page

appears. On this page, user can edit individual bookmarks, reorder bookmarks, and delete bookmarks. Editing is limited to user-created bookmarks.

- File Bookmark pages – Located on the home page, the File Bookmark panel lists each individual bookmarks. To edit bookmarks, tap the **Edit** button on the panel header and the Edit File Bookmark page will be displayed. On this page, user can edit individual bookmarks, reorder bookmarks, and delete bookmarks. Editing is limited to user-created file bookmarks.

- Preferences page – Located the home page is a Preferences button. When tapped, it displays the Preferences setting page, containing configuration options for changing username, delete cookies, delete session cookies and delete passwords.

- Error pages – Error pages that can be seen while using the features made available on the iPad are customized.

- Company logos – Most pages display a company logo. These pages are capable of displaying custom logos if uploaded from the admin GUI.

Table 6 lists the supported configurable options on the Apple iPad.

Table6: Configurable Options on the Apple iPad

Custom User Interface Options	Supported
Header Logo Image	Yes
Header Background Color	No
Sub-Header Background Color	No
Sub-Header Text Color	No
Start Page Message (Welcome message)	Yes
Bookmark Panel Arrangement	No
Enable/Disable Help Link	Yes
Window Size of Help Page	No
Show/Hide Preferences Toolbar	No
Show/Hide Session Counter	Yes
Browsing Toolbar Items	Yes
Post-Auth Sign-In Notification	Yes
Personalized Greetings	Yes
Show Copyright Notice in Footer	No

Related Documentation

- Customizing the Welcome Page

Defining Default Options for User Roles

You can define default options for all user roles, just as you can for delegated administrator roles. Default values are used for newly created roles or for roles where the session or UI

option check boxes are not selected in the User > User Roles > *UserName* > General > Overview window.

The default options include, but are not limited to:

- Session Options

  - **Session lifetime**—Define the idle timeout, maximum session length, and reminder time in minutes.

  - **Enable session timeout warning**—Determine whether to display warning and login page.

  - **Roaming Session**—Define level of mobility access.

  - **Persistent Session**—Define state across browser instances.

  - **Persistent password caching**—Define password state across sessions.

  - **Browser request follow-through**—Define response to browser session expiration.

  - **Idle timeout application activity**—Define system response to application session activity.

- UI Options

  - **Header**—Define the logo and background color.

  - **Sub-headers**—Define the background and text color.

  - **Start page**—Define which page appears after the user logs in.

  - **Bookmarks Panel Arrangement**—Define the panels that appear on the user's bookmark page.

  - **Help Page**—Display standard or custom help.

  - **User Toolbar**—Define the links that appear on a user's home page.

  - **Browsing toolbar**—Define the links that appear when a user is browsing an external web site.

  - **Personalized Greeting**—Display user's name and notification message on the user's welcome page.

  - **Bookmarks Panel Arrangement** Other—Show copyright notice.

Defining Default Options for User Roles

To define the default options for all user roles:

1. Select **Users > User Roles**.

2. Click **Default Options**.

3. Modify settings in the **Session Options**, **UI Options**, and **Custom Messages** tabs.

4. Click **Save Changes**. These become the new defaults for all new user roles.

If you do not want user roles to see the copyright notice, you can also clear the Show copyright notice and "Secured by Pulse Secure" label in footers check box for user roles, in general. That way, all subsequent roles you create do not allow the notice to appear on the end user UI.

Related Documentation

- Configuring General Role Options

- Role Restrictions

Customizing Messages

You can customize basic messages that may be displayed to your end users when they sign in to the device. You can change the message text, and you can add internationalized versions of the messages in Chinese (Simplified), Chinese (Traditional), French, German, Japanese, Korean, and Spanish, in addition to English.

To customize messages:

1. Select **Users > User Roles**.

2. Click **Default Options**.

3. Select the **Custom Messages** tab.

4. Select the language to use from the menu.

5. Enter your text in the Custom Message box, below the default message you want to override.

6. Click **Save Changes**.

7. Repeat the process to create messages in additional languages.

Related Documentation

- About Multi-Language Support for Connect Secure

- Localizing the User Interface

- Localizing Custom Sign-In and System Pages

Customizing UI Views for User Roles

You can use customization options on the Roles page to quickly view the settings that are associated with a specific role or set of roles. For instance, you can view all of the user roles and any Web bookmarks that you have associated with them. Additionally, you can use these customized views to easily link to the bookmarks and other configuration settings associated with a role.

To view a sub-set of data on the Roles page:

1. Click **Users > User Roles**.

2. Select an option from the View list at the top of the page. Table 7 zdescribes these options.

3. Select one of the following options from the For list:

- **All roles**—Displays the selected bookmarks for all user roles.

- **Selected roles**—Displays the selected bookmarks for the user roles you choose. If you select this option, select one or more of the check boxes in the Role list.

4. Click **Update**.

Table7: View Menu Options

Option	Description
Enabled Settings	Displays a graph outlining the remote access mechanisms and general options that you have enabled for the specified roles. Also displays links (the check marks) that you can use to access the corresponding remote access and general option configuration pages.
Restrictions	Displays Host Checker and Cache Cleaner restrictions that you have enabled for the specified roles. Also displays links you can use to access the corresponding Host Checker and Cache Cleaner configuration pages.
Meetings	Displays Pulse Collaboration settings that you have configured for the specified roles. Also displays links you can use to access the corresponding Pulse Collaboration configuration pages.
VPN Tunneling	Displays VPN Tunneling settings that you have configured for the specified roles. Also displays links you can use to access the corresponding VPN Tunneling configuration pages.
Role Mapping Rule & Realms	Displays the assigned authentication realms, role mapping rule conditions, and permissive merge settings for the specified roles. Also displays links you can use to access the corresponding realm and role mapping configuration pages.
Bookmarks: All	Displays the names and types of all of the bookmarks that you have enabled for the specified roles. Also displays links you can use to access the corresponding bookmark configuration pages. (Note that if you created a bookmark through a resource profile, the link appears in the Resource column. Otherwise, the link appears in the Bookmark column.)
Bookmarks: Web	Displays the Web bookmarks that you have enabled for the specified roles. Also displays links you can use to access the corresponding bookmark configuration pages. (Note that if you created a bookmark through a resource profile, the link appears in the Resource column. Otherwise, the link appears in the Web Bookmark column.)
Bookmarks: Files (Windows)	Displays the Windows File bookmarks that you have enabled for the specified roles. Also displays links you can use to access the corresponding bookmark configuration pages. (Note that if you created a bookmark through a resource profile, the link appears in the Resource column. Otherwise, the link appears in the Windows File Bookmark column.)
Bookmarks: Files (UNIX)	Displays the UNIX/NFS File bookmarks that you have enabled for the specified roles. Also displays links you can use to access the corresponding bookmark configuration pages. (Note that if you created a bookmark through a resource profile, the link appears in the Resource column. Otherwise, the link appears in the UNIX File Bookmark column.)
Bookmarks: Telnet	Displays the Telnet/SSH bookmarks that you have enabled for the specified roles. Also displays links you can use to access the corresponding bookmark configuration pages. (Note that if you created a bookmark through a resource profile, the link appears in the Resource column. Otherwise, the link appears in the Telnet/SSH Session column.)
Bookmarks: Terminal Services	Displays the Terminal Services bookmarks that you have enabled for the specified roles. Also displays links you can use to access the corresponding bookmark configuration pages. (Note that if you created a bookmark through a resource profile, the link appears in the Resource column. Otherwise, the link appears in the Terminal Services Session column.)
ACL Resource Policies: All	Displays the resource policies that are associated with the specified roles. Includes the type, name, description, action, and resources for each policy. Also displays links you can use to access the corresponding policy configuration pages.
ACL Resource Policies: Web	Displays the Web resource policies that are associated with the specified roles. Includes the type, name, description, action, and resources for each policy. Also displays links you can use to access the corresponding policy configuration pages.

ACL Resource Policies: Files (Windows)	Displays the Windows file resource policies that are associated with the specified roles. Includes the type, name, description, action, and resources for each policy. Also displays links you can use to access the corresponding policy configuration pages.
ACL Resource Policies: Files (UNIX)	Displays the UNIX file resource policies that are associated with the specified roles. Includes the type, name, description, action, and resources for each policy. Also displays links you can use to access the corresponding policy configuration pages.
ACL Resource Policies: SAM	Displays the JSAM and WSAM resource policies that are associated with the specified roles. Includes the type, name, description, action, and resources for each policy. Also displays links you can use to access the corresponding policy configuration pages.
ACL Resource Policies: Telnet	Displays the Telnet/SSH resource policies that are associated with the specified roles. Includes the type, name, description, action, and resources for each policy. Also displays links you can use to access the corresponding policy configuration pages.
ACL Resource Policies: Terminal Services	Displays the Terminal Services resource policies that are associated with the specified roles. Includes the type, name, description, action, and resources for each policy. Also displays links you can use to access the corresponding policy configuration pages.
ACL Resource Policies: VPN Tunneling	Displays the VPN Tunneling resource policies that are associated with the specified roles. Includes the type, name, description, action, and resources for each policy. Also displays links you can use to access the corresponding policy configuration pages.
Resource Profiles: All	Displays the resource profiles that are associated with the specified roles. Includes the type, name, bookmarks, and supporting policies for each profile. Also displays links you can use to access the corresponding resource profile configuration pages.
Resource Profiles: Web Applications	Displays the Web application resource profiles that are associated with the specified roles. Includes the name, bookmarks, and supporting policies for each profile. Also displays links you can use to access the corresponding resource profile configuration pages.
Resource Profiles: Web Hosted Java Applets	Displays the hosted Java applet resource profiles that are associated with the specified roles. Includes the name, bookmarks, and supporting policies for each profile. Also displays links you can use to access the corresponding resource profile configuration pages.
Resource Profiles: Files (Windows)	Displays the Windows file resource profiles that are associated with the specified roles. Includes the name, bookmarks, and supporting policies for each profile. Also displays links you can use to access the corresponding resource profile configuration pages.
Resource Profiles: Files (UNIX)	Displays the UNIX file resource profiles that are associated with the specified roles. Includes the name, bookmarks, and supporting policies for each profile. Also displays links you can use to access the corresponding resource profile configuration pages.
Resource Profiles: SAM Client Applications	Displays the JSAM and WSAM application resource profiles that are associated with the specified roles. Includes the name, bookmarks, and supporting policies for each profile. Also displays links you can use to access the corresponding resource profile configuration pages.
Resource Profiles: SAM WSAM destinations	Displays the WSAM destination resource profiles that are associated with the specified roles. Includes the name, bookmarks, and supporting policies for each profile. Also displays links you can use to access the corresponding resource profile configuration pages.
Resource Profiles: Telnet/SSH	Displays the Telnet/SSH resource profiles that are associated with the specified roles. Includes the name, bookmarks, and supporting policies for each profile. Also displays links you can use to access the corresponding resource profile configuration pages.
Resource Profiles: Terminal Services	Displays the Terminal Services resource profiles that are associated with the specified roles. Includes the name, bookmarks, and supporting policies for each profile. Also displays links you can use to access the corresponding resource profile configuration pages.

CHAPTER 5 Resource Profiles

Resource Profiles

A resource profile contains all of the resource policies, role assignments, and end-user bookmarks required to provide access to an individual resource. Resource profiles simplify resource configuration by consolidating the relevant settings for an individual resource into a single page within the admin console.

The system comes with two types of resource profiles:

- Standard resource profiles enable you to configure settings for a variety of resource types, such as web sites, client/server applications, directory servers, and terminal servers. When you use this method, you choose a profile type that corresponds to your individual resource and then provide details about the resource.

- Resource profile templates enable you to configure settings for specific applications. When you use this method, you choose a specific application (such as the Citrix NFuse version 4.0). Then, the system pre-populates a variety of values for you based on your chosen application and prompts you to configure additional settings as necessary.

Resource profiles are an integral part of the access management framework, and therefore are available on all Connect Secure products. However, you can only access resource profile types that correspond to your licensed features.

To create resource profiles, you:

- Create user roles through the **Users > User Roles** page of the admin console.

- Create resource profiles through the **Users > Resource Profiles** page of the admin console. When creating the resource profile, specify the resource, create autopolicies, associate the profile with user roles, and create bookmarks as necessary.

Related Documentation

- Resource Profile Components
- Resource Profile Templates

Resource Profile Components

Resource profiles contain the following components:

- **Resources**—When you are defining a resource profile, you must specify the individual resource that you want to configure (such as your company Intranet site or a Lotus Notes application). All other major settings within the profile branch from this resource. You can configure a variety of resource types, including web sites, client/server applications, directory servers, and terminal servers.

- **Autopolicies**—When you are defining a resource profile, you generally create autopolicies that establish the access requirements and other settings for the specified resource. The most common type of autopolicy enables access to the primary resource defined in the profile. Other policy types (such as compression and caching autopolicies) "fine-tune" how the system handles the data that it passes to and from the specified resource.

- **Roles**—When you are defining a resource profile, you generally associate the profile with user roles. The specified roles then inherit the autopolicies and (optionally) the bookmarks defined in the resource profile.

- **Bookmarks**—When you are defining a resource profile, you may optionally create a bookmark that links to the profile's primary resource (such as your company intranet's main page). You can also create additional bookmarks that link to various sites within the resource's domain (such as the Sales and Marketing intranet pages). The system displays these bookmarks to users who are assigned to the user roles that you specify.

Figure 17 shows how to configure resources using roles and resource policies. Note that to enable a bookmark for multiple user roles, you must manually re-create the bookmark and enable the appropriate access mechanism for each role. You must also use a variety of pages in the administrator console to create associated resource policies enabling access to the resource and other configuration options.

Figure17: Using Roles and Resource Policies to Configure Resources

General task:

How to accomplish the task using roles & resource policies:

Examples and corresponding administrator UI pages:

Enable access mechanism for role → Manually enable access mechanisms for each role

Enable Web resources for "Sales Role" (Users > User Roles > "Sales Role" > General > Overview page)

Enable Web resources for "Engineering Role" (Users > User Roles > "Sales Role" > General > Overview page)

Create bookmarks & associate with roles → Create & assign bookmarks separately for each role

Create bookmark to Intranet for "Sales Role" (Users > User Roles > "Sales Role" > Web > Bookmarks page)

Create bookmark to Intranet for "Engineering Role" (Users > User Roles > "Sales Role" > Web > Bookmarks page)

Create access & configuration policies → Create separate resource policies for each action.

Create access control policy enabling access to Intranet (Users > Resource Policies > Web > Access > Web ACLs page)

Create SSO policy that passes user credentials to Intranet server (Users > Resource Policies > Web > Remote SSO > Form POST page)

Create compression policy that compresses server content (Users > Resource Policies > Web > Compression page)

Figure 18 shows how to configure resources using resource profiles. Note that you can create a bookmark, associate it with multiple user roles, and create the associated autopolicies enabling access to the resource and other configuration options through a single section in the administrator console. Also note that the system automatically enables the appropriate access mechanism to the roles to which you assign the bookmark.

Figure18: Using Resource Profiles to Configure Resources
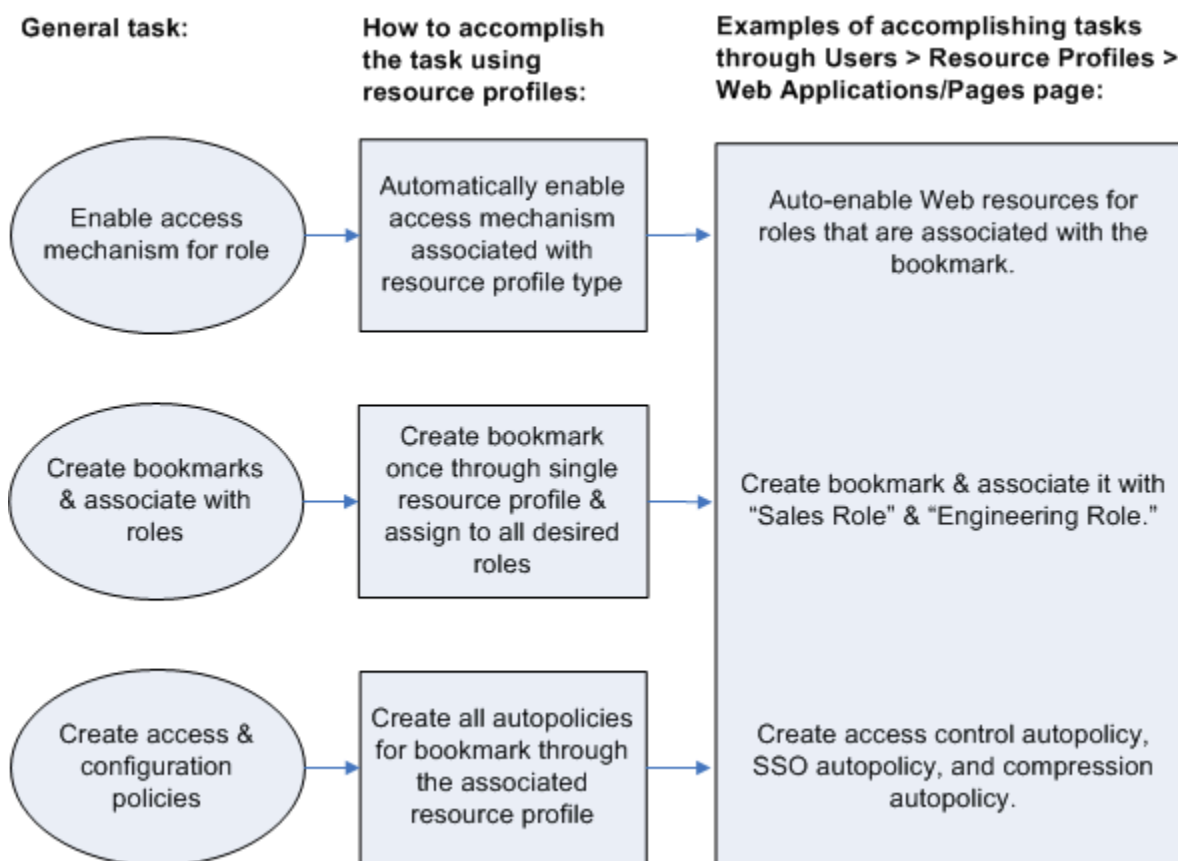
Related Documentation

- Defining a Resource Profile

- Defining Resource Profile Autopolicies

- Defining Resource Profile Roles

- Defining Resource Profile Bookmarks

Defining Resource Profile Resources

When you are defining a resource profile, you must specify the individual resource that you want to configure. Table 8 shows the dependency between the type of profile you choose and the resource you want to configure.

Table8: Resource Profile Types and Configuration Information

Use this type of resource profile	To configure this type of resource
Web application/pages	URLs to Web applications, Web servers, and Web pages; Java applets that are stored on third party servers.
Host Java applet	Java applets that you upload directly to the device.
File browsing	Windows and UNIX/NFS servers, shares, and file paths
SAM client application	Client/server applications
WSAM destination	Destination networks or servers
Telnet/SSH	Telnet or SSH servers
Terminal Services	Windows and Citrix terminal servers

NOTE: You cannot configure applications through VPN Tunneling using resource profiles. Instead, you must use roles and resource policies.

When defining resources, you can use Connect Secure variables, such as <user> to dynamically link users to the correct resources. For instance, you can specify the following Web resource in order to direct users to their own individual intranet pages:

http://yourcompany.intranet/<user>

If the resource field of two different resource profiles are identical and both resource profiles are mapped to the same role, a user might view a resource policy from one profile and a resource policy from the other resource profile. For example, consider the following:

Resource Profile #1:

Resource Profile Name: Intranet

Resource Profile resource: http://intranet.company.com

Resource Profile Web ACL: http://intranet.company.com/sales/*

Mapped to Role: Sales

Resource Profile #2:

Resource Profile Name: Intranet for Sales

Resource Profile resource: http://intranet.company.com

Resource Profile Web ACL: http://intranet.company.com/sales/docs/*

The end user that maps into the Sales role might see a bookmark name Intranet for Sales but the Web ACL enforcement will be http://intranet.company.com/sales/*.

This type of configuration is not supported.

Related Documentation

- Defining a Resource Profile

Defining Resource Profile Autopolicies

When you are defining a resource profile, you generally create autopolicies that establish the access requirements and other settings for the specified resource. The most common type of autopolicy enables access to the primary resource defined in the profile. Other policy types (such as compression and caching autopolicies) "fine-tune" how the system handles the data that it passes to and from the specified resource.

When creating resource profiles, the system only displays those autopolicies that are relevant to the resource profile type. For instance, you may choose to enable access to a client/server application through a WSAM resource profile. When you do, the system displays autopolicies that you can use to enable access to the specified application's server. On the other hand, the system does not display Java access control autopolicies, since Java settings do not apply to WSAM.

NOTE: When defining access policies, you must explicitly list each hostname address. The policy checking system does not append or use the default domain or search domains in the system network settings.

Additionally, the system consolidates all of the relevant autopolicy options in a single page of the user interface, enabling you to understand all of the configuration possibilities and requirements for any given resource type.

NOTE: Access control autopolicies are generally based on the primary resource that you define in the resource profile. If you change the profile's primary resource, however, the system does not necessarily update the corresponding autopolicies. You should re-evaluate your autopolicies after changing the profile's primary resource.

For administrators who are accustomed to using a pre-5.3 version of the Connect Secure product, note that autopolicies are resource policies. The system allows you to sort and order autopolicies along with standard resource policies in the Users > Resource Policies pages of the admin console. However, the system does not allow you to access more detailed configuration options for autopolicies through this section of the admin console. Instead, if you want to change the configuration of an autopolicy, you must access it through the appropriate resource profile.

For administrators who are accustomed to using a pre-5.3 version of the Connect Secure product, note that you can also automatically create resource policies by enabling the Auto-allow option at the role level. However, note that we recommend that you use autopolicies instead, since they directly correspond to the resource you are configuring rather than all resources of a particular type. (You may also choose to enable the Auto-allow option for a role-level feature and create autopolicies for

resources of the same type. When you do, the system creates policies for both and displays them in the appropriate resource policies page of the admin console.)

Related Documentation

- Defining a Resource Profile

- Defining Resource Profile Roles

- Defining Resource Profile Bookmarks

Defining Resource Profile Roles

Within a resource profile, you can assign user roles to the profile. For instance, you might create a resource profile specifying that members of the "Customers" role can access your company's Support Center, while members of the "Evaluators" role cannot. When you assign user roles to a resource profile, the roles inherit all of the autopolicies and bookmarks defined in the resource profile.

Since the resource profile framework does not include options for creating roles, you must create user roles before you can assign them to resource profiles. However, the resource profile framework does include some user role configuration options. For instance, if you assign a user role to a Web resource profile, but you have not enabled Web rewriting for the role, the system automatically enables it for you.

> NOTE: Note that you can assign roles to a resource profile through the role framework as well as the resource profile framework.

Related Documentation

- Defining a Resource Profile

- Defining Resource Profile Autopolicies

- Defining Resource Profile Bookmarks

Defining Resource Profile Bookmarks

When you create a resource profile, the system generally creates a bookmark that links to the profile's primary resource (such as your company intranet's main page). Optionally, you may also create additional bookmarks that link to various sites within the primary resource's domain (such as the Sales and Marketing intranet pages). When you create these bookmarks, you can assign them to user roles, thereby controlling which bookmarks users see when they sign into the end-user console.

> NOTE: WSAM and JSAM resource profiles do not include bookmarks, since the system cannot launch the applications specified in the resource profiles.

For example, you may create a resource profile that controls access to your company intranet. Within the profile, you may specify:

- Resource profile name: Your Intranet

- Primary resource: http://intranet.com

- Web access control autopolicy: Allow access to http://intranet.com:80/*

- Roles: Sales, Engineering

When you create this policy, the system automatically creates a bookmark called "Your Intranet" enabling access to http://intranet.com and displays the bookmark to members of the Sales and Engineering roles.

You may then choose to create the following additional bookmarks to associate with the resource profile:

- "Sales Intranet" bookmark: Creates a link to the http://intranet.com/sales page and displays the link to members of the Sales role.

- "Engineering Intranet" bookmark: Creates a link to the http://intranet.com/engineering page and displays the link to members of the Engineering role.

> NOTE: When configuring bookmarks, note that:
>
> - You can only assign bookmarks to roles that you have already associated with the resource profile—not all of the roles defined on the system. To change the list of roles associated with the resource profile, use settings in its Roles tab.
>
> - Bookmarks simply control which links the system displays to users—not which resources the users can access. For instance, in the example used above, a member of the Sales role would not see a link to the Engineering Intranet page, but he could access it by entering http://intranet.com/engineering his Web browser's address bar. Similarly, if you delete a bookmark, users can still access the resource defined in the profile.
>
> - The system allows you to create multiple bookmarks to the same resource. If you assign duplicate bookmarks to the same user role, however, the system Service only displays one of them to the users.
>
> - Bookmarks link to the primary resource that you define in the resource profile (or a sub-directory of the primary resource). If you change the profile's primary resource, the system updates the corresponding bookmarks accordingly.

Related Documentation

- Defining a Resource Profile

- Defining Resource Profile Autopolicies

- Defining Resource Profile Roles

Resource Profile Templates

Resource profile templates enable you to configure settings for specific applications. When you use this method, you choose a specific application (such as the Citrix NFuse version 4.0). Then, the system pre-populates a variety of values for you based on your chosen application and prompts you to configure additional settings as necessary.

Currently, the system includes templates for the following third-party applications:

- Citrix

- Lotus Notes

- Microsoft Outlook

- Microsoft Sharepoint

- NetBIOS file browsing

Related Documentation

- About Citrix Templates

- Creating Resource Profiles Using the Lotus iNotes Template

- Standard Application Support: MS Outlook

- Creating Resource Profiles Using the Microsoft Sharepoint Template

CHAPTER 6 Virtual Desktop Resource Profiles

Virtual Desktop Resource Profile Overview

In addition to standard resource profiles and resource profile templates, you can configure virtual desktops as resource profiles.

As with the other resource profiles, a virtual desktop profile contains all of the role assignments and end-user bookmarks required to provide access to an individual resource. Unlike other resource profile types, there is no resource policy to configure for virtual desktops due to the dynamic nature of virtual desktops. The IP address and port of the system is not known until the end user launches a session so dynamic ACLs are used.

Icons in the Virtual Desktops section on the end user's home page represent desktops defined by the administrator. Clicking the icon launches the session using the Virtual Desktop Infrastructure (VDI) architecture.

A few of the main features of virtual desktop resource profiles are:

- SSO so that the user can sign on without having to enter their credentials

- Dynamic ACLs

- Client delivery mechanism for end users who do not have the client already installed on their system

- Connection logging

Related Documentation

- Configuring a Citrix XenDesktop Resource Policy
- Configuring a VMware View Manager Resource Profile
- Defining Bookmarks for a Virtual Desktop Profile
- Configuring the Client Delivery
- Connecting to the Servers

Configuring a Citrix XenDesktop Resource Policy

The Citrix XenDesktop manages a pool of virtual desktops hosted on virtual machines and provides the connection management to those desktops. A list of XenDesktops is displayed to the end user as bookmarks. When a desktop is selected, the Citrix client is launched and the user can access that desktop.

To configure a Citrix XenDesktop profile:

1. Select **Users > Resource Profiles > Virtual Desktops**.

2. Click **New Profile**.

3. Select **Citrix XenDesktop** from the Type drop-down list.

4. Enter a name and description (optional) to identify this profile.

5. Enter the name or IP address and port of the connection broker using the format *ip:port*. For example,

   10.10.1.10:80
   xml.example.com:80

   You can enter more than one IP address. Place each address on a separate line.

6. Select the **Use SSL for connecting to the Server** check box if SSL is required to connect to the server.

7. Enter the username to connect to the connection broker or use the <USERNAME> session variable.

8. Enter the password:

   - To use a variable password to connect to the connection broker, select **Variable Password** and enter the variable in the form of <PASSWORD> or <PASSWORD@SEcAuthServer>.

   - Select **Password** to use a static password to connect to the connection broker and enter the user credential's password.

9. Enter the domain where the connection broker is located.

10. Select **Enable Java support** to specify a Java applet to use to associate with the resource profile. The system uses this applet to intermediate traffic or falls back to this applet when ActiveX is not available on the user's system.

11. Click **Save and Continue**.

12. Select the roles to which this profile applies and click **Add**.

    The Enabled Settings table under Users > User Roles also displays which roles have virtual desktops enabled.

13. Click **Save Changes**.

14. (Optional.) In the Bookmarks tab, modify the default bookmark created by the system and/or create new ones.

Related Documentation

Configuring a VMware View Manager Resource Profile

VMware View Manager, formerly VMware VDI, lets you run virtual desktops in a data center that provide end users a single view of all their applications and data in a personalized environment regardless of the device or location they log in from.

To configure a VMware View Manager profile:

1. Select **Users > Resource Profiles > Virtual Desktops**.

2. Click **New Profile**.

3. Select **VMware View Manager** from the Type drop-down list.

4. Enter a name and description (optional) to identify this profile.

5. Enter the name or IP address and port of the connection broker using the format *ip:port*. For example,

    10.10.1.10:80
    xml.example.com:80

   You can enter more than one IP address. Place each address on a separate line.

6. Select the **Use SSL for connecting to the Server** check box if SSL is required to connect to the server.

7. Enter the username to connect to the connection broker or use the <USERNAME> session variable.

8. Enter the password:

   - To use a variable password to connect to the connection broker, select **Variable Password** and enter the variable in the form of <PASSWORD> or <PASSWORD@SEcAuthServer>.

   - Select **Password** to use a static password to connect to the connection broker and enter the user credential's password.

9. Enter the domain where the View Manager server is located.

10. Click **Save and Continue**.

11. Select the roles to which this profile applies and click **Add**.

    The Enabled Settings table under Users > User Roles also displays which roles have virtual desktops enabled.

12. Click **Save Changes**.

13.   (Optional.) In the Bookmarks tab, modify the default bookmark created by the system and/or create new ones.

Related Documentation

- Virtual Desktop Resource Profile Overview

- Configuring a Citrix XenDesktop Resource Policy

- Defining Bookmarks for a Virtual Desktop Profile

- Configuring the Client Delivery

- Connecting to the Servers

Defining Bookmarks for a Virtual Desktop Profile

When you create a virtual desktop resource profile, the system automatically creates a bookmark that links to the server that you specified in the resource profile. The system allows you to modify this bookmark as well as create additional bookmarks to the same server.

These bookmarks are listed in the role bookmark pages (Users > User Roles > Role_Name > Virtual Desktop > Sessions) but you cannot add, modify or delete the bookmarks from the role bookmarks page. Bookmarks can only be added as part of the resource file.

To configure resource profile bookmarks for virtual desktop profiles:

1. Select **Users > Resource Profiles > Virtual Desktop**.

2. Click the name of the virtual desktop profile.

3. Click the Bookmark tab to modify an existing session bookmark. Or, click **New Bookmark** to create an additional session bookmark.

4. (Optional.) Change the name and description of the session bookmark. (By default, the system populates and names the session bookmark using the resource profile name.)

5. Specify whether all desktops or to a selected subset of desktops are available to the user.

   The desktop list is retrieved from the connection broker using the credentials defined in the profile resource page.

6. Enter the credentials used to log in to the actual VMware or XenDesktop machine. The system passes these credentials to the server so that users can sign on without having to manually enter their credentials.

7. Specify how the window should appear to the user during a session by configuring options in the Settings area of the bookmark configuration page.

   (XenDesktop) Under Preferred Client, you can select Automatic Detection, Citrix Client or Java. If you select Automatic Detection, the system checks to see if Citrix Client is present. If it is not present, the end user is given the choice to download the Citrix Client or to use the alternate client, Java ICA Client.

8. Allow users to access local resources such as printers and drives through the terminal session by configuring options in the Connect Devices area of the bookmark configuration page.

(VMware) **Enable MMR**—Redirect certain multimedia codecs running on the remote desktop to the local client for rendering of full-motion video and audio.

(VMware) **Allow Desktop Reset**—Allow users to reset their desktop without administrative assistance. For example, if the desktop hangs, there is currently no way for the user to perform a hard reboot of the desktop. This option allows the users to restart their own virtual desktops thereby reducing the dependency on the administrator or helpdesk.

9. Specify how the terminal emulation window should appear to the user during a terminal session by configuring options in the Desktop Settings area.

10. Specify the roles to which you want to display the session bookmarks if you are configuring the session bookmark through the resource profile pages, under Roles:

- **ALL selected roles**—Displays the session bookmark to all of the roles associated with the resource profile.

- **Subset of selected roles**—Displays the session bookmark to a subset of the roles associated with the resource profile. Then select roles from the ALL Selected Roles list and click Add to move them to the Subset of selected roles list.

11. Click **Save Changes**.

Related Documentation

- Defining Display Options for the Windows Terminal Services Session

- Virtual Desktop Resource Profile Overview

- Configuring a Citrix XenDesktop Resource Policy

- Configuring a VMware View Manager Resource Profile

- Configuring the Client Delivery

- Connecting to the Servers

Configuring the Client Delivery

You can use the Virtual Desktop Configuration page to define the client delivery mechanism for end-users who do not have the client. The process is similar for both Citrix XenDesktop and VMware View Manager.

1. Choose **System > Configuration> Virtual Desktops**. For Citrix XenDesktop, select **Citrix**.

2. Select **Download from Pulse Connect Secure** to download the client file from the system. Click Browse to locate the client file (.msi, .exe or .cab) and enter the version number.

3. Select **Download from a URL** to download the client file from the Internet. If desired, enter a new URL to override the default.

4. Check the **Access the URL through the Secure Gateway** check box if end users cannot directly access the specified Web page. Selecting this option allows users to use the secure gateway to access the URL.

5. Under Server Connection Timeout, enter the number of seconds to wait for the server to respond before timing out.

Related Documentation

- Virtual Desktop Resource Profile Overview

- Connecting to the Servers

Connecting to the Servers

When an end user clicks a desktop icon, the system passes credentials to the server based on the desktop profile.

For XenDesktop, the system authenticates to the Citrix DDC server using credentials defined in the desktop profile. If successful, the list of available desktops is returned by the DDC server and is represented as bookmarks to the end user. When an end user clicks a XenDesktop icon, the system retrieves the ICA from the XenDesktop server and presents a desktop session to the user.

When an end user clicks a VMware View Manager icon, the system authenticates to the View Manager using credentials defined in the desktop profile. If authentication is successful, a JSESSIONID cookie is returned by the View Manager, the system creates a tunnel using the cookie for the duration of the session.

If the desktop is unavailable, the client will continue to try to connect until the desktop is available or until a predefined timeout period occurs. An error message lets the user know the status, either that the system is retrying the connection or that the desktop is unavailable. Similarly if the desktop is already in use by another enduser, an error message is presented to the user.

User logs are updated to show which VM machines are assigned to each user. Username, realm, VM IP, port, connection type, pool and connection broker are logged with each message.

The Active Virtual Desktops Sessions page (System > Status > Virtual Desktop Sessions) lists the active connections, including the connection broker, the VM machine assigned to the user and the connection type.

Related Documentation

- Virtual Desktop Resource Profile Overview

- Configuring the Client Delivery

CHAPTER 7 Resource Policies

Resource Policies

A resource policy is a system rule that specifies resources and actions for a particular access feature. A resource is either a server or file that can be accessed through the system, and an action is to "allow" or "deny" a resource or to perform or not perform a function. Each access feature has one or more types of policies, which determine the system's response to a user request or how to enable an access feature (in the case of Email Client). You may also define detailed rules for a resource policy, which enable you to evaluate additional requirements for specific user requests.

You can create the following types of resource policies through the Resource Policies pages:

- **Web Resource Policies**—The Web resource policies specify the Web resources to which users may or may not browse. They also contain additional specifications such as header caching requirements, servers to which java applets can connect, code-signing certificates that the system should use to sign java applets, resources that the system should and should not rewrite, applications for which the system performs minimal intermediation, and single sign-on options.

- **File Resource Policies**—The file resource policies specify the Windows, UNIX, and NFS file resources to which users may or may not browse. hey also contain additional specifications such as file resources for which users need to provide additional credentials.

- **Secure Application Manager Resource Policies**—The Secure Application Manager resource policies allow or deny access to applications configured to use JSAM or WSAM to make socket connections.

- **Telnet/SSH Resource Policies**—The Telnet/SSH resource policies allow or deny access to the specified servers.

- **Terminal Services Policies**—The Terminal Services resource policies allow or deny access to the specified Windows servers or Citrix Metaframe servers.

- **VPN Tunneling Resource Policies**—The VPN Tunneling resource policies allow or deny access to the specified servers and specify IP address pools.

- **Secure Email Client Resource Policies**—The Secure Email Client access resource policy allows you to enable or disable e-mail client support. To allow end-users to open and save e-mail attachments of different document types in OWA and iNotes, select the OWA or iNotes type when defining a Web Application Resource Profile.

> **NOTE:** You can also create resource policies as part of the resource profile configuration process. In this case, the resource policies are called "advanced policies."

Resource policies are an integral part of the access management framework, and therefore are available on all Connect Secure products. However, you can access only resource policy types that correspond to your licensed features.

Related Documentation

- Resource Policy Components
- Resource Policy Evaluation
- Creating Detailed Rules for Resource Policies
- Customizing Resource Policy UI Views

Resource Policy Components

A resource policy contains the following information:

- **Resources**—A collection of resource names (URLs, hostnames, or IP address/netmask combinations) that specifies the resources to which the policy applies. You can specify a resource using a wildcard prefix to match hostnames. The default resource for a policy is star (*), meaning that the policy applies to all related resources.

- **Roles**—An optional list of user roles to which this policy applies. The default setting is to apply the policy to all roles.

- **Action**—The action for the system to take when a user requests the resource corresponding to the Resource list. An action may specify to allow or deny a resource or to perform or not perform an action, such as to rewrite Web content or allow Java socket connections.

- **Detailed Rules**—An optional list of elements that specifies resource details (such as a specific URL, directory path, file, or file type) to which you want to apply a different action or for which you want to evaluate conditions before applying the action. You can define one or more rules and specify the order in which the system evaluates them.

Specifying Resources for a Resource Policy

The system platform's engine that evaluates resource policies requires that the resources listed in a policy's Resources list follow a canonical format. This section describes the canonical formats available for specifying Web, file, and server resources. When a user tries to access a specific resource, the system compares the requested resource to the resources specified in the corresponding policies, starting with the first policy in a policy list. When the engine matches a requested resource to a resource specified in a policy's Resources list, it then evaluates further policy constraints and returns the appropriate action to the appliance (no further policies

are evaluated). If no policy applies, then the appliance evaluates the auto-allow bookmarks (if defined); otherwise the default action for the policy is returned.

> **NOTE:** You may not see the auto-allow option if you are using a new installation, if you use resource profiles rather than resource policies, or if an administrator has hidden the option.

General Notes About the Canonical Formats

Please note the following when using canonical formats:

- If a path component ends with forward-slash_star (/*), then it matches the leaf node and everything below. If the path component ends with forward-slash_ percent (/%), then it matches the leaf node and everything one-level below only. For example:

  /intranet/*matches:

  ```
  /intranet
  /intranet/home.html
  /intranet/elee/public/index.html
  ```

  /intranet/% matches:

  ```
  /intranet
  /intranet/home.html
  but NOT /intranet/elee/public/index.html
  ```

- A resource's hostname and IP address are passed to the policy engine at the same time. If a server in a policy's Resources list is specified as an IP address, then the evaluation is based on the IP address. Otherwise, the engine tries to match the two hostnames—it does not perform a reverse-DNS-lookup to determine the IP.

  > **NOTE:** You cannot specify a hostname for a VPN Tunneling resource policy. You can only specify an IP address.

- If a hostname is not fully qualified in the hosts file, such as "pulsesecure" instead of "intranet. pulsesecure.net", and you are accessing a hostname using the short name, then the engine performs the resource matching against the short name. If, however, the short name is not in the hosts file and the hostname resolution is done by DNS (by adding the domains listed in the Networks configuration page), then the fully qualified domain name (FQDN) is used for resource matching. In other words, for web resource policies a DNS lookup of the short name is performed. The result of the DNS lookup is a FQDN; the engine matches the FQDN with the ones entered in the UI.

Specifying Server Resources

When specifying server resources for Telnet/SSH, Terminal Services, or VPN Tunneling resource policies, note the following guidelines.

The canonical format is **[protocol://] host [:ports]**

The components are:

- Protocol (optional)—Possible case-insensitive values:

- tcp

- udp

- icmp

If the protocol is missing, then all protocols are assumed. If a protocol is specified, then the delimiter ":// " is required. No special characters are allowed.

> ℹ️ NOTE: Available only to VPN Tunneling policies. For other access feature resource policies, such as Secure Application Manager and Telnet/SSH, it is invalid to specify this component.

- Host (required)—Possible values:

  - IP address/Netmask—The IP address needs to be in the format: *a.b.c.d*

    The netmask may be in one of two formats:

    - Prefix: High order bits

    - IP: *a.b.c.d*

    For example: 10.11.149.2/24 or 10.11.149.2/255.255.255.0

    No special characters are allowed.

  - DNS Hostname—For example: www.juniper.com

    Table 9 shows the special characters allowed.

Table9: DNS Hostname Special Characters

*	Matches ALL characters
%	Matches any character except dot (.)
?	Matches exactly one character

> ℹ️ NOTE: You cannot specify a hostname for a VPN Tunneling resource policy. You can only specify an IP address.

  - Ports (optional)—Possible values are shown in Table 10.

Table10: Port Possible Values

*	Matches ALL ports; no other special characters are allowed
port[,port]*	A comma-delimited list of single ports. Valid port numbers are [1- 65535]. Do not enter a space between port numbers. You can specify up to 15 ports.
[port1]-[port2]	A range of ports, from port1 to port2, inclusive.

ℹ️ NOTE: You may mix port lists and port ranges, such as: 80,443,8080-8090, except for in VPN Tunneling where mixing of port lists and port ranges is not supported.

If the port is missing, then the default port 80 is assigned for http, 443 for https. For VPN Tunneling, if the port is missing then the default port http is *. If a port is specified, then the delimiter ":" is required. For example:

```
<username>.danastreet.net:5901-5910
tcp://10.10.149.149:22,23
tcp://10.11.0.10:80
udp://10.11.0.10:*
```

Resource Policy Evaluation

When the system receives a user request, it evaluates the resource policies corresponding to the type of request. When it processes the policy that corresponds to the requested resource, it applies the specified action to the request. This action is defined on the policy's General tab or Detailed Rules tab. For example, if a user requests a Web page, the system knows to use the Web resource policies. In the case of Web requests, the system always starts with the Web Rewriting policies (Selective Rewriting and Pass through Proxy) to determine whether or not to handle the request. If none of these policies applies (or none is defined), the system then evaluates the Web Access policies until it finds one that pertains to the requested resource.

The system evaluates a set of resource policies for an access feature from the top down, meaning that it starts with the policy numbered one and then continues down the policy list until it finds a matching policy. If you defined detailed rules for the matching policy, the system evaluates the rules from the top down, starting with the rule numbered one and stopping when it finds a matching resource in the rule's Resource list. Figure 19 illustrates the general steps of policy evaluation:

Figure19: Resource Policy Evaluation Steps
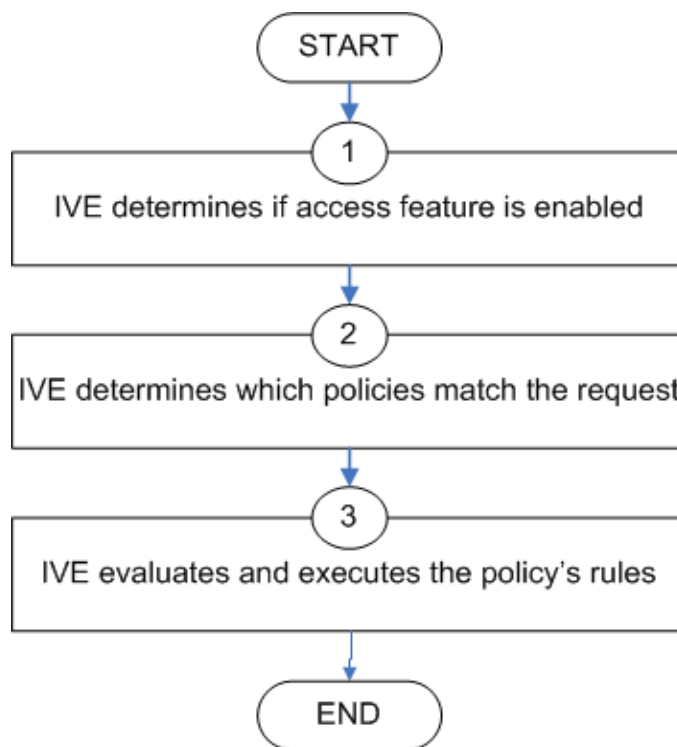
Details regarding each evaluation step:

1. The system receives a user request and evaluates the user's session role to determine if the corresponding access feature is enabled. A user's "session role" is based on either the role or roles to which the user is assigned during the authentication process. The access features enabled for a user are determined by an authentication realm's role mapping configuration.

2. The system determines which policies match the request. The system evaluates the resource policies related to the user request, sequentially processing each policy until finding the one whose resource list and designated roles match the request. (If you configure the system using resource profiles, the system evaluates the advanced policies that you configure as part of the resource profile.)

   The Web and file access features have more than one type of policy, so the system first determines the type of request (such as to a Web page, Java applet, or UNIX file) and then evaluates the policies related to the request. In the case of the Web access feature, the Rewriting policies are evaluated first for every Web request. The remaining five access features—Secure Application Manager, Secure Terminal Access, and Secure Email Client—have only one resource policy.

3. The system evaluates and executes the rules specified in the matching policies. You can configure policy rules to do two things:

   - Specify resources to which an action applies at a more granular level. For example, if you specify a Web server in the main policy settings for a Web Access resource policy, you

can define a detailed rule that specifies a particular path on this server and then change the action for this path.

- Require the user to meet specific conditions written as boolean expressions or custom expressions in order to apply the action.

4. The system stops processing resource policies as soon as the requested resource is found in a policy's Resource list or detailed rule.

> NOTE: If you use automatic (time-based) dynamic policy evaluation or you perform a manual policy evaluation, the system repeats the resource evaluation process described in this section.

Related Documentation

- User Roles Overview
- Creating Detailed Rules for Resource Policies
- Dynamic Policy Evaluation

Creating Detailed Rules for Resource Policies

The Web, file, Secure Application Manager, Telnet/SSH, and VPN Tunneling access features enable you to specify resource policies for individual Web, file, application, and telnet servers. The Email Client access features have one policy that applies globally. For this policy, you specify server settings that are used for every role that enables these access features. For all other access features, you can specify any number of resource polices, and for each, you can define one or more detailed rules.

A detailed rule is a an extension of a resource policy that may specify:

- Additional resource information—such as a specific path, directory, file, or file type—for resources listed on the General tab. Note that you may also specify the same resource list (as on the General tab) for a detailed rule if the only purpose of the detailed rule is to apply conditions to a user request.

- An action different from that specified on the General tab (although the options are the same).

- Conditions that must be true in order for the detailed rule to apply.

In many cases, the base resource policy—that is, the information specified on the General tab of a resource policy—provides sufficient access control for a resource:

If a user belonging to the (defined_roles) tries to access the (defined_resources), DO the specified (resource_action).

You may want to define one or more detailed rules for a policy when you want perform an action based on a combination of other information, which can include:

- A resource's properties, such as its header, content-type, or file type.

- A user's properties, such as the user's username and roles to which the user maps.

- A session's properties, such as the user's source IP or browser type, whether the user is running Host Checker or Cache Cleaner, the time of day, and certificate attributes.

Detailed rules add flexibility to resource access control by enabling you to leverage existing resource and permission information to specify different requirements for different users to whom the base resource policy applies.

Related Documentation

- Writing a Detailed Rule for Resource Policies

Writing a Detailed Rule for Resource Policies

Detailed rules add flexibility to resource access control by enabling you to leverage existing resource and permission information to specify different requirements for different users to whom the base resource policy applies.

To write a detailed rule for a resource policy:

1. On the New Policy page for a resource policy, enter the required resource and role information.

2. In the Action section, select **Use Detailed Rules** and then click **Save Changes**.

3. On the Detailed Rules tab, click **New Rule**.

4. On the Detailed Rule page:

   - In the Action section, specify:

     - **Disable SSO**—The system disables automatic SSO authentication for this user role and, instead, prompts the user for sign-in credentials.

     - **Basic Auth**—This option specifies that the system use the Basic Authentication Intermediation method to control SSO behavior.

       Enable Intermediation—Select the credentials to use. If this pull-down menu is blank, no basic authentication SSO settings are defined in the SSO General tab.

       Disable Intermediation—When you select this option, the system does not intermediate the challenge/response sequence.

       > ℹ️ NOTE: The system always intermediates requests to Web proxies that require basic authentication, even if you select Disable Intermediation.
       >
       > Although you are given an option to disable basic authentication intermediation, we do not recommend this option, as it is a very insecure authentication method and, in some cases, can transmit user credentials over the network in clear (unencrypted) text.

     - **NTLM**—This option specifies that the system use the Microsoft NTLM Intermediation method to control SSO behavior.

       Select the credentials to use. If this pull-down menu is blank, no NTLM SSO settings are defined in the SSO General tab.

Select the **Fallback to NTLM V1** option to try both NTLM V1 and NTLM V2. If you do not select this option, the system falls back only to NTLM V2. An intermediation page appear if SSO fails.

- **Kerberos**—This option specifies that the system use the Kerberos Intermediation method to control SSO behavior.

  Select the credentials to use. If this pull-down menu is blank, no kerberos SSO settings are defined in the SSO General tab.

  Select the **Fallback to NTLM V2** option to fallback only to NTLM V2 if kerberos fails. If you do not select this option, a Kerberos intermediation page appears if Kerberos SSO fails.

- **Constrained Delegation**—This option specifies that the system use the constrained delegation intermediation method to control SSO behavior.

  Select the credentials to use. If this pull-down menu is blank, no constrained delegation SSO settings are defined in the SSO General tab.

  Select the **Fallback to Kerberos** option fallback to Kerberos if constrained delegation fails. If you select this option, an intermediation page appears if constrained delegation fails. If you do not select this option and constrained delegation fails, an error page appears.

- In the Resources section, specify any of the following (required):

  - The same or a partial list of the resources specified on the General tab.

  - A specific path or file on the server(s) specified on the General tab, using wildcards when appropriate. For information about how to use wildcards within a Resources list, see the documentation for the corresponding resource policy.

  - A file type, preceded by a path if appropriate or just specify */*.file_extension to indicate files with the specified extension within any path on the server(s) specified on the General tab.

- In the Conditions section, specify one or more expressions to evaluate in order to perform the action (optional):

  - Boolean expressions: Using system variables, write one or more boolean expressions using the NOT, OR, or AND operators.

  - Custom expressions: Using the custom expression syntax, write one or more custom expressions.

    > *i* NOTE: You can use the <USER> substitution variable in ACLs for web pages, telnet, files, and SAM. You cannot use the variable in VPN Tunneling ACLs.
    >
    > When specifying a time condition, the specified time range cannot cross midnight. The workaround is to break the time range into two conditions.

- Click **Save Changes**.

5. On the Detailed Rules tab, order the rules according to how you want the system to evaluate them. Keep in mind that once the system matches the resource requested by the user to a

resource in a rule's Resource list, it performs the specified action and stops processing rules (and other resource policies).

Related Documentation

- Writing the Basic, NTLM and Kerberos Resources

- Using Custom Expressions in Rule Configuration

Customizing Resource Policy UI Views

You can limit which resource policies the system displays on any given resource policy page based on user roles. For instance, you can configure the Users > Resource Policies > Web page of the admin console to only display those resource policies that are assigned to the "Sales" user role.

To control which resource policies the system displays:

1. Navigate to **Users > Resource Policies >***Policy Type*.

2. From the Show all policies that apply to list, select **All Roles** or an individual role.

3. Click **Update**. The system displays resource policies that are assigned to the selected roles.

CHAPTER 8 Authentication and Directory Servers

AAA Server Overview

This topic includes the following information:

Understanding the Role of AAA Servers in the Pulse Secure access management framework

AAA stands for authentication, authorization, and accounting. A AAA server is a database that stores user credentials—username and password—and, in some cases, group information or other user attributes. The authentication results and group or user attribute information are used Pulse Secure access management framework policy decisions.

In the Pulse Secure access management framework, the sign-in page, realm, and AAA server configurations are associated. They determine user access and user role. A user submits credentials through a sign-in page, which specifies a realm, which is associated with a AAA server. If the access request meets the realm's authentication policy, the system forwards the user's credentials to the associated authentication server. The authentication server's job is to verify the user's identity. After verifying the user, the authentication server sends approval. If the realm also uses the server as a directory/attribute server, the AAA server sends the user's group information or other user attribute information. The access management framework then evaluates the realm's role-mapping rules to determine the user roles that apply to the session.

The Pulse Secure access management framework supports the following types of AAA servers:

- Local—You can create special purpose local databases to manually create user accounts, manage guest user access, permit anonymous access, or manage access based on digital certificates or MAC addresses.

- External (standards-based)—You can integrate standards-based LDAP and RADIUS servers with the access management framework. In addition to using the backend server for authentication, you can use LDAP group and RADIUS attribute information in role-mapping rules.

- External (other)—You can integrate compatible versions of popular third-party AAA servers with the access management framework. In addition to using the backend server for authentication, you can use Active Directory group information and SiteMinder attributes in role-mapping rules. In addition, you can use MDM device attributes in role mapping rules.

Table 11 is a reference of the AAA servers supported in Policy Secure and Connect Secure deployments.

Table11: Supported AAA Servers

	Policy Secure	Connect Secure
Local	*Local Authentication Server*[*], "Anonymous Server", "Certificate Server", *MAC Address Authentication Server* [*]Special features to manage guest users.	"Local Authentication Server"[**], "Anonymous Server", "Certificate Server", "SAML Server"[***] [**]No special features to manage guest users. [***]Supports an authentication server configuration when deployed as a SAML service provider. Different Connect Secure features support a local SAML server when deployed as a SAML identity provider.
External (standards-based)	"LDAP Server", "RADIUS Server"	"LDAP Server", "RADIUS Server"
External (other)	"Active Directory", "MDM Server", "NIS Server", "RSA ACE Server", "SiteMinder Server", *SQL Server*	"Active Directory", "MDM Server", "NIS Server", "RSA ACE Server", "SiteMinder Server"

Authentication Protocols Used by AAA Servers

Policy Secure supports multiple authentication protocols. The following authentication servers require the protocols listed:

- Local authentication servers—If the passwords are stored as hashed values, the protocols available are PAP and MS-CHAP v1 with or without EAP. If the passwords are stored as cleartext, CHAP and MD5-Challenge are also available.

- Active directory—The protocols available for inner authentication are PAP, MSCHAP and MS-CHAP v2, with or without EAP.

- LDAP—CHAP, EAP-MD5-Challenge, MS-CHAP v1, and MS-CHAP v2 protocols can be used with an LDAP authentication server only if the administration password is in cleartext. By default, challenge response protocols are disabled for LDAP servers. Use these protocols only with noninteractive devices (for example, phones), as password management is not possible if these protocols are used for authentication.

- Anonymous authentication server—The anonymous authentication server is not supported with the open protocols.

AAA Server Configuration Task Summary

To integrate an authentication server:

1. Configure the authentication server. Select **Authentication > Authentication > Auth. Servers** page and complete the authentication server configuration.

2. Create an authentication realm. Select **Users > User Realms** or **Administrators > Admin Realms** and select the authentication server when you complete the authentication realm configuration.

Related Documentation

- *Understanding Authentication Realms*

- Understanding Authentication Realms

Using the Local Authentication Server

This topic describes the local authentication server. It includes the following information:

- Local Authentication Server Overview
- Configuring the Local Authentication Server
- Creating User Accounts
- Managing User Accounts
- Creating Administrator User Accounts
- Using the Admin User Sign-In Page to Manage the Local Authentication Users Table

Local Authentication Server Overview

This section provides an overview of the feature and its limitations. It includes the following sections:

- Understanding the Local Authentication Server

Understanding the Local Authentication Server

The local authentication server is an authentication database that is built in to the system. Therefore, it is considered a "local" server in contrast to a third-party enterprise AAA server that is connected over the network.

Typically, you create local user accounts for temporary users who do not have accounts on your enterprise AAA servers. Temporary users include lab users or guests, but you might find the local authentication server useful to create temporary accounts for users who are normally verified by an enterprise AAA server that you plan to disable.

You also use the local authentication server to create accounts for administrator users, such as system administrators.

NOTE: Although it is common practice to use the local authentication server for administrator accounts, it does not preclude you from using any of the supported third-party enterprise AAA servers in your administrator access management framework.

Configuring the Local Authentication Server

You can create multiple local authentication server instances. When you define a new local authentication server, you must give the server a unique name and configure options for passwords and guest users.

To create a local authentication server:

1. Select **Authentication > Auth. Servers**.

2. Select **Local Authentication** and click **New Server** to display the configuration page.

3. Complete the configuration as described in Table 12.

4. Save the configuration.

Table12: Local Authentication Server Configuration Guidelines

Settings	Guidelines
Name	Specify a name that is useful to you.
Password Options	
Minimum length	Specify a number of characters. The valid range is 0-99. 6 is the default.
Maximum length	Specify a number of characters. The valid range is 0-99. 8 is the default. The maximum length cannot be less than the minimum length.
Minimum digits	Specify the number of digits required in a password. Do not require more digits than the value of the maximum length option.

Minimum letters	Specify the number of letters required in a password. Do not require more letters than the value of the maximum length option. If you enable the previous option, the combined total of the two options cannot exceed that of the value specified in the maximum length option.
Uppercase and lowercase required	Select this option if you want all passwords to contain a mixture of uppercase and lowercase letters. NOTE: Require passwords to contain at least two letters if you also require a mix of uppercase and lowercase letters.
Different from username	Select this option if the password cannot equal the username.
Different from previous password	Select this option if a new password cannot equal the previous password.
Stored as cleartext	Select this option if you are using open authentication protocol sets. CHAP and EAP-MD5-Challenge work with local authentication servers only if you select this option. NOTE: Be aware of the security implications of storing passwords as cleartext.
Password Management	
Allow users to change passwords	Select this option if you want users to be able to change their passwords. NOTE: In addition to selecting local authentication password management options, you must select the **Enable Password Management** option for the associated realm authentication policy.
Force password change	Select this option to specify the number of days after which a password expires. The default is 64 days.
Prompt users to change password	Select this option to specify when to prompt the user to change passwords.

Creating User Accounts

You use the Users page to create local authentication server user accounts. A user account includes a username and password to be used for authentication, as well as other information used for records and account management.

To create a local user account:

1. Select **Authentication > Auth. Servers**.

2. Select the local authentication server to which you want to add a user account.

3. Click the **Users** tab.

4. Click **New** to display the configuration page.

5. Complete the configuration as described in Table 13.

6. Save the configuration.

Table13: User Account Configuration Guidelines

Settings	Guidelines
Username	Do not include "~~" in a username. NOTE: You cannot change a username after you create the account.

Full Name	Specify the user's full name.
Password	Specify a password. Make sure that the password you enter conforms to the password options specified on the local authentication server configuration page.
Confirm password	Confirm the password.
One-time use	Select this option to limit the user to one login. After one successful login, the user's login state is set to disabled, and the user receives an error message when attempting subsequent sign-ins. However, you can manually reset this option to allow the same user to log in again. If you do not select this option, the user account is subject to the specified start and end time for the account.
Enabled	Select this check box if it is not already selected. If the one-time use option has been implemented, this option is listed as **Disabled** after the user has logged in successfully. If a permanent or one-time user is logged in and you disable this option, the user is immediately logged out of the system and receives an error message.
Require user to change password	Select this option to force users to change their passwords at the next login. NOTE: If you force the user to change passwords, you must also enable the local authentication password management options.

Managing User Accounts

You use the Users page to list, modify, and delete local authentication server user accounts.

To manage a user account:

1. Select **Authentication > Auth. Servers**.

2. Click the link for the authentication server you want to manage.

3. Click the **Users** tab to display the user accounts table.

   The user accounts table includes entries for the accounts that have been created. The Last Sign-in Statistic column shows the last successful sign-in date and time for each user, the user's IP address, and the agent or browser type and version.

4. Use the controls to search for users and manage user accounts:

   - To search for a specific user, enter a username in the Show users named box and click **Update**.

     > TIP: You can use an asterisk (*) as a wildcard, where * represents any number of zero or more characters. For example, to search for all usernames that contain the letters jo, enter ***jo***. The search is case-sensitive. To display the entire list of accounts again, type ***** or delete the field's contents and click **Update**.

   - To limit the number of users displayed on the page, enter a number in the Show *N* users box and click **Update**.

   - To edit the user account configuration, click the link in the Username column to display the Update Local User Account page.

   - To terminate the user session and delete the account, select the box next to the user account record and click **Delete**.

Select a user to display the user account configuration page. You can use this page to modify the account settings, or to disable or quarantine the account.

Creating Administrator User Accounts

You use the Admin Users page to create a special admin user account that enables the account holder to manage the local authentication server users table. These special admin users can sign in to a special page that enables them to create, modify, and delete user accounts.

To create a special admin user account:

1. Select **Authentication > Auth. Servers > System Local**.

2. Click the **Admin Users** tab to display the configuration page.

3. Specify a username, select an authentication realm, and click **Add** to create the administrator user.

4. Save the configuration.

Using the Admin User Sign-In Page to Manage the Local Authentication Users Table

The special admin users created using the feature shown in the previous section can manage the local authentication server accounts table. For example, if an admin user named adminuser is provisioned to manage user accounts for the Users realm, when adminuser signs into the Users realm sign-in page, a User Admin button appears on the toolbar at the top of the page. Figure 20 shows the toolbar.

Figure20: Sign-In Page Toolbar
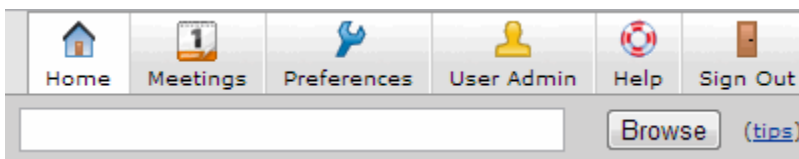
The special admin user can click the **User Admin** button to display the User Admin page, which shows the local authentication server users table. Figure 21

Figure21: User Admin Page
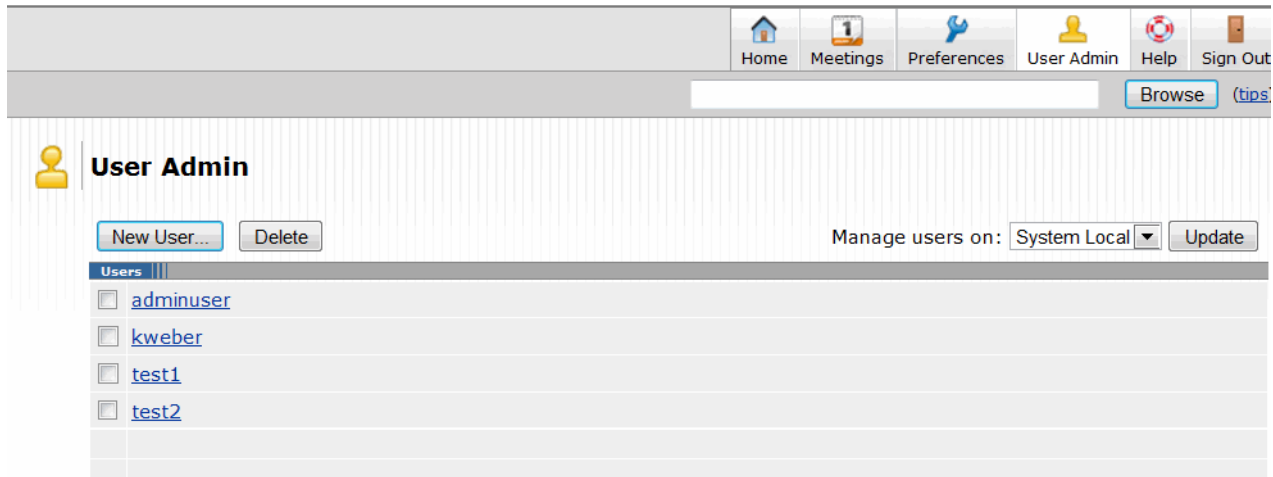
The special admin user can select accounts and delete them, and can create user accounts. The same account management guidelines apply as when using the admin console for creating and modifying user records.

Figure 22 shows the New Local User configuration page.

Figure22: New Local User Configuration Page

Related Documentation

- AAA Server Overview

- *Specifying Password Access Restrictions*

- Specifying Password Access Restrictions

Using Active Directory

This topic describes integration with the Microsoft® Windows® platform Active Directory™ service. It includes the following information:

- Microsoft Windows Platform Active Directory Service Overview
- Configuring Authentication and Authorization with Active Directory Service (Standard Mode)
- Configuring Authentication and Authorization with Active Directory Service (Legacy Mode)
- Displaying the User Accounts Table
- Troubleshooting the Active Directory Service (Standard Mode)

Microsoft Windows Platform Active Directory Service Overview

This section describes support for using Policy Secure and Connect Secure with the Active Directory AAA service. It includes the following sections:

- Understanding Active Directory
- Active Directory Feature Support
- Interoperability Requirements and Limitations
- Understanding the Active Directory Standard Configuration and Active Directory Legacy Mode Configuration

Understanding Active Directory

Active Directory is a directory service used in Windows domain networks. It is included in most Windows server operating systems. Enterprise servers that run Active Directory are called domain controllers. An Active Directory domain controller authenticates and authorizes users and computers in a Windows domain network.

When you use Active Directory as the authentication and authorization service for your Pulse Secure access management framework, users can sign in to Policy Secure and Connect Secure using the same username and password they use to access their Windows desktops. You can also use Active Directory group information in role mapping rules.

Active Directory Feature Support

Pulse Secure access management framework supports the following Active Directory features:

- Honors trust relationships in Active Directory and Windows NT environments.
- Supports Domain Local Groups, Domain Global Groups, and Universal Groups defined in the Active Directory forest.
- Supports use of Kerberos, NTLMv2, and NTMLv1 authentication protocols.
- Supports user principal name (UPN) format for usernames. This support is available for Web login, Policy Secure at Layer 3 (JUAC), and EAP-MS-CHAP v2.

Interoperability Requirements and Limitations

The following limitations apply to interoperability with Active Directory:

- The Pulse Secure access management framework uses Active Directory security groups, not distribution groups. Security groups allow you to use one type of group for not only assigning rights and permissions, but also as a distribution list for e-mail.

- Each Active Directory configuration you create for the Pulse Secure access management framework should use a different and unique machine account name.

- If the current Active Directory domain controller is not reachable, the user or machine authentication requests fail for a few seconds (less than 2 minutes) before attempting to authenticate users with another domain controller in the Active Directory domain.

- We do not support Active Directory implementations that use the equal sign operator (=) in a group name, such as: "\=THIRD FLOOR GROUP". The Pulse Secure access management framework authentication process involves search operations that use the equal sign operator (=) when parsing server catalogs to retrieve group names, usernames and domain names, as well as user_SID and domain_SID values. You might encounter unexpected behavior that can affect normal processing of authentication services if a group name configured on your Active Directory server includes an equal sign operator (=).

- Active Directory versions Windows 2008 R2 and later use a dynamic port range. The default start port is 49152 and the default end port is 65535. Therefore, if there is a firewall between the Pulse Secure client service and the Active Directory Service, you must increase the remote procedure call (RPC) port range on the firewall. See .Microsoft Knowledge Base article 929851.

- The Pulse Secure password management feature, which enables users to change their Active Directory passwords through the Pulse Secure service Web server, is not supported for users of trusted domains that do not trust the domain specified in the Pulse Active Directory configuration.

- UPN format for user login is not supported for MS-CHAP v2.

Understanding the Active Directory Standard Configuration and Active Directory Legacy Mode Configuration

Policy Secure and Connect Secure supports interoperability with Active Directory using two configurations:

- Active Directory standard configuration. Recommended. Supports interoperability with any version of Active Directory, and is the required configuration mode to support authentication using MS-CHAP v2 with Windows 2008 R2 Active Directory Service. Machine authentication, for example, uses MS-CHAP v2.

- Active Directory Legacy Mode configuration. Supports interoperability with Active Directory versions Microsoft 2003 or earlier. You might choose to use the Active Directory Legacy Mode configuration as your primary configuration if you require role-mapping rules to use "domain local groups" of trusted child domains. The Active Directory standard configuration does not support "domain local groups" for authorization. You can also use an LDAP server as the directory server if you want to use domain local groups in role mapping.

We recommend configuring and using Active Directory standard configuration as your primary configuration and Active Directory Legacy Mode as a fallback configuration. If you are

experiencing issues with the standard Active Directory configuration, you can use the **Switch to Active Directory Legacy Mode** at the bottom of the Active Directory authentication server configuration page to activate the Active Directory Legacy Mode configuration. The fallback configuration saves you from having to reconfigure your authentication server, server catalog, and role mapping rules.

Configuring Authentication and Authorization with Active Directory Service (Standard Mode)

To configure integration with Active Directory Service (standard mode):

1. Select **Authentication > Auth. Servers**.

2. Select **Active Directory / Windows NT** and click **New Server** to display the configuration page.

3. Select **Active Directory** mode and complete the configuration as described in Table 14.

4. Save the configuration.

Table14: Active Directory Mode Settings

Settings	Guidelines
Mode	
	Select one of the following modes:
	• **Active Directory**—For recent versions of Windows Server.
	• **Active Directory Legacy Mode**—For Windows Server 2003 and earlier.
	This table describes Active Directory mode.
Base Configuration	
Name	Specify a name to identify the server within the system.
Domain	Specify the NetBIOS domain name for the Active Directory domain.
	The system uses DNS to discover domain controllers in the Active Directory forest. It sends authentication requests to the domain controller at the closest site. Ensure that your DNS servers are configured to resolve the Active Directory domain controller fully qualified domain name (FQDN) and service (SRV) records.
	Do not create an Active Directory and an Active Directory Legacy Mode configuration with the same domain and computer name. They must have different computer names.
Kerberos Realm	Specify the FQDN of the Active Directory domain. For example, if "pulsesecure" is the domain name (NetBIOS name), then pulsesecure.net is the Kerberos realm name.
Domain Join Configuration	

Username	Specify a username that has permission to join computers to the Active Directory domain. Use the "Delegate Control" workflow in Active Directory to assign the following user account permissions to the username or to a group to which the user belongs: • Write • Write All Properties • Change Password • Reset Password • Validate Write to DNS hostname • Read and write DNS host attributes • Delete Computer Objects • Create Computer Objects
Password	Specify the password for the special user.
Save Credentials	If this setting is not enabled, the credentials entered will be destroyed after successfully joining the domain. This option is useful when managing clusters. For example, you might want to save the credentials for a cluster node you have yet to add. If you do not enable this option, you must manually enter the credentials when you add the new cluster node.
Container Name	Specify the container path in Active Directory in which to create the machine account. Changing this field triggers a domain rejoin action. The default is **Computers**, which is a standard container created during installation of the AD server. The AD Computers container is the default location for new computer accounts created in the domain. If desired, you may specify a different container or OU. To specify nested containers, use a forward slash (*I*) as the container separator. For example: **outer OU/inner OU**. NOTE: Do not use backslashes in the path. Using backslashes causes an Invalid DN Syntax error.
Computer Name	Specify the machine account name. The default computer name is derived from the license hardware in the following format: 0161MT2L00K2C0. We recommend the Computer Name string contain no more than 14 characters to avoid potential issues with the AD/NT server. Do not include the '$' character.
Update Join Status / Reset Join	The following colors are used to indicate status: • Gray. The Domain Join action has not been attempted. This is the default status that appears when you are using the page to create a new Active Directory configuration. • Yellow. Attempting to join the Active Directory domain. This is the default status that appears after saving configuration settings or when any domain join settings are changed in an existing configuration. • Green. The attempt was successful. This status indicates that this server can now be used to authenticate users. • Red. The attempt to join the Active Directory domain was not successful. Click **Update Join** to get the latest join status of nodes. If the status appears persistently red, click **Reset Join** to reinitiate the domain join process. The Reset Join action requires Active Directory administrator credentials. NOTE: • For cluster nodes, you might need to click Update Join multiple times to obtain the latest join status of nodes. • Transient network issues might also cause the join status indicator to appear red. Before restarting the join process, ensure that it is not caused by network issues. Make sure your DNS servers can resolve queries to the Active Directory domain controller and that the Active Directory credentials are valid and have the appropriate permissions.

Additional Options

Authentication Protocol	The system attempts authentication using the protocols you have enabled in the order shown on the configuration page. For example, if you have selected the check boxes for Kerberos and NTLMv2, the system sends the credentials to Kerberos. If Kerberos succeeds, the system does not send the credentials to NTLMv2. If Kerberos is not supported or fails, the system uses NTLMv2 as the next protocol in order.
	Kerberos. Select this option to enable the Kerberos authentication protocol. Kerberos is the most secure method and is required for Kerberos single sign-on authentication. Kerberos must be enabled if you plan to use Pulse Secure client single sign-on or browser-based agentless single sign-on (SPNEGO).
	Enable NTLM protocol. Select this option to enable NTLM if you plan to use any of the following features:
	• Machine authentication using OAC, Pulse Secure client, or Windows native 802.1x supplicants.
	• MS-CHAP-based authentication protocols for any 802.1x supplicants.
	• User password management.
	• Role mapping rules based on group membership.
	If you enable NTLM, select one of the following versions:
	• **NTLMv2**. This protocol is moderately secure. It is required for machine authentication and MS-CHAP v2 based 802.1x authentication protocols.
	• **NTLMv1**. This protocol is comparatively less secure. It might be required for compatibility with existing legacy servers, MS-CHAP based servers, and MS-CHAP based 802.1x authentication protocols.
Trusted domain lookup	**Contact trusted domains**. Select this option to contact domain controllers of trusted domains directly without proxying authentication requests and group membership checks through the domain controller.
	If this option is not selected:
	• Network contact with trusted domains is not permitted, but pass-through authentication using the primary domain is still permitted.
	• Trusted domain user's group lookup for Kerberos SSO and SPNEGO authentication does not work even though user authentication succeeds.
	• Trusted domain user's password-based authentication does not work.
	• Only groups from the domain in which this system is a member are available for use in role mapping when a group search is performed in the server catalog window.
	NOTE: If you want to restrict trusted domain users and computers (machine authentication) from logging in when this option is not selected, you can define a custom expression based on the ntdomain variable and use it in role mapping rules. For example, if Policy Secure or Connect Secure belongs to the domain named Corporate, you can define a custom expression as **ntdomain=Corporate** and use the custom expression in the role mapping rule of the realm.
Domain Connections	**Maximum simultaneous connections per domain**. Enter the maximum number of simultaneous domain connections (1 to 10).
	This field specifies the maximum number of simultaneous connections that the auth daemon should open to the domain controller of one domain. A value of greater than 1 can improve the scalability with simultaneous authentication requests. However, this field value should be judiciously used, especially if trusted domain setting is enabled. This value dictates how many authentication processes are created per domain. For example: if the maximum domain connection is configured as 4 and there are 5 trusted domains, there could be as many as 5*4+1 = 21 auth processes. Hence if there are many trusted domains, the domain connection value needs to be controlled by the administrator, failing which there could be too many auth processes created only for AD authentication purpose.
	By default, this field value is set to 2 if trusted domain setting enabled. If trusted domain is not enabled then the default value is set to 5.
	NOTE: If **Contact trusted domains** is enabled, a value above 6 may degrade overall system performance.
SPNEGO Single Sign On	This feature is available only on Policy Secure.
	Enable SPNEGO. Select this option to support SPNEGO SSO.

	Keytab Upload. Select this option to use the controls to upload the SPNEGO keytab. The keytab must be generated on the Active Directory Service for the SPN. It must match the FQDN used to access this device.
Machine account password change	**Enable periodic password change of machine account**. Select this option to change the domain machine account password for this configuration.
	Change machine password frequency. Specify a frequency in days. For example, every 30 days.
User Record Synchronization	
This feature is available only on Connect Secure.	
Enable User Record Synchronization	Select this option to retain the bookmarks and individual preferences regardless of which system you log in to.
Logical Auth Server Name	Specify a logical authentication server name.
Save Changes?	

After you have saved your configuration, the "Mode" section is removed from the top of the configuration page and the "Active Directory Selection" section is included at the bottom of the page. Click the **Switch to Active Directory Legacy Mode** button to display the Active Directory Legacy Mode configuration page. The settings that are applicable to both configurations are populated in the Active Directory Legacy Mode configuration page. You must complete the remaining settings.

Configuring Authentication and Authorization with Active Directory Service (Legacy Mode)

To create an Active Directory Legacy Mode configuration:

1. Select **Authentication > Auth. Servers**.

2. Select **Active Directory / Windows NT** and click **New Server** to display the configuration page.

3. Select **Active Directory Legacy Mode** to display the Active Directory Legacy Mode configuration page.

4. Complete the configuration as described in Table 15.

5. Save the configuration.

Table 15 describes the Active Directory Legacy Mode configuration.

Table15: Active Directory Legacy Mode Settings

Settings	Guidelines
Mode	
	Select one of the following modes:
	• **Active Directory**—For recent versions of Windows Server.
	• **Active Directory Legacy Mode**—For Windows Server 2003 and earlier.
	This table describes Active Directory Legacy Mode.
Server	
Name	Specify a name to identify the server within the system.

Primary Domain Controller	Specify the name or IP address for the primary domain controller or Active Directory server.
Backup Domain Controller	(Optional) Specify the IP address of your backup domain controller or Active Directory server.
Domain	Specify the NetBios domain name for the organization. The NetBios name can be found within Active Directory, under Users and Computers. Right-click the domain name, and select **Properties**.
	The system uses DNS to discover domain controllers in the Active Directory forest. It sends authentication requests to the domain controller at the closest site. Ensure that your DNS servers are configured to resolve Active Directory domain controller FQDNs and SRV records.
	Do not create an Active Directory and an Active Directory Legacy Mode configuration with the same domain and computer name. They must have different computer names.
	You may, however, create two Active Directory Legacy Mode configurations with the same domain and the same computer name.
Allow domain to be specified as part of username	Select this option to allow users to sign in by entering a domain name in the Username box in the format: *domain\username*.
Allow trusted domains	Select this option to get group information from all trusted domains within a forest.
Domain Controller is a Windows 2008 server	Select this option if the domain controller is a Windows 2008 server. The Windows 2008 server has several enhancements to the Active Directory Server, which is now called Active Directory Domain Services.
Administrator	
Admin Username	Specify an administrator username for the AD or NT server. Make sure the administrator you specify is a domain administrator in the same domain as the AD or NT server. Do not include a domain name with the server administrator username in the Admin Username box.
Admin Password	Specify the corresponding password.
Additional Options	
Authentication Protocol	The system attempts authentication using the protocols you have enabled in the order shown on the configuration page. For example, if you have selected the check boxes for Kerberos and NTLMv2, the system sends the credentials to Kerberos. If Kerberos succeeds, the system does not send the credentials to NTLMv2. If Kerberos is not supported or fails, the system uses NTLMv2 as the next protocol in order.
	Kerberos. This protocol is the most secure method and is required for Kerberos single sign-on authentication. Kerberos must be enabled if you plan to use Pulse Secure client single sign-on or browser-based agentless single sign-on (SPNEGO).
	NTLMv2. This protocol is moderately secure. It is required for machine authentication and MS-CHAP v2 based 802.1x authentication protocols.
	NTLMv1. This protocol is comparatively less secure. It might be required for compatibility with existing legacy servers, MS-CHAP based servers, and MS-CHAP based 802.1x authentication protocols.
Kerberos Realm Name	Select one of the following options to specify the Kerberos realm name:
	• **Use LDAP to get Kerberos realm name**. Select this option to retrieve the Kerberos realm name from the Active Directory server using the specified administrator credentials.
	• **Specify Kerberos realm name**. Select this option to specify the realm name in the corresponding text box. Specify the FQDN of the Active Directory domain. For example, if "pulsesecure" is the domain name (NetBIOS name), then pulsesecure.net is the Kerberos realm name.

SPNEGO Single Sign On	This option is available only on Policy Secure.
	Enable SPNEGO. Select this option to support SPNEGO SSO.
	Keytab Upload. Use the controls to upload the SPENEGO keytab. The keytab must be generated on the Active Directory Service for the SPN. It must match the FQDN used to access this device.
Group Search With LDAP	**Enable Group Search With LDAP**. Select this option to use LDAP to retrieve group membership information. If you enable group search with LDAP, specify one of the following LDAP protocols:
	• **LDAP**. Select this option to use unsecure LDAP.
	• **LDAPS**. Select this option to use secure LDAP (LDAP communication using an SSL tunnel). You enable LDAPS by installing a properly formatted certificate. If you select LDAPS, optionally select **Validate Server Certificate** to validate this certificate.
	NOTE: Active directory group lookup using LDAP fails if a user from a different domain tree attempts to log in. Native Active Directory group lookup works correctly.
	Active directory group lookup using LDAP fails if the same username is present in both the parent and child domain and if only the child domain user is part of a group. Native Active Directory group lookup works correctly.
	Admin DN. Specify the administrator Distinguished Name. You can click the **Test Configuration** button at the bottom of the page to automatically populate this field.
	Base Search DN. Specify the point to start searching for the user. For example, **dc=pulsesecure,dc=com**. You can click the **Test Configuration** button at the bottom of the page to automatically populate this field.
	Nested Group Level. Specify how many levels within a group to search for the user. Higher numbers create longer queries or search time.
Advanced Options	
User may belong to Domain Local Groups across trust boundaries	Specify if the trusted domains' Domain Local Groups must be searched to determine users' group memberships. This might impact performance.
Container Name	Specify the Container path in Active Directory in which to create the machine account. Changing this field triggers a domain rejoin action.
	The default is **Computers**, which is a standard container created during installation of the AD server. The AD Computers container is the default location for new computer accounts created in the domain.
	If desired, you may specify a different container or OU. To specify nested containers, use a forward slash (**/**) as the container separator. For example: **outer OU/inner OU**.
	NOTE: Do not use backslashes in the path. Using backslashes causes an Invalid DN Syntax error.
Computer Name	The computer name is pre-filled with an entry in the format of vcNNNNHHHHHHHH, where, in an IVS system, the NNNN is the IVS ID (assuming you have an IVS license) and the HHHHHHHH is a hex representation of the IP address of the Pulse Secure device. You can change this to a unique name, either the one provided by default or one of your own choosing to more easily identify your systems in the Active Directory. In a non-IVS system, the first six characters of the name will be 'vc0000' because there is no IVS ID to display. For example, the name could be something like vc0000a1018dF2' for a non-IVS system. The following entry conforms with the required format: **IVE_123**.
	In a clustered environment with the same AD authentication server, this name is also unique among all cluster nodes, and this configuration section displays all of the identifiers for all attached cluster nodes.

User Record Synchronization

This feature is available only on Connect Secure.

Enable User Record Synchronization	Select this option to retain the bookmarks and individual preferences regardless of which system you log in to.
Logical Auth Server Name	Specify a logical authentication server name.
Save Changes?	

After you have saved your configuration, the "Mode" section is removed from the top of the configuration page and the "Active Directory Selection" section is included at the bottom of the page. Click the **Switch to Active Directory mode** button to display the Active Directory standard mode configuration page. The settings that are applicable to both configurations are populated in the Active Directory standard mode configuration page. You must complete the remaining settings.

Displaying the User Accounts Table

To display user accounts:

1. Select **Authentication > Auth. Servers**.

2. Click the link for the authentication server you want to manage.

3. Click the **Users** tab to display the user accounts table.

   The user accounts table includes entries for the accounts that have been created. The Last Sign-in Statistic column shows the last successful sign-in date and time for each user, the user's IP address, and the agent or browser type and version.

4. Use the controls to search for users and manage user accounts:

   - To search for a specific user, enter a username in the Show users named field and click **Update**.

     TIP: You can use an asterisk (*) as a wildcard, where * represents any number of zero or more characters. For example, to search for all usernames that contain the letters jo, enter ***jo***. The search is case-sensitive. To display the entire list of accounts again, type ***** or delete the field's contents and click **Update**.

   - To limit the number of users displayed on the page, enter a number in the Show N users field and click **Update**.

   - To terminate their user session and delete the account, select the check box next to the user account record and click **Delete**.

Troubleshooting the Active Directory Service (Standard Mode)

To troubleshoot the Active Directory Service (standard mode):

1. Select **Authentication > Auth. Servers >** *AD Server name* **> Troubleshooting**.

2. Select the appropriate functions described in Table 16.

Table16: Active Directory Server Troubleshooting Functions

Function	Description

Basic Verification	Verifies whether the domain is properly joined and if the **winbindd** service is running. The number of winbindd processes is displayed, along with the ongoing CPU and memory usage for each process. For example, if user authentication is slow or fails randomly, use this function to check the number of **winbindd** processes and the CPU, memory and file descriptor usage. Select **Restart AD Services** to correct faulty processes.
Test User Authentication	Prompts for a username and password and attempts to log in. If successful, the groups the user belongs to are displayed. Only the regular password authentication is done (MS-CHAPv2 and other types of authentication are not available).
Test User Password Change	Prompts for a username and the old and new password for a user and attempts to change the password on the AD server.
List Domain Info	Lists each domain and all trusted domains. Selecting a domain lists each Domain Controller for the domain, its IP address, and whether it is reachable. For example, if user authentication fails consistently and the domain is shown as successfully joined in the AD Server Settings page, the domain trust may be broken. Use this function to check the trusted domains. Also, if the domain join fails consistently or user authentication to a trusted domain fails consistently, the domain might not be reachable or the DNS configuration may be incorrect. Use this function to verify whether the domains and trusted domain are reachable.
Change Machine Password	Sends a request to the domain controller to change the machine password. A confirmation prompt is displayed to confirm the change.
Restart AD Services	Restarts the **winbindd** process, which may restore proper authentication, specifically during load and longevity scenarios. A confirmation prompt is displayed to confirm the restart (users cannot log in during a restart).
Reset Join	Reinitiates the domain join process. A confirmation prompt is displayed to confirm the reset and allows you to clear the Samba cache and keytab files before the reset. This is the same function shown on the AD server's Settings page and requires Active Directory administrator credentials. For example, if user group changes are not reflected in the user authentication, run this function with **Clear Samba Cache** enabled.
Samba Diagnostics Logs	Displays Diagnostic Logs page where you can download the Samba logs.
Load Output	Displays up to the last 500 lines of the troubleshooting output for the current session.
Save Output File	Saves all the troubleshooting messages for the current session.
Clear Output File	Erases all the troubleshooting messages saved in the output file (they cannot be retrieved).

Related Documentation

- AAA Server Overview

- *Using Kerberos SSO*

- Understanding Multidomain User Authentication

- Understanding Active Directory and Windows NT Group Information Support

Understanding Multidomain User Authentication

This topic provides an overview of multidomain user authentication with Active Directory and Windows NT. It includes the following information:

- Multi-Domain User Authentication Overview

- Windows NT User Normalization
- Kerberos Support
- Windows NT4 Support

Multi-Domain User Authentication Overview

The Pulse Secure access management framework allows for multidomain Active Directory and Windows NT authentication. The system authenticates users in the domain that you configure, users in child domains, and users in all domains trusted by the configured domain.

Users in the default domain can sign into the system using just their username, or the default domain and the username in the format default-domain\username.

When you enable trusted domain authentication, users in trusted or child domains can sign in using the name of the trusted or child domain and the username in the format trusted-domain\username. Note that enabling trusted domain authentication adds to the server response time.

Windows NT User Normalization

To support multidomain authentication, the Pulse Secure access management framework uses "normalized" Windows NT credentials when it contacts an Active Directory or Windows NT4 domain controller for authentication. Normalized Windows NT credentials include both the domain name and the username: domain\username. Regardless of how the user signs in (either using just a username or using the domain\username format), the access management framework always processes the username in domain\username format.

When a user signs in using only their username, the access management framework normalizes their Windows NT credentials as default-domain\username. Authentication succeeds only if the user is a member of the default domain.

When a user signs in using the domain\username format, the access management framework attempts to authenticate the user as a member of the domain the user specifies. Authentication succeeds only if the user-specified domain is a trusted or child domain of the default domain. If the user specifies an invalid or untrusted domain, authentication fails.

Two variables, <NTUser> and <NTDomain>, allow you to individually refer to Windows NT domain and username values. The system populates these two variables with the Windows NT domain and username information.

In role mapping rules, when you specify **USER = someusername**, the system treats this rule semantically as **NTUser = someusername AND NTDomain = defaultdomain**.

Kerberos Support

We recommend you configure the Pulse Secure access management framework to use the Kerberos authentication protocol with Windows domain controllers. When a user logs in to the system, the system performs Kerberos authentication and attempts to fetch the Kerberos realm name for the domain controller, as well as all child and trusted realms, using LDAP calls.

You can use Kerberos differently. You can specify the Kerberos realm name when configuring an Active Directory authentication server. We do not recommend this method for two reasons:

- You cannot specify more than one realm name. The system cannot then authenticate against child or trusted realms of the realm you specify.

- If you misspell the realm name, the system cannot authenticate users against the proper realm.

Windows NT4 Support

The Pulse Secure access management framework does not support Kerberos-based authentication in Windows NT4 domain controllers. The system uses NTLM with a backend Windows NT4 domain controller.

Related Documentation

- Using Active Directory

- Understanding Active Directory and Windows NT Group Information Support

Understanding Active Directory and Windows NT Group Information Support

This topic describes support for polling group information from Active Directory and Windows NT servers. It includes the following information:

- Active Directory Group Information Overview

- Windows NT4 Group Information Overview

Active Directory Group Information Overview

The Pulse Secure access management framework supports user group lookup in Domain Local, Domain Global, and Universal groups in the default domain, child domains, and all trusted domains. The system obtains group membership using one of three methods that have different capabilities:

- Group information in User's Security Context—Returns information about the user's Domain Global groups.

- Group information obtained using LDAP search calls—Returns information about the user's Domain Global groups and about the user's Universal groups if the access management framework queries the Global Catalog Server.

- Group information using native RPC calls—Returns information about the user's Domain Local Group.

With respect to role-mapping rules, the system attempts group lookup in the following order:

- Checks for all Domain Global groups using the user's security context.

- Performs an LDAP query to determine the user's group membership.

- Performs an RPC lookup to determine the user's Domain Local group membership.

Windows NT4 Group Information Overview

The Pulse Secure access management framework supports group lookup in the Domain Local and Domain Global groups created in the default domain, as well as all child and other trusted domains. The system obtains group membership using:

- Domain Global group information from the user's security context.

- Domain Local information using RPC calls.

In the Windows NT4 environment, the system does not use LDAP-based search calls.

Related Documentation

- Using Active Directory

Join Domain for Active Directory-based Authentication Server Without Using a Domain Admin Account

With Active Directory on Windows Server 2000, Windows Server 2003 and Windows Server 2008, the system can *join domain* (for an Active Directory based Authentication server) without using a domain administrator account.

1. Identify the user or group.

   - Identify the user or group to be granted permissions to perform AD operations on behalf of the system. This can be any pre-existing user or group, or it can be a new one that you create.

   - If you use a group, be sure to add a user to the group with which you will configure the Active Directory authentication server.

2. Open the Active Directory Users and Computers window.

   1. Click **Start > Programs > Administrative Tools > Active Directory Users and Computers**.

   2. In the Active Directory Users and Computers interface, click **View**, and then select **Advanced Features**.

3. Open the Access Control Settings for Computers dialog box.

   1. In the Active Directory Users and Computers window, right-click **Computers** and then click **Properties**. The Computer Properties dialog appears.

   2. Click the **Security** tab and then click **Advanced**. The Advanced Security Settings for Computers dialog box appears.

3. Grant the user (or group) permission to Create Computer Objects and Delete Computer Objects in the Computers container.

   1. In the Advanced Security Settings for Computers dialog box, click **Add**. The Select User, Computer, or Group dialog box appears.

   2. Select or manually enter the name of the user or group you want to grant permissions to perform AD operations on behalf of the system, and then click **OK**.

   3. In the Permission Entry for Computers dialog box that appears, click **This object only** in the Apply onto list.

4. In the Permissions list, select the **Create Computer Objects** check box and the **Delete Computer Objects** check box, and then click **OK**.

5. Grant the user (or group) permission to reset passwords and modify permissions on objects.

    1. In the Active Directory Users and Computers window, right-click **Users** and then click **Properties**. The User Properties dialog box appears.

    2. Click the **Security** tab and then click **Advanced**. The Advanced Security Settings for Users dialog box appears.

    3. Click **Add**. The Select User, Computer, or Group dialog box appears.

    4. Click or manually enter the name of the user (or group) you want to grant permissions to perform AD operations on behalf of the system, and then click **OK**.

    5. In the Permission Entry for Users dialog box, click **User objects** in the Apply onto list.

    6. n the Permissions list, select the **Reset Password** check box, and then click OK.

7. Verify there are no Deny entries that would affect the user (or group).

8. Exit the Advanced Security Settings for Users dialog box.

    1. In the Advanced Security Settings for Users dialog box, click **OK**.

    2. In the Users Properties dialog box, click **OK**.

3. Exit the Active Directory Users and Computers interface.

4. Configure the Active Directory Authentication Server.

    1. In the Administrator section of the Active Directory authentication server configuration, enter the username you identified in Step 1 above and their password in the Admin Username and Admin Password fields respectively, and click **Save Changes**.

    2. Restart the system services.

Related Documentation

- Using Active Directory

Using the Anonymous Server

This topic describes integration with the anonymous server. It includes the following information:

- Anonymous Server Overview
- Configuring Authentication with the Anonymous Server
- Monitoring Anonymous User Sessions

Anonymous Server Overview

This section describes support for using Policy Secure and Connect Secure with the anonymous server. It includes the following sections:

- Understanding the Anonymous Server
- Anonymous Server Feature Support

- Interoperability Requirements and Limitations

Understanding the Anonymous Server

The anonymous server is a local authentication server that allows any user to access the system without providing a username and password.

Instead, when a user enters the URL of a sign-in page that is configured to authenticate against an anonymous server, the Pulse Secure access management framework bypasses the standard sign-in page and immediately displays the welcome page to the user.

Anonymous Server Feature Support

Pulse Secure access management framework supports the following anonymous server features:

- Enables guest access without username or password

- Supports Host Checker scans before allowing a guest device to connect to the network

- Supports firewall enforcement roles and policies to limit the resources available to the guest user

Interoperability Requirements and Limitations

The following limitations apply to the anonymous server configuration and logging:

- You can add only one anonymous server configuration.

- You cannot create an administrator realm that uses the anonymous server. Anonymous administration is not allowed.

- During configuration, you must choose the anonymous server as both the authentication server and the directory or attribute server in the Users > User Realms > General tab.

- When creating role mapping rules through the Users > User Realms > Role Mapping tab, the Pulse Secure access management framework does not allow you to create mapping rules that apply to specific users (such as "Joe"), because the anonymous server does not collect username information. You can only create role mapping rules based on a default username (*), certificate attributes, or custom expressions.

- For security reasons, you might want to limit the number of users who sign in through an anonymous server at any given time. To do this, use the option on the Users > User Realms > [Realm] > Authentication Policy > Limits tab (where [Realm] is the realm that is configured to use the anonymous server to authenticate users).

Configuring Authentication with the Anonymous Server

To configure authentication with the anonymous server:

1. Select **Authentication** > **Auth. Servers**.

2. Select **Anonymous Server** and click **New Server** to display the configuration page.

3. Complete the configuration as described in Table 17.

4. Save the configuration.

Table17: Anonymous Server Settings

Settings	Guidelines
Name	Specify a name to identify the server within the system.
User Record Synchronization	This feature is available only on Connect Secure.
Enable User Record Synchronization	Select this option to retain the bookmarks and individual preferences regardless of which system you log in to.
Logical Auth Server Name	Specify a logical authentication server name.

Monitoring Anonymous User Sessions

The purpose of the anonymous server is to enable unauthenticated access. Therefore, the system does not maintain session tables, and the Anonymous Server configuration page does not have a corresponding Users tab. The system does maintain user access logs for anonymous access. The username is recorded in the user access log as "AnonUser1234". If the user is logging in using the agentless access method, the user access log records the host's IP address. You can view the User Access Log file by navigating to System > Log/Monitoring.

Figure 23 shows the User Access Log Page.

Figure23: User Access Log

Related Documentation

- AAA Server Overview

- *Creating an Authentication Realm*

- Creating an Authentication Realm

- KB: How to reduce the impact of an unreachable authentication server

- KB: Configuring Realm Selection

Using the Certificate Server

This topic describes integration with the certificate server. It includes the following information:

- Certificate Server Overview
- Configuring Authentication with the Certificate Server
- Displaying the User Accounts Table

Certificate Server Overview

This section describes support for using Policy Secure and Connect Secure with the certificate server. It includes the following sections:

- Understanding the Certificate Server
- Feature Support
- Interoperability Requirements and Limitations

Understanding the Certificate Server

The certificate server is a local server that allows user authentication based on the digital certificate presented by the user without any other user credentials.

When you use a certificate server, the user experience is similar to anonymous authentication. If the certificate is secured through a hardware or a software token or through a password, the certificate server authentication is very useful. The certificate contains the full distinguished name (DN) and the system extracts the values from the DN and uses it for role mapping rules, authentication policies, and role restrictions.

Feature Support

The Pulse Secure access management framework supports the following certificate server features:

- Certificate directory services to retrieve user attributes in role mapping rules, authentication policies, and role restrictions.

- Load CA-created certificates on the system.

- Load multiple certificates from different CAs for use with different authentication realms.

Interoperability Requirements and Limitations

If you choose a certificate attribute with more than one value, the system uses the first matched value. For example, if you enter <certDN.OU> and the user has two values for the attribute (ou=management, ou=sales), the system uses the "management" value.

To use all values, add the SEP attribute to the variable. For example, if you enter <certDN.OUT SEP=":"> the system uses "management:sales".

Configuring Authentication with the Certificate Server

To configure authentication with the certificate server:

1. Select **Authentication** > **Auth. Servers**.

2. Select **Certificate Server** and click **New Server** to display the configuration page.

3. Complete the configuration as described in Table 18.

4. Save the configuration.

Table18: Certificate Server Settings

Settings	Guidelines
Name	Specify a name to identify the server within the system.
User Name Template	Specify a username template. Specify how the system should construct a username. You may use any combination of certificate variables contained in angle brackets and plain text. NOTE: This value populates the <USER> and <USERNAME> session variables for use throughout the rest of the system configuration.
User Record Synchronization	This applies only to Connect Secure.
Enable User Record Synchronization	Select this option to retain the bookmarks and individual preferences regardless of which system you log in to.
Logical Auth Server Name	Specify a logical authentication server name.

Displaying the User Accounts Table

To display user accounts:

1. Select **Authentication > Auth. Servers**.

2. Click the link for the authentication server you want to manage.

3. Click the **Users** tab to display the user accounts table.

   The user accounts table includes entries for the accounts that have been created. The Last Sign-in Statistic column shows the last successful sign-in date and time for each user, the user's IP address, and the agent or browser type and version.

4. Use the controls to search for users and manage user accounts:

   - To search for a specific user, enter a username in the Show users named box and click **Update**.

> **TIP:** You can use an asterisk (*) as a wildcard, where * represents any number of zero or more characters. For example, to search for all usernames that contain the letters jo, enter ***jo***. The search is case-sensitive. To display the entire list of accounts again, type ***** or delete the field's contents and click **Update**.

- To limit the number of users displayed on the page, enter a number in the Show *N* users box and click **Update**.

- To terminate the user session and delete the account, select the check box next to the user account record and click **Delete**.

Related Documentation

- AAA Server Overview

- KB: Connect Secure Troubleshooting: Certificate Authentication with Windows 7 Clients

- How to: Connect Secure: Using Certificates

Using an LDAP Server

This topic describes integration with the LDAP server. It includes the following information:

- LDAP Server Overview

- Configuring Authentication with an LDAP Server

- Displaying the User Accounts Table

LDAP Server Overview

This section describes support for using Policy Secure and Connect Secure with the LDAP server. It includes the following sections:

- Understanding LDAP Server

- LDAP Feature Support

- Interoperability Requirements and Limitations

Understanding LDAP Server

Lightweight Directory Access Protocol (LDAP) facilitates the access of online directory services. The Internet Engineering Task Force (IETF) designed and specified LDAP as a better way to make use of X.500 directories, having found the original Directory Access Protocol (DAP) too complex for average Internet clients to use. LDAP is a relatively simple protocol for updating and searching directories running over TCP/IP.

LDAP directory consists of a collection of attributes with a name, known as a distinguished name (DN). Each of the entry's attributes, known as a relative distinguished name (RDN), has a type and one or more values. The types are typically mnemonic strings, such as CN for common name. The valid values for each field depend on the types.

The full DN is constructed by stringing together RDNs from most specific to least specific, separated by commas, as shown in the following example:

cn=Bob_Employee, ou= account_mgr, o=sales, dc=Acme,dc=com.

LDAP Feature Support

Pulse Secure access management framework supports the following LDAP features:

- LDAP directory services to retrieve user attributes and group membership in role mapping rules

- Encrypted connections to the LDAP server using LDAP over SSL (LDAPS) or Start Transport Layer Security (TLS)

- Password management feature enabling users who access an LDAP server to manage their passwords using the policies defined on the LDAP server

- Fine-grained password policy (FGP) for Active Directory 2008

Interoperability Requirements and Limitations

The following limitations apply to interoperability with LDAP:

- By default, challenge response protocols are disabled for LDAP servers. Use these protocols only with noninteractive devices (for example, phones), as password management is not possible if these protocols are used for authentication.

- To use the CHAP, EAP-MD5-Challenge, MS-CHAP-V1, and MS-CHAP-V2 protocols, the LDAP server must store the user's password in cleartext.

- Backup LDAP servers must be the same version as the primary LDAP server. Also, we recommend that you specify the IP address of a backup LDAP server instead of its hostname, which might accelerate failover processing by eliminating the need to resolve the hostname to an IP address.

Configuring Authentication with an LDAP Server

To configure authentication with an LDAP server:

1. Select **Authentication** > **Auth. Servers**.

2. Select **LDAP Server** and click **New Server** to display the configuration page.

3. Complete the configuration as described in Table 19.

4. Save the configuration.

Table19: LDAP Server Settings

Settings	Guidelines
Name	Specify a name to identify the server within the system.
LDAP Server	Specify the LDAP server name or the IP address.

LDAP Port	Specify the LDAP port for the LDAP server.
	Default port number: 389 (unencrypted connection)
	Default port number: 636 (SSL connection)
Backup LDAP Server1	(Optional) Specify the parameters for backup LDAP server1.
	The specified backup LDAP server is used for failover processing. The authentication request is first routed to the primary LDAP server, and then to the specified backup servers if the primary server is unreachable.
Backup LDAP Port1	Specify the parameters for backup LDAP port1.
Backup LDAP Server2	(Optional) Specify the parameters for backup LDAP server2.
Backup LDAP Port2	Specify the parameters for backup LDAP port2.
LDAP Server Type	Select the backend LDAP server type from the following choices: • **Generic** • **Active Directory** • **iPlanet** • **Novell eDirectory** • **Profiler** (*Policy Secure only*)
Connection	Select one of the following options for the connection to the LDAP server: • **Unencrypted**– The device sends the username and password to the LDAP Directory Service in cleartext. • **LDAPS**– The device encrypts the data in the LDAP authentication session using the Secure Socket Layer (SSL) protocol before sending it to the LDAP Directory Service. • **Start TLS**– The device allows both secure and plain requests against an LDAP server on a single connection. NOTE: • If you select **LDAPS** or **Start TLS**, the Validate Certificate option is displayed for the configured LDAP server(s) and its referral servers. Select this option if the SSL connection uses digital certificate security. • If you enable validation for the referral servers, make sure your network DNS supports reverse lookup zone. • If you want to verify the server certificates, the root CA and Intermediate CAs must be imported as trusted CAs.
Connection Timeout (seconds)	Specify the time to wait for connection to the primary LDAP server, and then to each backup LDAP server. Default: 15 seconds
Search Timeout (seconds)	Specify the time to wait for search results from a connected LDAP server.
Test Connection	(Optional) To verify the connection between Pulse Secure client and LDAP servers, click the **Test Connection** button. NOTE: We recommend using the Test Connection function only after saving changes on the LDAP Server Configuration page.
Authentication required?	

Authentication required to search LDAP	Select this option to require authentication when performing search or password management operations. NOTE: • If you use Active Directory, you must select the **Authentication required to search LDAP** check box and provide the full DN and password of an account that can reach Active Directory. • You can enable password management on any LDAP server. This feature enables users who authenticate through an LDAP server to manage their passwords through the system using the policies defined on the LDAP server. To enable password management on any LDAP server, you must provide an administrator account (with write privileges to the directory) for the administrator DN.
Admin DN	Specify the administrator DN for queries to the LDAP directory.
Password	Specify the password for the LDAP server.
Finding user entries	
Base DN	Specify the base DN under which the users are located. For example, dc=sales,dc=acme, dc=com.
Filter	Specify a unique variable that can be used to do a fine search in the tree. For example, samAccountname=<username> or cn=<username>. NOTE: • Include <username> in the filter to use the username entered on the sign-in page for the search. • Specify a filter that returns 0 or 1 user DNs per user; the device uses the first DN returned if more than 1 DN is returned.
Remove Domain from Windows users names? (Applicable to UAC only)	
Strip domain from Windows username	Select this option to pass the username without the domain name to the LDAP server.
Enable Challenge-Response open protocols? (Applicable to UAC only)	
Enable Challenge-Response open protocols	Select this option if you want to use a challenge-response protocol for authentication. NOTE: By default, these protocols are disabled for LDAP servers because account management is not possible.
Determining group membership	
Base DN	Specify the base DN to search for user groups.
Filter	Specify a unique variable which can be used to do a fine search in the tree. For example, samAccountname=<username> or cn=<GROUPNAME>.
Member Attribute	Specify all the members of a static group. For example, member or uniquemember (iPlanet specific).
Reverse group search	Select this option to start the search from the member instead of the group. This option is available only for Active Directory server types.
Query Attribute	Specify an LDAP query that returns the members of a dynamic group. For example, memberURL.
Nested Group Level	Specify how many levels within a group to search for the user. NOTE: The higher the number, the longer the query time, so we recommend that you specify to perform the search no more than two levels deep.
Nested Group Search	Select one of the following options: • **Nested groups in Server Catalog**–This option is faster because it can search within the implicit boundaries of the nested group. • **Search all nested groups**–With this option, the device searches the Server Catalog first. If the device finds no match in the catalog, then it queries LDAP to determine if a group member is a subgroup.

Displaying the User Accounts Table

To display user accounts:

1. Select **Authentication > Auth. Servers**.

2. Click the link for the authentication server you want to manage.

3. Click the **Users** tab to display the user accounts table.

   The user accounts table includes entries for the accounts that have been created. The Last Sign-in Statistic column shows the last successful sign-in date and time for each user, the user's IP address, and the agent or browser type and version.

4. Use the controls to search for users and manage user accounts:

   - To search for a specific user, enter a username in the Show users named box and click **Update**.

     TIP: You can use an asterisk (*) as a wildcard, where * represents any number of zero or more characters. For example, to search for all usernames that contain the letters jo, enter ***jo***. The search is case-sensitive. To display the entire list of accounts again, type ***** or delete the field's contents and click **Update**.

   - To limit the number of users displayed on the page, enter a number in the Show *N* users box and click **Update**.

   - To terminate the user session and delete the account, select the check box next to the user account record and click **Delete**.

Related Documentation

- AAA Server Overview

- Using the LDAP Password Management Feature

- *Using the LDAP Server Catalog*

Using the LDAP Password Management Feature

This topic describes support and limitations for LDAP password management. It includes the following information:

- LDAP Password Management Feature Overview

- Enabling LDAP Password Management

- LDAP Password Management Support

- LDAP Password Management for Windows AD Versions

- Troubleshooting LDAP Password Management

LDAP Password Management Feature Overview

The password management feature enables users who access an LDAP server to manage their passwords through the Pulse Secure access management framework using the policies

defined on the LDAP server. For example, if a user tries to sign in to the system with an LDAP password that is about to expire, the system notices the user through the interface, and then passes the user's response back to the LDAP server without requiring the user to sign in to the LDAP server separately.

Users, administrators, and help desk administrators who work in environments where passwords have set expiration times may find the password management feature very helpful. If users are not informed that their passwords are about to expire, they can change them themselves through the system rather than call the help desk.

Once this feature is enabled, the system performs a series of queries to determine user account information, such as when the user's password was last set, whether the account is expired, and so on. The Pulse Secure access management framework does this by using its internal LDAP or Samba client. Many servers, such as Microsoft Active Directory or Sun iPlanet, offer an Administrative Console to configure account and password options.

LDAP-based password management works with only three types of LDAP servers:

- Microsoft Active Directory. For Active Directory, password policy attributes can be configured in the user entry container level or any organization level above the user container. If these attributes are configured at multiple levels, the level closest to the user node takes precedence. The password management feature is not supported on the Active Directory Global Catalog because password policy attributes are not fully populated in the Active Directory Global Catalog.

- For Active Directory 2008, the Pulse Secure access management framework supports the Fine Grained Password Policy (FGP) configured in the AD user container.

- Sun Microsystems iPlanet

- Novell eDirectory

LDAP-based password management does not work on generic LDAP servers such as OpenLDAP.

The system relies on the back-end server to pinpoint the cause of error when a password change operation fails. However, although LDAP servers may report errors accurately to human operators, they do not always do so when communicating programmatically to systems. Therefore, reported errors might be generic or cryptic.

The system does not support customized password policies.

Enabling LDAP Password Management

To enable password management, you must first create an instance of the LDAP server. Next, you associate the LDAP server with the applicable realms. Finally, you select the enable password management feature at the realm level.

LDAP Password Management Support

The Pulse Secure access management framework supports password management with the following LDAP directories:

- Microsoft Active Directory/Windows NT

- Sun iPlanet

- Novell eDirectory

- Generic LDAP directories, such as IBM Secure Directory and OpenLDAP

Table 20 describes supported password management functions, their corresponding function names in the individual LDAP directories, and any additional relevant details. These functions must be set through the LDAP server itself before the system can pass the corresponding messages, functions, and restrictions to end users.

The Active Directory attribute names shown are specific to the Domain Security Policy object. Similar attributes for the corresponding functions are used for the Active Directory 2008 Fine-Grained Password Policy. Refer to Microsoft documentation for details.

When authenticating against a generic LDAP server, such as IBM Secure Directory, the system supports only authentication and allowing users to change their passwords. Password management functions are not supported when the CHAP family protocols are used for authentication. All functions are available when the JUAC protocol is used for authentication (Policy Secure only).

Table20: Supported Password Management Functions

Function	Active Directory	iPlanet	eDirectory	Generic
Authenticate user	unicodePwd	userPassword	userPassword	userPassword
Allow user to change password if enabled	Server tells us in bind response (uses ntSecurityDescriptor)	If passwordChange == ON	If passwordAllowChange == TRUE	Yes
Log out user after password change	Yes	Yes	Yes	Yes
Force password change at next login	If pwdLastSet == 0	If passwordMustChange == ON	If pwdMustChange == TRUE	-
Expired password notification	userAccountControl== 0x80000	If Bind Response includes OID 2.16.840.1.113730.3.4.4 == 0	Check date/time value	-
Password expiration notification (in X days/hours)	if pwdLastSet - now() < maxPwdAge - 14 days (Read from domain attributes) (The system displays warning if less than 14 days)	If Bind Response includes control OID 2.16.840.1.113730.3.4.5 (contains date/time) (The system displays warning if less than 14 days)	If now() - passwordExpirationTime< 14 days (The system displays warning if less than 14 days)	-

Disallow authentication if "account disabled/locked"	userAccountControl== 0x2 (Disabled) accountExpires userAccountControl == 0x10 (Locked) lockoutTime	Bind ErrorCode: 53 "Account Inactivated" Bind Error Code: 19 "Exceed Password Retry Limit"	Bind ErrorCode: 53 "Account Expired" Bind ErrorCode: 53 "Login Lockout"
Honor "password history"	Server tells us in bind response	Server tells us in bind response	Server tells us in bind response
Enforce "minimum password length"	If set, the system displays message telling user minPwdLength	If set, the system displays message telling user passwordMinLength	If set, the system displays message telling user passwordMinimumLength
Disallow user from changing password too soon	If pwdLastSet - now() < minPwdAge, then we disallow	If passwordMinAge > 0, then if now() is earlier than passwordAllowChangeTime, then we disallow	Server tells us in bind response
Honor "password complexity"	If pwdProperties == 0x1, then enabled. Complexity means the new password does not contain username, first or last name, and must contain characters from 3 of the following 4 categories: English uppercase, English lowercase, Digits, and Non-alphabetic characters (ex. !, $, %)	Server tells us in bind response	Server tells us in bind response

Note the following expected behavior:

- When you select the User must change password after reset option on the iPlanet server, you must also reset the user password before this function takes effect. This issue is a limitation of iPlanet.

- The system displays a warning about password expiration only if the password is scheduled to expire in 14 days or less. The system displays the message during each sign-in attempt. The warning message contains the remaining number of days, hours, and minutes that the user has to change the password before it expires on the server. The default value is 14 days, but you can change it on the password configuration page of the admin console.

LDAP Password Management for Windows AD Versions

The Pulse Secure access management framework supports password management with the following Windows servers:

- Microsoft Active Directory 2008

- Microsoft Active Directory 2003

- Windows NT 4.0

Table 21 describes supported password management functions. These functions are not supported for a layer 2 connection when CHAP, MS-CHAP, or PAP are used as authentication protocols.

Table21: AD/NT Password Management Matrix

Function	Active Directory	Active Directory 2003	Active Directory 2008 FGP	Windows NT
Authenticate user	Yes	Yes	Yes	Yes
Allow user to change password if licensed and if enabled	Yes	Yes	Yes	Yes
Log out user after password change	Yes	Yes	Yes	Yes
Force password change at next login	Yes	Yes	Yes	Yes
Password expired notification	Yes	Yes	Yes	Yes
Account disabled	Yes	Yes	Yes	Yes
Account expired	Yes	Yes	Yes	Yes

Note the following expected behavior:

- Changes on the Active Directory domain security policy can take 5 minutes or longer to propagate among Active Directory domain controllers. Additionally, this information does not propagate to the domain controller on which it was originally configured for the same time period. This issue is a limitation of Active Directory.

- When changing passwords in Active Directory using LDAP, the system automatically switches to LDAPS, even if LDAPS is not the configured LDAP method. To support LDAPS on the Active Directory server, you must install a valid SSL certificate into the server's personal certificate store. The certificate must be signed by a trusted CA, and the CN in the certificate's Subject field must contain the exact hostname of the Active Directory server, (for example: adsrv1.company.com). To install the certificate, select the Certificates Snap-In in the Microsoft Management Console (MMC).

- The Account Expires option in the User Account Properties tab only changes when the account expires, not when the password expires. Microsoft Active Directory calculates the password expiration using the Maximum Password Age and Password Last Set values retrieved from the User object and Fine-Grained Password Policy objects or the Domain Security Policy LDAP objects.

- The system displays a warning about password expiration only if the password is scheduled to expire in 14 days or less. The system displays the message during each sign-in attempt. The warning message contains the remaining number of days, hours, and minutes that the user has to change the password before it expires on the server. The default value is 14 days, but you can change it on the password configuration page of the admin console.

Troubleshooting LDAP Password Management

When you troubleshoot, provide any pertinent system logs, server logs, configuration information, and a TCP trace from the system. If you are using LDAPS, switch to the "Unencrypted" LDAP option LDAP server configuration while taking the LDAP TCP traces.

Related Documentation

- Using an LDAP Server

Configuring LDAP Search Attributes for Meeting Creators

Use options in the Meetings tab to specify individual LDAP attributes that a meeting creator may use to search for users when scheduling a meeting.

To configure Pulse Collaboration search attributes:

1. Select **Authentication > Auth. Servers**.

2. Click an LDAP server name (or create an LDAP server and then save it), and then choose the Meetings tab.

3. In the User Name field, enter the username attribute for this server. For example, enter SamAccountName for an Active Directory server or uid for an iPlanet server.

4. In the Email Address field, enter the e-mail attribute for this server.

5. In the Display Name, Attributes field, enter any additional LDAP attributes whose contents you want to allow meeting creators to view (optional). (For example, to help the meeting creator easily distinguish between multiple invitees with the same name, you may want to expose an attribute that identifies the departments of individual users.) Enter the additional attributes one per line using the format: DisplayName,AttributeName. You may enter up to 10 attributes.

6. Click **Save Changes**.

Related Documentation

- Pulse Collaboration Overview

Using an MDM Server

This topic describes integration with the mobile device management (MDM) servers. It includes the following information:

- Understanding MDM Integration
- Feature Support
- Configuring an MDM Server
- Display the Active Users Page

Understanding MDM Integration

MDM vendors provide enrollment and posture assessment services that prompt employees to enter data about their mobile devices. When the user installs the MDM application on the device and completes enrollment, the MDM pushes the device certificate to the device. After enrollment, the MDM maintains a database record that includes information about the enrollee—attributes related to device identity, user identity, and posture assessment against MDM policies.

The Pulse Secure access management framework MDM authentication server configuration determines includes details on how the system communicates with the MDM Web RESTful API service and how it derives the device identifier from the certificates presented by endpoints.

After you have configured the MDM authentication server, you can configure a realm that uses the MDM data for authorization, and you can use MDM device attributes in the role mapping rules that are the basis for your network access and resource access policies.

Feature Support

The Pulse Secure device access management framework supports integration with the following MDM solutions:

- AirWatch
- MobileIron

Configuring an MDM Server

The authentication server configuration is used by the system to communicate with the MDM. In the device access management framework, the MDM server is used as the device authorization server.

To configure the authentication server:

1. Select **Authentication > Auth Servers** to navigate to the authentication server configuration pages.

2. Select **MDM Server** and click **New Server** to display the configuration page.

3. Complete the configuration as described in Table 22.

4. Save the configuration.

Table22: Authentication Server Configuration Guidelines

Settings	Guidelines
Name	Specify a name for the configuration.
Type	Select one of the following options: - **AirWatch** - **MobileIron**
Server	
Server Url	Specify the URL for the MDM server. This is the URL the MDM has instructed you to use to access its RESTful Web API (also called a RESTful Web service). NOTE: You must configure your firewalls to allow communication between these two nodes over port 443.
Viewer Url	Specify the URL for the MDM report viewer. This URL is used for links from the Active Users page to the MDM report viewer.

Request Timeout	Specify a timeout period (0-60 seconds) for queries to the MDM server. The default is 15 seconds.
Administrator	
Username	Specify the username for an account that has privileges to access the MDM RESTful Web API.
Password	Specify the corresponding password.
Tenant Code	*AirWatch only*. Copy and paste the AirWatch API tenant code.
Device Identifier	
Device identity	*Policy Secure only*. Select an option on whether to require that the MDM certificate is presented by the endpoint when signing in: • **Require**—Require that the device certificate pushed to client devices during enrollment be used at sign-in. If this option is selected, and the client device does not have a certificate, authorization fails. Use this option when you require endpoints to adhere to your certificate security requirements. • **Use Certificate if present**—Use the certificate to derive the device ID if the certificate is presented at sign-in, but do not reject authentication if the certificate is not present. You can use this option in conjunction with a role mapping rule and a remediation VLAN to identify devices that have not perfected MDM enrollment. • **Always Use MAC address**—In some cases, the MDM certificate might be configured without a device identifier. When the endpoint uses an 802.1x framework to authenticate, Policy Secure can obtain the MAC address from the RADIUS return attribute callingStationID. The system can then use the MAC address as the device identifier. NOTE: This option is not present in Connect Secure. A device certificate is required to determine device identity.
ID Template	Construct a template to derive the device identifier from the certificate attributes. The template can contain textual characters as well as variables for substitution. The variables are the same as those used in role mapping custom expressions and policy conditions. Enclose variables in angle brackets like this <variable>. For example, suppose the certificate DN is: CN=<DEVICE_UDID>, uid=<USER_ID>, o=Company. With this configuration, the certificate could identify both the user and the device. In this example, the device ID template is <certDN.CN>.
ID Type	Select the device identifier type that matches the selection in the MDM SCEP certificate configuration: • **UUID**—The device Universal Unique Identifier. This is the key device identifier supported by MobileIron MDM. • **Serial Number**—The device serial number. • **UDID** —The device Unique Device Identifier. This is the key device identifier supported by AirWatch MDM.

Display the Active Users Page

The Active Users page lists data about current sessions, including access to realms that use the MDM server for device authorization.

To display the Active Users page, select **Systems > Active Users**.

Figure 24 shows the Active Users page for Policy Secure. The page for Connect Secure is similar.

Figure24: Active Users Page

☑	!	User ▼	Realm	Roles	Signed in	Signed in IP	MAC Address	Device Details	Agent Type	Agent Version	Endpoint Security Status
☐		admindb	Admin Users	.Administrators	2013/8/26 17:05:15	172.21.140.88			Windows 7 FireFox		**Not Applicable**
☐		atsang	Cert	iOS, Attack, Sunnyvale, Users	2013/8/26 17:08:09		24-ab-81-4e-d9-4c	🖥			▶ **Not Compliant - Not Running** (**Logs**)

Show users named: *   Show 200 users   Update

[Delete Session...] [Delete All Sessions...] [Refresh Roles] [Disable All Users...]

Number of Users: 2

TIP: Click the icon in the Device Details column to navigate to the MDM report viewer page for the device.

Related Documentation

- AAA Server Overview

- Understanding the Device Access Management Framework

Using an NIS Server

This topic describes integration with the NIS server. It includes the following information:

- NIS Server Overview

- Configuring Authentication with an NIS Server

- Displaying the User Accounts Table

NIS Server Overview

This section describes support for using Policy Secure and Connect Secure with the NIS server. It includes the following sections:

- Understanding NIS Server

- Feature Support

- Interoperability Requirements and Limitations

Understanding NIS Server

Network Information Service (NIS) is an authentication server that allows a central server to manage password authentication, hosts, services, and so on.

When you use an NIS server as the authentication and authorization service for your Pulse Secure access management framework, users can sign in to Policy Secure and Connect Secure using the same username and password that is used for the NIS server.

Feature Support

Pulse Secure access management framework supports the following NIS server features:

- Password management feature enables users who access an NIS server to manage their policies defined on the NIS server.

- Integrates NIS map data for passwords, groups, and hosts with corresponding objects in Active Directory.

- Allows migration of NIS domains to Active Directory.

Interoperability Requirements and Limitations

The following limitations apply when defining and monitoring an NIS server instance:

- You can only use NIS authentication with the system if your passwords are stored on the NIS server using Crypt or MD5 formats.

- You can only add one NIS server configuration to the system, but you can use that configuration to authenticate any number of realms.

- The username submitted to the system cannot contain two consecutive tilde symbols (~~).

Configuring Authentication with an NIS Server

To configure authentication with the NIS server:

1. Select **Authentication** > **Auth. Servers**.

2. Select **NIS Server** and click **New Server** to display the configuration page.

3. Complete the configuration as described in Table 23.

4. Save the configuration.

Table23: NIS Server Settings

Settings	Guidelines
Name	Specify a name to identify the server within the system.
NIS Server	Specify the name or IP address of the NIS server.
NIS Domain	Specify the domain name for the NIS server.
User Record Synchronization	This feature is available only on Connect Secure.
Enable User Record Synchronization	Select this option to retain the bookmarks and individual preferences regardless of which system you log in to.
Logical Auth Server Name	Specify a logical authentication server name.

Displaying the User Accounts Table

To display user accounts:

1. Select **Authentication > Auth. Servers**.

2. Click the link for the authentication server you want to manage.

3. Click the **Users** tab to display the user accounts table.

Figure 25 shows the Users page for Policy Secure. The Users page for Connect Secure is similar.

The user accounts table includes entries for the accounts that have been created. The Last Sign-in Statistic column shows the last successful sign-in date and time for each user, the user's IP address, and the agent or browser type and version.

4. Use the controls to search for users and manage user accounts:

- To search for a specific user, enter a username in the Show users named box and click **Update**.

  TIP: You can use an asterisk (*) as a wildcard, where * represents any number of zero or more characters. For example, to search for all usernames that contain the letters jo, enter ***jo***. The search is case-sensitive. To display the entire list of accounts again, type ***** or delete the field's contents and click **Update**.

- To limit the number of users displayed on the page, enter a number in the Show *N* users box and click **Update**.

- To terminate the user session and delete the account, select the check box next to the user account record and click **Delete**.

Figure25: User Accounts Table

Related Documentation

- AAA Server Overview

Using a RADIUS Server

This topic describes integration with the RADIUS server. It includes the following information:

- RADIUS Server Overview
- Configuring Authentication with a RADIUS Server
- Displaying the User Accounts Table

RADIUS Server Overview

This section describes support for using an external RADIUS server. It includes the following sections:

Understanding RADIUS Server

A Remote Authentication Dial-In User Service (RADIUS) server is a type of server that allows you to centralize authentication and accounting for users.

The following authentication schemes are supported:

- Access-Request–The user enters the username and password to request access to RADIUS server.

- Access-Accept–The user is authenticated.

- Access-Reject–The user is not authenticated and is prompted to reenter the username and password, or access is denied.

- Access-Challenge–A challenge is issued by the RADIUS server. The challenge collects additional data from the user.

Feature Support

Pulse Secure access management framework supports the following RADIUS features:

- RADIUS authentication.

- RADIUS attributes that can be used in role mapping.

- RADIUS directory services to retrieve user attributes in role-mapping rules.

- RADIUS accounting to track the services and the network resources used.

- RADIUS proxy to configure your external RADIUS server as an inner or outer proxy target. When you specify RADIUS proxy, some fields in the RADIUS server configuration page are not applicable. This feature is supported only on Policy Secure.

- RADIUS Disconnect messages. This feature is applicable for Connect Secure.

Using Challenge Expressions

The Pulse Secure access management framework supports the RSA Authentication Manager using the RADIUS protocol and a SecurID token (available from Security Dynamics). If you use SecurID to authenticate users, they must supply a user ID and the concatenation of a PIN and a token value.

When you define a RADIUS server, the Pulse Secure access management framework allows administrators to use hard-coded (default) challenge expressions that support Defender 4.0 and some RADIUS server implementations (such as Steel-Belted RADIUS and RSA RADIUS) or to enter custom challenge expressions that allow the system to work with many different RADIUS implementations and new versions of the RADIUS server, such as Defender 5.0. The system

looks for the response in the Access-Challenge packet from the server and issues an appropriate Next Token, New PIN, or Generic Passcode challenge to the user.

Using CASQUE Authentication

CASQUE authentication uses a token-based challenge/response authentication mechanism employing a CASQUE player installed on the client system. Once configured with CASQUE authentication, the RADIUS server issues a challenge with a response matching the custom challenge expression (:([0-9a-zA-Z/+=]+):). The system then generates an intermediate page that automatically launches the CASQUE player installed on the user's system.

PassGo Defender

If you are using a PassGo Defender RADIUS server, the user sign-in process is as follows:

1. The user is prompted for and enters a username and password.

2. The username and encrypted password are sent over the network to the RADIUS server.

3. The RADIUS server sends a unique challenge string to the system. The system displays this challenge string to the user.

4. The user enters the challenge string in a Defender token and the token generates a response string.

5. The user enters the response string on the system and clicks **Sign In**.

Using RADIUS Attributes

Table 24 describes the RADIUS attributes that are supported in RADIUS role-mapping.

Table24: RADIUS Attributes

Attribute	Description
ARAP-Challenge-Response	Contains the response to the challenge of a dial-in client. Sent in an Access-Accept packet with Framed-Protocol of ARAP.
ARAP-Features	Includes password information that the network access server (NAS) must send to the user in an ARAP feature flags packet. Sent in an Access-Accept packet with Framed- Protocol of ARAP.
ARAP-Password	Appears in an Access-Request packet containing a Framed-Protocol of ARAP. Only one of User-Password, CHAP-Password, or ARAP-Password must be included in an Access-Request, or one or more EAP-Messages.
ARAP-Security	Identifies the ARAP security module to be used in an Access-Challenge packet.
ARAP-Security-Data	Contains the actual security module challenge or response, and is in Access-Challenge and Access-Request packets.
ARAP-Zone-Access	Indicates how to use the ARAP zone list for the user.
Access-Accept	Provides specific configuration information necessary to begin delivery of service to the user.
Access-Challenge	Sends the user a challenge requiring a response, and the RADIUS server must respond to the Access-Request by transmitting a packet with the Code field set to 11 (Access-Challenge).
Access-Reject	Transmits a packet with the Code field set to 3 (Access-Reject) if any value of the received Attributes is not acceptable.

Access-Request	Conveys information specifying user access to a specific NAS, and any special services requested for that user.
Accounting-Request	Conveys information used to provide accounting for a service provided to a user.
Accounting-Response	Acknowledges that the Accounting-Request has been received and recorded successfully.
Acct-Authentic	Indicates how the user was authenticated, whether by RADIUS, the NAS itself, or another remote authentication protocol.
Acct-Delay-Time	Indicates how many seconds the client has been trying to send this record.
Acct-Input-Gigawords	Indicates how many times the Acct-Input-Octets counter has wrapped around 2^{32} over the course of this service being provided.
Acct-Input-Octets	Indicates how many octets have been received from the port during the current session.
Acct-Input-Packets	Indicates how many packets have been received from the port during the session provided to a Framed User.
Acct-Interim-Interval	Indicates the number of seconds between each interim update in seconds for this specific session.
Acct-Link-Count	Indicates the count of links known to have been in a given multilink session at the time the accounting record is generated.
Acct-Multi-Session-Id	Indicates a unique Accounting ID to make it easy to link together multiple related sessions in a log file.
Acct-Output-Gigawords	Indicates how many times the Acct-Output-Octets counter has wrapped around 2^{32} during the current session.
Acct-Output-Octets	Indicates how many octets have been sent to the port during this session.
Acct-Output-Packets	Indicates how many packets have been sent to the port during this session to a Framed User.
Acct-Session-Id	Indicates a unique Accounting ID to make it easy to match start and stop records in a log file.
Acct-Session-Time	Indicates how many seconds the user has received service.
Acct-Status-Type	Indicates whether this Accounting-Request marks the beginning of the user service (Start) or the end (Stop).
Acct-Terminate-Cause	Indicates how the session was terminated.
Acct-Tunnel-Connection	Indicates the identifier assigned to the tunnel session.
Acct-Tunnel-Packets-Lost	Indicates the number of packets lost on a given link.
CHAP-Challenge	Contains the Challenge Handshake Authentication Protocol (CHAP) challenge sent by the NAS to a PPP CHAP user.
CHAP-Password	Indicates the response value provided by a PPP CHAP user in response to the challenge.
Callback-Id	Indicates the name of a location to be called, to be interpreted by the NAS.
Callback-Number	The dialing string to be used for callback.
Called-Station-Id	Allows the NAS to send the phone number that the user called, using Dialed Number Identification Service (DNIS) or similar technology.
Calling-Station-Id	Allows the NAS to send the phone number that the call came from, using Automatic Number Identification (ANI) or similar technology.
Class	Sent by the server to the client in an Access-Accept and then sent unmodified by the client to the accounting server as part of the Accounting-Request packet, if accounting is supported.
Configuration-Token	Used in large distributed authentication networks based on proxy.
Connect-Info	Sent from the NAS to indicate the nature of the user's connection.

Event-Timestamp	Records the time that this event occurred on the NAS, in seconds since January 1, 1970 00:00 UTC.
Filter-Id	Indicates the name of the filter list for this user.
Framed-AppleTalk-Link	Indicates the AppleTalk network number used for the serial link to the user, which is another AppleTalk router.
Framed-AppleTalk-Network	Indicates the AppleTalk Network number which the NAS can probe to allocate an AppleTalk node for the user.
Framed-AppleTalk-Zone	Indicates the AppleTalk Default Zone to be used for this user.
Framed-Compression	Indicates the compression protocol to be used for the link.
Framed-IP-Address	Indicates the address to be configured for the user.
Framed-IP-Netmask	Indicates the IP netmask to be configured for the user when the user is a router to a network.
Framed-IPv6-Pool	Contains the name of an assigned pool used to assign an IPv6 prefix for the user.
Framed-IPv6-Prefix	Indicates an IPv6 prefix (and corresponding route) to be configured for the user.
Framed-IPv6-Route	Indicates the routing information to be configured for the user on the NAS.
Framed-Interface-Id	Indicates the IPv6 interface identifier to be configured for the user.
Framed-IPX-Network	Indicates the IPX Network number to be configured for the user.
Framed-MTU	Indicates the maximum transmission unit to be configured for the user, when it is not negotiated by some other means (such as PPP).
Framed-Pool	Indicates the name of an assigned address pool used to assign an address for the user.
Framed-Protocol	Indicates the framing to be used for framed access.
Framed-Route	Indicates the routing information to be configured for the user on the NAS.
Framed-Routing	Indicates the routing method for the user, when the user is a router to a network.
Idle-Timeout	Sets the maximum number of consecutive seconds of idle connection allowed to the user before termination of the session or prompt.
Keep-Alives	Uses SNMP instead of keepalives.
Login-IP-Host	Indicates the system with which to connect the user when the Login-Service Attribute is included.
Login-IPv6-Host	Indicates the system with which to connect the user when the Login-Service Attribute is included.
Login-LAT-Group	Contains a string identifying the LAT group codes that this user is authorized to use.
Login-LAT-Node	Indicates the node with which the user is to be automatically connected by LAT.
Login-LAT-Port	Indicates the port with which the user is to be connected by LAT.
Login-LAT-Service	Indicates the system with which the user is to be connected by LAT.
Login-Service	Indicates the service to use to connect the user to the login host.
Login-TCP-Port	Indicates the TCP port with which the user is to be connected when the Login-Service Attribute is also present.
MS-ARAP-Challenge	Only present in an Access-Request packet containing a Framed-Protocol Attribute with the value 3 (ARAP).
MS-ARAP-Password-Change-Reason	Indicates the reason for a server-initiated password change.
MS-Acct-Auth-Type	Represents the method used to authenticate the dial-up user.
MS-Acct-EAP-Type	Represents the Extensible Authentication Protocol (EAP) type used to authenticate the dial-up user.

MS-BAP-Usage	Describes whether the use of BAP is allowed, disallowed, or required on new multilink calls.
MS-CHAP-CPW-1	Allows the user to change password if it has expired.
MS-CHAP-CPW-2	Allows the user to change password if it has expired.
MS-CHAP-Challenge	Contains the challenge sent by a NAS to a MS-CHAP user.
MS-CHAP-Domain	Indicates the Windows NT domain in which the user was authenticated.
MS-CHAP-Error	Contains error data related to the preceding MS-CHAP exchange.
MS-CHAP-LM-Enc-PW	Contains the new Windows NT password encrypted with the old LAN Manager password hash.
MS-CHAP-MPPE-Keys	Contains two session keys for use by the Microsoft Point-to-Point Encryption (MPPE).
MS-CHAP-NT-Enc-PW	Contains the new Windows NT password encrypted with the old Windows NT password hash.
MS-CHAP-Response	Contains the response value provided by a PPP MS-CHAP user in response to the challenge.
MS-CHAP2-CPW	Allows the user to change password if it has expired.
MS-CHAP2-Response	Contains the response value provided by an MS- CHAP-V2 peer in response to the challenge.
MS-CHAP2-Success	Contains a 42-octet authenticator response string.
MS-Filter	Transmits traffic filters.
MS-Link-Drop-Time-Limit	Indicates the length of time (in seconds) that a link must be underutilized before it is dropped.
MS-Link-Utilization-Threshold	Represents the percentage of available bandwidth utilization below which the link must fall before the link is eligible for termination.
MS-MPPE-Encryption-Policy	Signifies whether the use of encryption is allowed or required.
MS-MPPE-Encryption-Types	Signifies the types of encryption available for use with MPPE.
MS-MPPE-Recv-Key	Contains a session key for use by the MPPE.
MS-MPPE-Send-Key	Contains a session key for use by the MPPE.
MS-New-ARAP-Password	Transmits the new ARAP password during an ARAP password change operation.
MS-Old-ARAP-Password	Transmits the old ARAP password during an ARAP password change operation.
MS-Primary-DNS-Server	Indicates the address of the primary domain name server (DNS) server to be used by the PPP peer.
MS-Primary-NBNS-Server	Indicates the address of the primary NetBIOS name server (NBNS) server to be used by the PPP peer.
MS-RAS-Vendor	Indicates the manufacturer of the RADIUS client machine.
MS-RAS-Version	Indicates the version of the RADIUS client software.
MS-Secondary-DNS-Server	Indicates the address of the secondary DNS server to be used by the PPP peer.
MS-Secondary-NBNS-Server	Indicates the address of the secondary DNS server to be used by the PPP peer.
Message-Authenticator	Signs Access-Requests to prevent spoofing Access-Requests using CHAP, ARAP, or EAP authentication methods.

NAS-IP-Address	Indicates the identifying IP address of the NAS that is requesting authentication of the user, and must be unique to the NAS within the scope of the RADIUS server.
NAS-IPv6-Address	Indicates the identifying IPv6 Address of the NAS that is requesting authentication of the user, and must be unique to the NAS within the scope of the RADIUS server.
NAS-Identifier	Contains a string identifying the NAS originating the Access-Request.
NAS-Port	Indicates the physical port number of the NAS that is authenticating the user.
NAS-Port-Id	Contains a text string that identifies the port of the NAS that is authenticating the user.
NAS-Port-Type	Indicates the type of the physical port of the NAS that is authenticating the user.
Password-Retry	Indicates how many authentication attempts a user is allowed to attempt before being disconnected.
Port-Limit	Sets the maximum number of ports to be provided to the user by the NAS.
Prompt	Indicates to the NAS whether it should echo the user's response as it is entered, or not echo it.
Proxy-State	Indicates that a proxy server can send this attribute to another server when forwarding an Access-Request. The attribute must be returned unmodified in the Access-Accept, Access-Reject or Access-Challenge.
Reply-Message	Indicates that the text that can be displayed to the user.
Service-Type	Indicates the type of service the user has requested, or the type of service to be provided.
Session-Timeout	Sets the maximum number of seconds of service to be provided to the user before termination of the session or prompt.
State	Indicates that the packet must have only zero or one State Attribute. Usage of the State Attribute is implementation dependent.
Telephone-number	Using the Calling-Station-Id and Called-Station-Id RADIUS attributes, authorization and subsequent tunnel attributes can be based on the phone number originating the call, or the number being called.
Termination-Action	Indicates the action the NAS should take when the specified service is completed.
Tunnel-Assignment-ID	Indicates to the tunnel initiator the particular tunnel to which a session is to be assigned.
Tunnel-Client-Auth-ID	Specifies the name used by the tunnel initiator during the authentication phase of tunnel establishment.
Tunnel-Client-Endpoint	Contains the address of the initiator end of the tunnel.
Tunnel-Link-Reject	Indicates the rejection of the establishment of a new link in an existing tunnel.
Tunnel-Link-Start	Marks the creation of a tunnel link.
Tunnel-Link-Stop	Marks the destruction of a tunnel link.
Tunnel-Medium-Type	Indicates the transport medium to use when creating a tunnel for those protocols (such as L2TP) that can operate over multiple transports.
Tunnel-Medium-Type	Indicates the transport medium to use when creating a tunnel for those protocols (such as L2TP) that can operate over multiple transports.
Tunnel-Password	Specifies a password used to access a remote server.
Tunnel-Preference	Indicates that if RADIUS server returns more than one set of tunneling attributes to the tunnel initiator, you should include this attribute in each set to indicate the relative preference assigned to each tunnel.
Tunnel-Private-Group-ID	Indicates the group ID for a particular tunneled session.
Tunnel-Reject	Marks the rejection of the establishment of a tunnel with another node.
Tunnel-Server-Auth-ID	Specifies the name used by the tunnel terminator during the authentication phase of tunnel establishment.

Tunnel-Server-Endpoint	Indicates the address of the server end of the tunnel.
Tunnel-Start	Marks the establishment of a tunnel with another node.
Tunnel-Stop	Marks the destruction of a tunnel to or from another node.
Tunnel-Type	Indicates the tunneling protocol(s) to be used (in the case of a tunnel initiator) or the tunneling protocol in use (in the case of a tunnel terminator).
User-Name	Indicates the name of the user to be authenticated.
User-Password	Indicates the password of the user to be authenticated, or the user's input following an Access-Challenge.

Understanding RADIUS Accounting

You can configure the device to send session start and stop messages to a RADIUS accounting server. The device sends a user-session start message after the user successfully signs in and the device maps to a role.

Whenever a user session is terminated, the device sends a user-session stop message to the accounting server. A user session is terminated whenever the user:

- Manually signs out

- Times out because of either inactivity or exceeding the maximum session length

- Is denied access because of Host Checker role-level restrictions

- Is manually forced out by an administrator as a result of dynamic policy evaluation

> NOTE: If users are signed into a device cluster, the RADIUS accounting messages might show the users signing in to one node and signing out of another.

Table 25 describes the attributes that are common to start and stop messages. Table 26 describes the attributes that are unique to start messages. Table 27 describes the attributes that are unique to stop messages.

Table25: Attributes Common to Start and Stop Messages

Attribute	Description
User-Name (1)	Specifies the string that the device administrator specifies during RADIUS server configuration.
NAS-IP-Address (4)	Specifies the device's IP address.
NAS-Port (5)	The device sets this attribute to 0 if the user signed in using an internal port, or 1 if an external port is used.
Framed-IP-Address (8)	Specifies the user's source IP address.
NAS-Identifier (32)	Specifies the configured name for the device client under the RADIUS server configuration.
Acct-Status-Type (40)	The device sets this attribute to 1 for a start message, or 2 for a stop message in a user-session or a subsession.

Acct-Session-Id (44)	Specifies the unique accounting ID that matches start and stop messages corresponding to a user-session or to a subsession.
Acct-Multi-Session-Id (50)	Specifies the unique accounting ID that you can use to link together multiple related sessions. Each linked session must have a unique Acct-Session-Id and the same Acct-Multi-Session-Id.
Acct-Link-Count (51)	Specifies the count of links in a multilink session at the time the system generates the accounting record.

Table26: Start Attributes

Attribute	Description
Acct-Authentic (45)	The device sets this attribute to: • RADIUS—if the user is authenticated to a RADIUS server. • Local—if the user is authenticated to a local authentication server. • Remote—if the user is authenticated through any other RADIUS server.

Table27: Stop Attributes

Attribute	Description
Acct-Session-Time (46)	Specifies the duration of the user-session or the subsession.
Acct-Terminate-Cause (49)	The device uses one of the following values to specify the event that caused the termination of a user session or a subsession: • User Request (1) – User manually signs out. • Idle Timeout (4) – User is idle and times out. • Session Timeout (5) – User's maximum session times out. • Admin Reset (6) – User is forced out from active users page.

Interoperability Requirements and Limitations

You must configure the third-party RADIUS server to communicate with the Pulse Secure access management framework.

On the RADIUS server, configure the following settings:

• Hostname.

• Network IP address.

• Client type, if applicable. If this option is available, select **Single Transaction Server** or its equivalent.

• Type of encryption for authenticating client communication. This choice should correspond to the client type.

• Shared secret.

The following are the requirements and limitations for Interim update feature:

- If you want a server to receive interim accounting messages, you can statically configure an interim value on the client, in which case, the locally configured value overrides any value that might be included in the RADIUS Access-Accept message.

- The octet count reported in the accounting messages is the cumulative total since the beginning of the user session.

- The interim update byte count is only supported based on a user session, not on SAM or NC sessions.

Configuring Authentication with a RADIUS Server

To configure authentication with the RADIUS server:

1. Select **Authentication** > **Auth. Servers**.

2. Select **RADIUS Server** and click **New Server** to display the configuration page.

3. Complete the configuration as described in Table 28.

4. Save the configuration.

Table28: RADIUS Server Settings

Settings	Guidelines
Name	Specify a name to identify the server within the system.
NAS-Identifier	Specify the name that identifies the Network Access Server (NAS) client to the RADIUS server. NOTE: • If you do not specify the NAS identifier, the value specified in the Hostname field on the System > Network > Overview page of the administrator console is used. • If you use the RADIUS proxy feature, the NAS-Identifier field is not used. Proxy passes on the entire RADIUS packet including the NAS identifier from the client.
Primary Server	
Radius Server	Specify the name or IP address of the RADIUS server.
Authentication Port	Specify the authentication port value for the RADIUS server. Default port number: 1812, 1645 (legacy servers)
NAS-IP-Address	Specify the NAS IP address. NOTE: • If you leave this field empty, the internal IP address is passed to RADIUS requests. • If you configure the NAS IP address, then the system passes the value regardless of which cluster node sends the requests. • If you use the RADIUS proxy feature, this field is not used. Proxy passes on the entire RADIUS packet including the NAS IP address from the client.
Timeout (seconds)	Specify the interval of time to wait for a response from the RADIUS server before timing out the connection.
Retries	Specify the number of times to try to make a connection after the first attempt fails.

Users authenticate using tokens or one-time passwords.	Select this option to prompt the user for a token instead of a password.
	For example, you can use this option to dynamically prompt for a password or token based on sign-in policies by configuring two instances of the same authentication server. You can use one instance for wireless users with this option enabled and that prompts the user for a token, and another instance for wired users with this option disabled and that prompts the user for a password.
	NOTE: If you are using RADIUS proxy feature, this option is not used.
Backup Server (required only if Backup server exists)	
Radius Server	Specify the secondary RADIUS server.
	The authentication request is first routed to the primary RADIUS server, then to the specified backup server if the primary server is unreachable.
	Accounting messages are sent to the RADIUS server by each cluster node without consolidation.
	RADIUS accounting follows these assumptions:
	• If the cluster is active/passive, all users are connected to one node at a time.
	• If the cluster is active/active and does not use a balancer, users are connected to different nodes but are static.
	• If the cluster is active/active and uses a balancer, the balancer usually enforces a persistent source IP. In this case, users are always connected to the same node.
	NOTE: RADIUS does not support load balancing.
Authentication Port	Specify the authentication port.
Shared Secret	Specify the shared secret.
Accounting Port	Specify the accounting port.
Radius Accounting	
User-Name	Specify the user information to the RADIUS accounting server.
	You can enter any of the applicable session variables. Applicable variables include those that are set the time after the user signs in and maps to a role.
	The default variables for this field are as follows:
	• USER: Logs the username to the accounting server.
	• REALM: Logs the realm to the accounting server.
	• ROLE SEP=",": Logs the list of comma-separated roles assigned to the user.
	• ROLE: Logs the role to the accounting server.
	NOTE: If you assign the user to more than one role, the system separates them with commas.
Interim Update Interval (minutes)	Select this option to achieve more precise billing for long-lived session clients and during network failure.
	NOTE:
	• If you are using the RADIUS proxy feature, the fields in this section are not used.
	• The minimum interim update interval is 15 minutes. The data statistics (bytes in and bytes out) for RADIUS accounting might not be sent for a J-SAM/W-SAM/NC session if the session is less than 30 seconds long and the applications keep the connections open all the time.
Use NC assigned IP Address for FRAMED-IP-ADDRESS attribute value in Radius Accounting	Select the Use NC assigned IP Address for FRAMED-IP-ADDRESS attribute value in Radius Accounting check box to use the IP address returned from Connect Secure for the Framed-IP-Address attribute. Two IP addresses are recorded: one prior to authenticating with Connect Secure, and one returned by VPN Tunneling after authentication. Select this option to use the VPN Tunneling IP address for the Framed-IP-Address attribute instead of the pre-authenticated (original) IP address.

Radius Disconnect

This feature is applicable for Connect Secure

Enable processing of Radius Disconnect Requests	Select this option to process Radius Disconnect Requests. The Radius Disconnect requests received from the backend Radius server will terminate sessions that match the attributes in the request. NOTE: You must not configure multiple RADIUS authentication servers with the same backend server details. The Radius attributes that are used for session identification are: • Framed-IP-Address (for sessions with VPN Tunnel only) • Acct-Session-Session-Id • Acct-Multi-Session-Id • User-Name

Custom challenge expressions

This feature is applicable for Policy Secure

(Optional) Three types of challenge expressions exist with each automatically set to its pre-populated default. The custom option allows the administrator to configure the actual string pattern to match for any of the three modes. To add a custom expression, select the check box for the appropriate challenge expression type, and add a custom expression in the associated text box.

NOTE:

• If you use SecureID to authenticate users, then provide the user ID and the concatenation of PIN and the token value.
• When using CASQUE authentication, specify:([0-9a-zA-Z/+=]+): as the custom expression for the Generic Login Challenge Expression.
• If you are using the RADIUS proxy feature, the fields in this section are not used.

Next Token	Specify the appropriate Next Token.
New PIN	Specify the New PIN.
Generic Login	Specify the Generic Login challenge to the user.

Custom Radius Rules

This feature is applicable for Connect Secure

(Optional) Click **New Radius Rule** to add a custom challenge rule that determines the action to take for an incoming packet. When a user enters his or her username and password, the initial authorization request is sent to the server. The server may respond with a Challenge or Reject packet. In the **Add Custom Radius Challenge Rule** window, you select the packet type (Challenge or Reject) and then specify what action to take. For example, you can show a login page with a specific error message to the user, or automatically send an ACCESS-REQUEST packet back to the server.

To create a custom challenge rule:

1. Select the incoming packet type:
   - Access Challenge—sent by the RADIUS server requesting more information in order to allow access
   - Access Reject—sent by the RADIUS server rejecting access

2. Specify an expression to evaluate, based on the Radius attribute, and click Add. If you specify more than one expression, the expressions are "ANDed" together. To remove an expression, click the delete icon next to the expression.

3. Choose the action to take by selecting one of the following radio buttons:
   - show NEW PIN page—user must enter a new PIN for the token
   - show NEXT TOKEN page—user must enter the next tokencode
   - show GENERIC LOGIN page—display an additional page to the user in response to an Access Challenge sent by the server. Sometimes a Radius server returns a Challenge packet and requires the user to enter additional information to continue the login process. For example, a server receives the initial username and password and sends an SMS message to the user's mobile phone with a one-time password (OTP). The user enters the OTP in the generic login page.
   - show user login page with error—display the standard login page with an embedded error message. This option lets you bypass the standard message string sent by Connect Secure and display a custom error message to the user. Enter your custom message in the Error Message text box. There is no maximum character limit for this message.
   - send ACCESS REQUEST with additional attributes—send an ACCESS-REQUEST packet with the specified attribute/value pair(s). Select an attribute, enter its value and click Add. To delete an attribute, click the delete icon next to the attribute/value pair. You must set User-Password to <PASSWORD> otherwise an "Invalid username or password" message appears.

4. Click **Save Changes** to save your edits, then click Close to close this window.

   Your custom rules appear in the table under the Custom Radius Authentication Rule section. To delete a rule, select the check box next to the rule and click Delete.

Displaying the User Accounts Table

To display user accounts:

1. Select **Authentication** > **Auth. Servers**.

2. Click the link for the authentication server you want to manage.

3. Click the **Users** tab to display the user accounts table.

   The user accounts table includes entries for the accounts that have been created. The Last Sign-in Statistic column shows the last successful sign-in date and time for each user, the user's IP address, and the agent or browser type and version.

4. Use the controls to search for users and manage user accounts:

   - To search for a specific user, enter a username in the Show users named box and click **Update**.

     TIP: You can use an asterisk (*) as a wildcard, where * represents any number of zero or more characters. For example, to search for all usernames that contain the letters jo, enter ***jo***. The search is case-sensitive. To display the entire list of accounts again, type ***** or delete the field's contents and click **Update**.

- To limit the number of users displayed on the page, enter a number in the Show *N* users box and click **Update**.

- To terminate the user session and delete the account, select the check box next to the user account record and click **Delete**.

Related Documentation

- AAA Server Overview

Using an ACE Server

This topic describes integration with an ACE Server (now named RSA Authentication Manager). It includes the following information:

- RSA Authentication Manager Overview

- Configuring Authentication with RSA Authentication Manager

- Displaying the User Accounts Table

RSA Authentication Manager Overview

This section describes support for using Policy Secure and Connect Secure with an ACE Server (now named RSA Authentication Manager). It includes the following sections:

- Understanding RSA Authentication Manager

- Feature Support

- Interoperability Requirements and Limitations

Understanding RSA Authentication Manager

RSA Authentication Manager (formerly known as ACE/Server) is an authentication and authorization server that allows user authentication based on credentials from the RSA SecurID® product from RSA Security Inc.

When you use RSA Authentication Manager as the authentication and authorization service for your Pulse Secure access management framework, users can sign in to Policy Secure and Connect Secure using the same username and password stored in the backend server.

Table 29 describes RSA SecurID hardware token and software token user sign-in methods.

Table29: Sign-in Methods

Method	Action
Using a hardware token and the standard system sign-in page	The user browses to the standard system sign-in page, and then enters the username and password (consisting of the concatenation of the PIN and the RSA SecurID hardware token's current value). The system then forwards the user's credentials to the authentication server.
Using a software token and the custom SoftID system sign-in page	The user browses to the SoftID custom sign-in page. Then, using the SoftID plug-in, the user enters the username and PIN. The SoftID plug-in generates a passphrase by concatenating the user's PIN and token and passes the passphrase to the authentication server.

If the RSA Authentication Manager positively authenticates the user, the user gains access to the system. Otherwise, the RSA Authentication Manager:

- Denies the user access to the system.

- Prompts the user to generate a new PIN (New PIN mode) if the user is signing in to the system for the first time. Users see different prompts depending on the method they use to sign in.

  If the user signs in using the SoftID plug-in, then the RSA prompts the user to create a new pin; otherwise Policy Secure and Connect Secure prompt the user to create a new PIN.

- Prompts the user to enter the next token (Next Token mode) if the token entered by the user is out of sync with the token expected by RSA Authentication Manager. Next Token mode is transparent to users signing in using a SoftID token. The RSA SecurID software passes the token through the system to RSA Authentication Manager without user interaction.

- Redirects the user to the standard system sign-in page (SoftID only) if the user tries to sign-in to the RSA SecurID Authentication page on a computer that does not have the SecurID software installed.

Feature Support

Pulse Secure access management framework supports the following RSA Authentication Manager features:

- New PIN mode

- Next-token mode

- Data Encryption Standard (DES)/ Secure Dial-In (SDI) encryption

- Advanced Encryption Standard (AES) encryption

- Slave Authentication Manager support

- Name locking

- Clustering

Interoperability Requirements and Limitations

The following limitations apply when defining and monitoring an RSA Authentication Manager instance:

- You can only add one RSA Authentication Manager configuration to the system, but you can use that configuration to authenticate any number of realms.

- You cannot customize the load balancing algorithm.

- When you enter the New PIN or Next Token mode, enter the required information within three minutes. Otherwise, the system cancels the transaction and notifies the user to reenter the credentials.

- The system can handle a maximum of 200 RSA Authentication Manager transactions at any given time. A transaction only lasts as long as is required to authenticate against the RSA Authentication Manager.

For example, when a user signs into the system, the RSA Authentication Manager transaction is initiated when the user submits the request for authentication and ends once the RSA Authentication Manager has finished processing the request. The user may then keep his or her session open, even though the RSA Authentication Manager transaction is closed.

Configuring Authentication with RSA Authentication Manager

To configure authentication with an ACE server:

1. Select **Authentication** > **Auth. Servers**.

2. Select **ACE Server** and click **New Server** to display the configuration page.

3. Complete the configuration as described in .

4. Save the configuration.

Table30: ACE Server Settings

Settings	Guidelines
Name	Specify a name to identify the server within the system.
ACE Port	Specify the default port of the authentication server.
	NOTE: If no port is specified in the sdconf.rec file, the default port is used.
Configuration File	
Current config file	Specify the RSA Authentication Manager configuration file.
	NOTE: You must update this file on the device anytime you make changes to the source file.
Imported on	Display the date on which the config file is imported.
Import new config file	Use the Choose File button to upload the sdconf.rec configuration file.
Node Verification File	
Node	Save the configuration to redisplay the configuration page. The updated page includes a section that lists a timestamp for the negotiation of the node secret between the system and the backend RSA server. The negotiation and verification automatically occurs after first successful login. Do not expect entries in the table until at least one user has authenticated successfully.
User Record Synchronization	This feature is available only on Connect Secure.
Enable User Record Synchronization	Select this option to retain the bookmarks and individual preferences regardless of which system you log in to.
Logical Auth Server Name	Specify a logical authentication server name.

Displaying the User Accounts Table

To display user accounts:

1. Select **Authentication > Auth. Servers**.

2. Click the link for the authentication server you want to manage.

3. Click the **Users** tab to display the user accounts table.

   The user accounts table includes entries for the accounts that have been created. The Last Sign-in Statistic column shows the last successful sign-in date and time for each user, the user's IP address, and the agent or browser type and version.

4. Use the controls to search for users and manage user accounts:

   - To search for a specific user, enter a username in the Show users named box and click **Update**.

     TIP: You can use an asterisk (*) as a wildcard, where * represents any number of zero or more characters. For example, to search for all usernames that contain the letters jo, enter ***jo***. The search is case-sensitive. To display the entire list of accounts again, type ***** or delete the field's contents and click **Update**.

   - To limit the number of users displayed on the page, enter a number in the Show *N* users box and click **Update**.

   - To terminate the user session and delete the account, select the check box next to the user account record and click **Delete**.

Related Documentation

   - AAA Server Overview

   - How to: Connect Secure: Implementing Authentication with RSA SecurID

   - KB: Troubleshooting Error AD-37 while importing RSA sdconf.rec

   - KB: Seamless authentication integration with RSA SecurID

Using the SAML Server

This topic describes the local SAML authentication server. It includes the following information:

   - SAML Server Overview

   - Configuring Authentication with the SAML Server

   - Displaying the User Accounts Table

SAML Server Overview

This section describes support for using the local Connect Secure SAML authentication server. It includes the following sections:

   - Understanding SAML

   - SAML Feature Support

   - Interoperability Requirements and Limitations

Understanding SAML

SAML is an XML-based framework for communicating user authentication, entitlement, and attribute information. The standard defines the XML-based assertions, protocols, bindings, and profiles used in communication between SAML entities. SAML is used primarily to implement Web browser single sign-on (SSO). SAML enables businesses to leverage an identity-based security system like Connect Secure to enforce secure access to web sites and other resources without prompting the user with more than one authentication challenge.

For complete details on the SAML standard, see the OASIS web site:

http://www.oasis-open.org/committees/tc_home.php?wg_abbrev=security

SAML Feature Support

When deployed as SAML service provider, Connect Secure runs a local SAML server that relies on the SAML identity provider authentication and attribute assertions when users attempt to sign in to Connect Secure. Note that authentication is only part of the Connect Secure security system. The access management framework determines access to the system and protected resources.

Connect Secure supports:

- HTTP Redirect binding for sending AuthnRequests

- HTTP Redirect binding for sending/receiving SingleLogout requests/responses

- HTTP POST and HTTP Artifact bindings for receiving SAML responses

- RequestedAuthnContext context class specifications

Interoperability Requirements and Limitations

Before you begin:

- Check to see whether the SAML identity provider implements SAML 2.0 or SAML 1.1.

- Check to see whether the SAML identity provider uses HTTP POST or HTTP Artifact bindings for SAML assertions.

- Check to see whether the SAML identity provider has published a SAML metadata file that defines its configuration. If the SAML identity provider metadata file is available, configuration is simpler and less prone to error.

- Complete the system-wide SAML settings if you have not already done so. Select **System > Configuration > SAML > Settings**. For details, see "Configuring Global SAML Settings".

- Add metadata for the SAML identity provider to the metadata provider list if you have not already done so. Select **System > Configuration > SAML**. For details, see "Managing SAML Metadata Files".

The sign-in URL for which a session needs to be established for Connect Secure as a service provider is identified by the RelayState parameter (HTTP URL parameter for artifact and HTML form parameter for POST.) In a service provider initiated case, the system populates RelayState as an HTTP URL parameter while sending AuthnRequest. In the IdP-Initiated scenario (Connect Secure is a service provider and there is a third-party IdP), the IdP must be configured to set the appropriate Sign-in URL of Connect Secure in the RelayState parameter

of the HTML form containing the SAML response. For more information, see the SAML 2.0 specification.

Configuring Authentication with the SAML Server

To configure the SAML server:

1. Select **Authentication > Auth. Servers**.

2. Select **SAML Server** and click **New Server** to display the configuration page.

3. Complete the configuration as described in .

4. Save the configuration.

Table31: SAML Service Provider Profile

Settings	Guidellines
Name	Specify a name to identify the server instance.
Settings	
SAML Version	Select **2.0** or **1.1**, depending on the SAML version used by the SAML IdP.
SA Entity Id	This value is prepopulated. It is generated by the system, based on the value for the Host FQDN for SAML setting on the System > Configuration > SAML > Settings page.
Configuration Mode	Select **Manual** or **Metadata**. If a metadata file or location is available from the SAML identity provider, use the metadata option to make configuration simpler and less prone to error. To upload or set the location for the published metadata file, select **System > Configuration > SAML** and click the **New Metadata Provider** button.
Identity Provider Entity ID	The identity provider entity ID is sent as the Issuer value in the assertion generated by the SAML identity provider.
	If you use the metadata option, this setting can be completed by selecting the identity provider entity ID from the list. The list is populated by the identity provider entities defined in metadata files added to the System > Configuration > SAML page.
	If you complete this setting manually, specify the Issuer value in assertions generated by the SAML identity provider. Typically, you ask the SAML identity provider administrator for this setting.
Identity Provider Single Sign On Service URL	The identity provider SSO service URL is a URL provisioned by the SAML identity provider. The setting is required to support service-provider-initiated SSO. If missing, the system cannot successfully redirect the user request.
	If you use the metadata option, this setting can be completed by selecting the SSO service URL from the list. The list is populated by the identity provider entities defined in metadata files added to the System > Configuration > SAML page.
	If you complete this setting manually, ask the SAML identity provider administrator for this setting.

User Name Template	Specify how the system is to derive the username from the assertion. If the field is left blank, it uses the string received in the NameID field of the incoming assertion as the username.
	If you choose a certificate attribute with more than one value, the system uses the first matched value. For example, if you enter <certDN.OU> and the user has two values for the attribute (ou=management, ou=sales), the system uses "management". To use all values, add the SEP attribute to the variable. For example, if you enter <certDN.OUT SEP=":">, the system uses "management:sales". The attributes received in the attribute statement in the incoming assertion are saved under userAttr. These variables can also be used with angle brackets and plain text. If the username cannot be generated using the specified template, the login fails. If the NameID filed of the incoming assertion is of type X509Nameformat, then the individual fields can be extracted using system variable "assertionNameDN".
Allowed Clock Skew (minutes)	Specify the maximum allowed difference in time between the system clock and the SAML identity provider server clock.
	NOTE: SAML is a time sensitive protocol. The time-based validity of a SAML assertion is determined by the SAML identity provider. If the SAML identity provider and SAML service provider clocks are askew, the assertion can be determined invalid, and you will receive the following error:
	"SAML Transferred failed. Please contact your system administrator. Detail: Failure: No valid assertion found in SAML response."
	We recommend you use NTP to ensure the clocks are synchronized and that you set an Allowed Clock Skew value that accommodates any expected or permissible skew.
Support Single Logout	Single logout is a mechanism provided by SAML for logging out a particular user from all the sessions created by the identity provider. Select this option if the system must receive and send a single logout request for the peer SAML identity provider.
	If you use the metadata option, the Single Logout Service URL setting can be completed by selecting the SLO service URL from the list. The list is populated by the identity provider entities defined in metadata files added to the System > Configuration > SAML page. The system sends Single Logout requests to this URL.
	In addition, if you use the metadata option, the Single Logout Response URL setting is completed based on your selection for Single Logout Service URL. If the identity provider has left this setting empty in its metadata file, the system sends the Single Logout response to the SLO service URL.
	If you complete these settings manually, ask the SAML identity provider administrator for guidance.
	The Support Single Logout service for the identity provider must present a valid certificate.

SSO Method

Artifact	When configured to use the Artifact binding, the system contacts the Artifact Resolution Service (ARS) to fetch the assertion using SOAP protocol. If the ARS is hosted on a HTTPS URL, then the certificate presented by the ARS is verified by the system. For this verification to pass successfully, the CA of the server certificate issued to the identity provider ARS must be added to the trusted server CA on the system. Complete the following settings to configure SAML using the HTTP Artifact binding: • Source ID. Enter the source ID for the identity provider ARS. Source ID is Base64-encoded, 20-byte identifier for the identity provider ARS. If left blank, this value is generated by the system. • Source Artifact Resolution Service URL. For metadata-based configuration, this field is completed automatically from the metadata file and is not configurable. For manual configurations, enter the URL of the service to which the SP ACS is to send ArtifactResolve requests. ArtifactResolve requests are used to fetch the assertion from the artifact received by it. • SOAP Client Authentication. Select **HTTP Basic** or **SSL Client Certificate** and complete the related settings. If you use an SSL client certificate, select a certificate from the device certificate list. • Select Device Certificate for Signing. Select the device certificate the system uses to sign the AuthnRequest sent to the identity provider SSO service. If you do not select a certificate, the system does not sign AuthnRequest. • Select Device Certificate for Encryption. Select the device certificate the system uses to decrypt encrypted data received in the SAML response. The public key associated with the device certificate is used by the identity provider for encryption.
POST	When configured to use the POST binding, the system uses a response signing certificate to verify the signature in the incoming response or assertion. The certificate file must be in PEM or DER format. The certificate you select should be the same certificate used by the identity provider to sign SAML responses. Complete the following settings to configure SAML using the HTTP POST binding: • Response Signing Certificate. If you use the metadata-based configuration option, select a certificate from the list. The list is populated by the identity provider entities defined in metadata files added to the System > Configuration > SAML page. If you configure these settings manually, browse to and upload the certificate to be used to validate the signature in the incoming response or assertion. If no certificate is specified, the certificate embedded in the response is used. • Enable Signing Certificate status checking. Select this option to check the validity of the signing certificate before verifying the signature. This setting applies to any certificate used for signature verification. If this option is enabled, the response will be rejected if the certificate is revoked, expired, or untrusted. If this option is selected, the certificate CA must be added to the Trusted Client CA store. If this option is not enabled then the certificate is used without any checks. • Select Device Certificate for Signing. Select the device certificate the system uses to sign the AuthnRequest sent to the identity provider SSO service. If you do not select a certificate, the system does not sign AuthnRequest. • Select Device Certificate for Encryption. Select the device certificate the system uses to decrypt encrypted data received in the SAML response. The public key associated with the device certificate is used by the identity provider for encryption.

Authentication Context Classes	Use the **Add** and **Remove** buttons to select authentication context classes to be sent in the authentication requests to the SAML identity provider. These are included in the RequestedAuthnContext element.
	In the OASIS standard, an authentication context is defined as "the information, additional to the authentication assertion itself, that the relying party may require before it makes an entitlements decision with respect to an authentication assertion."
	This feature supports all authentication context classes specified in the SAML 2.0 OASIS Authn Context specification.

	For example, if you select X509, the system sends the following context:
	
	`<samlp:RequestedAuthnContext>`
	`    <saml:AuthnContextClassRef xmlns:saml="urn:oasis:names:tc:SAML:2.0:assertion">`
	`            urn:oasis:names:tc:SAML:2.0:ac:classes:X509</saml:AuthnContextClassRef>`
	`</samlp:RequestedAuthnContext>`

	In response, the SAML IdP sends the context data along with the authentication results. The system stores the context data in the session cache and as a system variable named samlAuthnContextClass. The system variable can be used in role mapping rules and resource policy detailed rules.
	Specify a comparison attribute within the RequestedAuthnContext element. The comparison attribute specifies the relative strengths of the authentication context classes specified in the request and the authentication methods offered by a SAML IdP. The following values defined in the SAML 2.0 OASIS core specification can be selected: • **exact**——Requires the resulting authentication context in the authentication statement to be the exact match of at least one of the authentication contexts specified. • **minimum**—Requires the resulting authentication context in the authentication statement to be at least as strong as one of the authentication contexts specified. • **maximum**—Requires the resulting authentication context in the authentication statement to be stronger than any one of the authentication contexts specified. • **better**—Requires the resulting authentication context in the authentication statement to be as strong as possible without exceeding the strength of at least one of the authentication contexts specified. Select the same value that is configured on the SAML IdP. If none is specified in the SAML IdP configuration, the implicit default is exact.
Service Provider Metadata Settings	
Metadata Validity	Enter the number of days the metadata is valid. Valid values are 0 to 9999. 0 specifies the metadata does not expire.
Do Not Publish SA Metadata	Select this option if you do not want to publish the metadata at the location specified by the Entity ID field.
Download Metadata	This button appears only after you have saved the authentication server configuration. Use this button to download the metadata of the current SAML service provider.
User Record Synchronization	
Enable User Record Synchronization	Allow users to retain their bookmarks and individual preferences regardless of which device they log in to.
Logical Auth Server Name	Specify the server name if you have enabled user record synchronization.

Displaying the User Accounts Table

To display user accounts:

1. Select **Authentication > Auth. Servers**.

2. Click the link for the authentication server you want to manage.

3. Click the **Users** tab to display the user accounts table.

   The user accounts table includes entries for the accounts that have been created. The Last Sign-in Statistic column shows the last successful sign-in date and time for each user, the user's IP address, and the agent or browser type and version.

4. Use the controls to search for users and manage user accounts:

   - To search for a specific user, enter a username in the Show users named box and click **Update**.

     TIP: You can use an asterisk (*) as a wildcard, where * represents any number of zero or more characters. For example, to search for all usernames that contain the letters jo, enter ***jo***. The search is case-sensitive. To display the entire list of accounts again, type ***** or delete the field's contents and click **Update**.

   - To limit the number of users displayed on the page, enter a number in the Show *N* users box and click **Update**.

   - To terminate the user session and delete the account, select the box next to the user account record and click **Delete**.

Related Documentation

   - Understanding SAML 2.0

   - SAML 2.0 Supported Features Reference

   - Using SAML AuthnContext Class Variables in Role Mapping and Web ACL Rules

   - Understanding SAML 1.1

   - Understanding SAML 1.1 Profiles

   - Understanding SAML 1.1 Assertions

Using a SiteMinder Server

This topic describes integration with the SiteMinder server. It includes the following information:

   - SiteMinder Server Overview

   - Configuring the Back-End SiteMinder Server

   - Configuring Authentication with a SiteMinder Server

   - Displaying the User Accounts Table

SiteMinder Server Overview

This section describes support for using Policy Secure and Connect Secure with the SiteMinder server. It includes the following sections:

   - Understanding SiteMinder Server

- Feature Support

- Interoperability Requirements and Limitations

Understanding SiteMinder Server

CA SiteMinder server is an authentication and authorization server.

When you configure the Pulse Secure access management framework to authenticate users with a SiteMinder policy server, the system passes the user's credentials to SiteMinder during authentication. Once SiteMinder receives the credentials, it may use standard username and password authentication, RSA Authentication Manager SecurID tokens, or client-side certificates to authenticate the credentials.

The system also passes a protected resource URL to SiteMinder during authentication to determine which SiteMinder realm it should use to authenticate the user. When the system passes the protected resource URL, SiteMinder authorizes the user's URL against the realm that is associated with the resource and allows the user to seamlessly access any resources whose protection levels are equal to or less than the URL that was passed.

Feature Support

The Pulse Secure access management framework supports the following SiteMinder features:

- Single Sign-on Using SMSESSION Cookies
- Automatic Sign-In
- Authentication Schemes

Single Sign-on Using SMSESSION Cookies

The Pulse Secure access management framework enables single sign-on (SSO) to SiteMinder-protected resources using SMSESSION cookies. An SMSESSION cookie is a security token that encapsulates SiteMinder session information. Depending on your configuration, either the SiteMinder Web agent or the system creates an SMSESSION cookie and then posts the cookie to the following locations so the user does not have to reauthenticate to access additional resources.

- Pulse Secure access management framework-If the user tries to access a SiteMinder resource within the session (for example, from the system file browsing page), the system passes its cached SMSESSION cookie to the Web agent for authentication.

- The user's Web browser-If the user tries to access a SiteMinder resource from outside the session (for example, when using a protected resource on a standard agent), SiteMinder uses the cached SMSESSION cookie stored in the user's Web browser to authenticate/authorize the user.

Automatic Sign-In

If you enable the Automatic Sign-In option, the system can use an SMSESSION cookie generated by another agent to enable single sign-on from a SiteMinder resource. When a user accesses the system sign-in page with an SMSESSION cookie, the system verifies the SMSESSION cookie. Upon successful verification, the system establishes a session for the user. You can use the following authentication mechanisms when you enable automatic sign-in through the system:

- Custom agent-The system authenticates the user against the policy server and generates a SMSESSION cookie. When you select this option, you can enable SSO on other SiteMinder agents that use the same policy server. To enable SSO on these agents, update each of them to accept third-party cookies. If you select this option and the user enters his system session with an SMSESSION cookie, the system attempts automatic sign-in when the user enters the session.

- HTML form post-The system posts credentials to a standard Web agent that you have already configured. The Web agent then creates SMSESSION cookies. If you select this option, you cannot use SecurID New Pin and Next Token modes or client-side certificate authentication. If you select this option and the user enters his session with an SMSESSION cookie, the system attempts automatic sign-in when the user enters the session.

- Delegated authentication-The system delegates authentication to a standard agent. If this option is enabled, the system tries to determine the FCC URL associated with the protected resource. The system then redirects the user to the FCC URL with the system sign-in URL as the target. Upon successful authentication, the user is redirected back to the system with an SMSESSION cookie and the system does an automatic sign-in for the user.

Authentication Schemes

The Pulse Secure access management framework works with the following types of SiteMinder authentication schemes:

- Basic username and password authentication—The user's name and password are passed to the SiteMinder policy server. The policy server authenticates them to another server for authentication.

- RSA Authentication Manager SecurID token authentication—The SiteMinder policy server authenticates users based on a username and password generated by an RSA Authentication Manager SecurID token.

- Client-side certificate authentication—The SiteMinder policy server authenticates users based on their client-side certificate credentials. If you choose this authentication method, the Web browser displays a list of client certificates from which users can select. If you choose to authenticate users with this method, you must import the client certificate through the System > Certificates > Trusted Client CAs tab.

Interoperability Requirements and Limitations

The following requirements and limitations apply:

- The Automatic Sign-in feature is not supported for administrator roles. This feature is only available for end users.

- If you use the Authenticate using custom agent option, update all other Web agents to accept the device generated cookie, and apply a software patch to all other Web agents.

- Pulse Policy Secure supports SiteMinder server version 6.0, version 5.5, and version 12.0. If you run older agents than the supported agents, you might experience cookie validation problems, including crossed log entries and intermittent user timeouts.

- You can choose which SiteMinder server version you want to support when you create a server instance. You can choose version 5.5, which supports both versions 5.5 and 6.0, or you can choose version 6.0, which supports only version 6.0, or version 12.0. There is no

difference in the SiteMinder authentication server functionality based on which version you select. This option only controls the version of the SDK to use. We recommend you match the compatibility mode with the version of the policy server.

- When you use SiteMinder to authenticate, the primary and backup policy servers must run the same SiteMinder server software version. A mixed deployment (where the primary server runs a different server software version than the backup) is not supported.

- SiteMinder does not store the IP address in the SMSESSION cookie, and therefore cannot pass it to the system.

- SiteMinder sends the SMSESSION cookie to the system as a persistent cookie. To maximize security, the system resets the persistent cookie as a session cookie once authentication is complete.

- When you use SiteMinder to authenticate, the Pulse Secure access management framework disregards any system session and idle timeouts and uses session and idle timeouts set through the SiteMinder realm instead.

- When you use SiteMinder to authenticate, users must access the system using a fully qualified domain name. This is because the SiteMinder SMSESSION cookie is only sent for the domain for which it is configured. If users access the system using an IP address, they might receive an authentication failure and will be prompted to authenticate again.

- You can update all of your standard Web agents to the appropriate Siteminder Agent Quarterly Maintenance Release (QMR) to accept the cookies. If you are running SiteMinder version 5 Web agents, use the QMR5 hot fix. The system is compatible with version 5.x and later SiteMinder agents. Older versions of SiteMinder agents are susceptible to cookie validation failures.

- You can set the Accept Third Party Cookie attribute (AcceptTPCookie) to yes in the Web agent's configuration file (webagent.conf) or to 1 in the Windows Registry for the IIS Web server. The location of the attribute depends on the SiteMinder version and Web server you are using. Refer to the documentation provided with your SiteMinder server.

Configuring the Back-End SiteMinder Server

The following sections do not give complete SiteMinder configuration instructions—they are only intended to help you make SiteMinder work with the Pulse Secure access management framework. For in-depth SiteMinder configuration information, refer to the documentation provided with your SiteMinder policy server.

- Configuring the SiteMinder Agent
- Configuring the Authentication Scheme
- Configuring the SiteMinder Domain
- Configuring the SiteMinder Realm
- Configuring a Rule or Response Pair to Pass Usernames

Configuring the SiteMinder Agent

A SiteMinder agent filters user requests to enforce access controls. For instance, when a user requests a protected resource, the agent prompts the user for credentials based on an authentication scheme and sends the credentials to a SiteMinder policy server. A Web agent is

simply an agent that works with a Web server. When configuring SiteMinder to work with the Pulse Secure access management framework, you must configure the system as a Web agent in most cases.

If you select the Delegate authentication to a standard agent option, you must set the following options in the agent configuration object of the standard Web agent to host the FCC URL:

- <EncryptAgentName=no>

- <FCCCompatMode=no>

To configure the system as a Web agent on the SiteMinder policy server:

1. In the SiteMinder Administration interface, click the **System** tab.

2. Right-click **Agents** and select **Create Agent**.

3. Enter a name for the Web agent and a description. You must enter this name when creating a SiteMinder realm.

4. Select the **Support 5.x agents** option for compatibility with the system.

5. Under Agent Type, select **SiteMinder** and then select **Web Agent** from the list. This setting is required for compatibility with the system.

6. Under IP Address or Hostname, enter the name or IP address of the system.

7. In the Shared Secret box, enter and confirm a secret for the Web agent. Note that you must enter this secret when configuring the system.

8. Click **OK**.

Configuring the Authentication Scheme

Within SiteMinder, an authentication scheme is a way to collect user credentials and determine the identity of a user. You may create different authentication schemes and associate different protection levels with each. For example, you may create two schemes—one that authenticates users based solely on the users' client-side certificates and provides them a low protection level, and a second that uses RSA Authentication Manager SecurID token authentication and provides users a higher protection level.

To configure a SiteMinder authentication scheme:

1. In the SiteMinder Administration interface, select the **System** tab.

2. Right-click **Authentication Schemes** and select **Create Authentication Scheme**.

3. Enter a name for the scheme and (optionally) a description. You must enter this name when configuring the SiteMinder realm.

4. Under Authentication Scheme, select one of the following options:

   - **Basic Template**

   - **HTML Form Template**

   - **SecurID HTML Form Template**-If you are using SecurID authentication, you must choose SecurID HTML Form Template (instead of SecurID Template). Choosing this option enables the Policy Server to send ACE sign-in failure codes.

- **X509 Client Cert Template**

- **X509 Client Cert and Basic Authentication**

> ⓘ NOTE:
> - You must select HTML Form Template to handle reauthentication.
>
> - If you select X509 Client Cert Template or X509 Client Cert and Basic Authentication, you must import the certificate through System > Certificates > Trusted Client CAs.

5. Enter a protection level for the scheme. Note that this protection level carries over to the SiteMinder realm that you associate with this scheme.

6. Select **Password Policies Enabled for this Authentication Scheme** if you want to reauthenticate users who request resources with a higher protection level than they are authorized to access.

7. Select **Scheme Setup** tab, and enter the options required by your authentication scheme type.

   If you want the system to reauthenticate users who request resources with a higher protection level than they are authorized to access, you must enter the following settings:

   - Under Server Name, enter the hostname (for example, **sales.yourcompany.net**).

   - Select the **Use SSL Connection** check box.

   - Under Target, enter the sign-in URL defined plus the parameter "ive=1" (for example, /highproturl?ive=1). The system must have a sign-in policy that uses ***/highproturl** as the sign-in URL and only uses the corresponding SiteMinder authentication realm.

   - Clear the **Allow Form Authentication Scheme to Save Credentials** check box.

   - Leave **Additional Attribute List** empty.

8. Click **OK**.

   If you change a SiteMinder authentication scheme on the policy server, you must flush the cache using the Flush Cache option on the Advanced tab.

Configuring the SiteMinder Domain

Within SiteMinder, a policy domain is a logical grouping of resources associated with one or more user directories. Policy domains contain realms, responses, and policies. When configuring the Pulse Secure access management framework to work with SiteMinder, you must give user access to a SiteMinder resource within a realm, and then group the realm into a domain.

To configure a SiteMinder domain:

1. Select the **System** tab, right-click **Domains** and select **Create Domain**, or click **Domains** and select an existing SiteMinder domain.

2. Add a realm to the domain.

Configuring the SiteMinder Realm

Within SiteMinder, a realm is a cluster of resources within a policy domain grouped together according to security requirements. When configuring SiteMinder to work with the Pulse Secure access management framework, you must define realms to determine which resources the users might access.

To configure a SiteMinder Realm:

1. In the SiteMinder Administration interface, select the **Domains** tab.

2. Expand the domain that you created.

3. Right-click **Realms** and select **Create Realm**.

4. In the Agent field, select the Web agent that you created.

5. In the Resource Filter field, enter a protected resource. This resource inherits the protection level specified in the corresponding authentication scheme.

   For the default protection level, enter **/ive-authentication**. You must enter this resource when configuring the system. If you use sign-in policies with nondefault URLs such as */nete or */cert, you must have corresponding resource filters in the SiteMinder configuration.

6. From the Authentication Schemes list, select the scheme that you created.

7. Click **OK**.

Configuring a Rule or Response Pair to Pass Usernames

Within SiteMinder, you can use rules to trigger responses during authentication or authorization. A response passes DN attributes, static text, or customized active responses from the SiteMinder policy server to a SiteMinder agent. When you configure SiteMinder to work with the Pulse Secure access management framework, you must create a rule that triggers when a user successfully authenticates. Then, you must create a corresponding response that passes the user's username to the system Web agent.

To create a new rule:

1. In the SiteMinder Administration interface, select the **Domains** tab.

2. Expand the domain that you created and then expand **Realms**.

3. Right-click the realm that you created and select **Create Rule under Realm**.

4. Enter a name and (optionally) description for the rule.

5. Under Action, select **Authentication Events** and then select **OnAuthAccept** from the drop down list.

6. Select **Enabled**.

7. Click **OK**.

To create a new response:

1. In the SiteMinder Administration interface, select the **Domains** tab.

2. Expand the domain that you created.

3. Right-click **Responses** and select **Create Response**.

4. Enter a name and (optionally) a description for the response.

5. Select **SiteMinder** and then select the Web agent.

6. Click **Create**.

7. From the Attribute list, select **WebAgent-HTTP-Header-Variable**.

8. Under Attribute Kind, select **Static**.

9. Under Variable Name, enter **IVEUSERNAME**.

10. Under Variable Value, enter a username.

11. Click **OK**.

Configuring Authentication with a SiteMinder Server

To configure authentication with SiteMinder server:

1. Select **Authentication** > **Auth. Servers**.

2. Select **SiteMinder Server** and click **New Server**.

3. Complete the configuration as described in Table 32.

4. Save the configuration.

   After you have saved the configuration, the page that is redisplayed includes an Advanced tab.

5. Click the **Advanced** tab to display the configuration page.

6. Complete the configuration as described in Table 33.

7. Save the configuration.

Table32: SiteMinder Server Settings

Settings	Guidelines
Name	Specify a name to identify the server within the system.
Policy Server	Specify name or IP address of the policy server.
Backup Server(s)	(Optional) Specify a comma-delimited list of backup policy servers.
Failover Mode?	Select one of the following failover mode options: • **Yes**–The device uses the main policy server unless it fails. • **No**–The device does the load balancing among all the specified policy servers.
Agent Name	Specify the agent name configured on the policy server.
Secret	Specify the shared secret configured on the policy server. The value is case sensitive.

Compatible with	Select a SiteMinder server version.
	• **5.5 Policy Servers**–Supports version 5.5 and version 6.0. This is the default.
	• **6.0 Policy Servers**–Supports only version 6.0 of the SiteMinder server API.
	• **12.0 Policy Servers**–Supports only version 12.0.
On logout, redirect to	Specify a URL to which users are redirected when they sign out of the device (optional). If you leave this field empty, users see the default sign-in page.
	The On logout, redirect to setting is included in the product release for backwards-compatibility. We strongly recommend that you use the customizable sign-in pages feature instead.
Protected Resource	Specify a default protected resource. If you do not create sign-in policies, the system uses this default URL to set the user's protection level for the session. The system also uses this default URL if you select the **Automatic Sign-In** option. If your users are signing in to the "*" URL (default device sign-in page), enter any URL ("/ive-authentication" is the default) to set the protection level to the default value. If you do create sign-in policies, the device uses those sign-in policies instead of this default URL.
	You must enter a forward slash (/) at the beginning of the resource. For example, enter **/local-authentication**.
Resource Action	Displays the resource action configured on the back-end SiteMinder server.
Users authenticate using tokens or one-time passwords	Select this option if you want the device to prompt the user for a token instead of a password; that is, if users submit tokens or one-time use passwords to the device.
	For example, you can use this option to dynamically prompt for a password or token based on sign-in policies by configuring two instances of the same authentication server. You can use one instance for wireless users who have this option enabled and it prompts the user for a token, and another instance for wired users who have this option disabled and it prompts the user for a password.
	This feature is available only on Policy Secure.
Server Catalog	Use the Server Catalog button to display the Server Catalog in a new window. Add the SiteMinder user attributes (such as the cookiename) that you want to use for role mapping.
SMSESSION cookie settings	
When sending cookies to the end-user's browser	Specify the cookie domain for either the end user or the device. A cookie domain is a domain in which the user's cookies are active. For example, the system sends cookies to the user's browser in this domain.
	Multiple domains should use a leading period and be comma-separated. For example, **.sales.myorg.com**, **.marketing.myorg.com**.
	Domain names are case-sensitive. You cannot use wildcard characters. For example, if you define ".pulsesecure.net" the user must access the device as "http://ive.pulsesecure.net" to ensure that his SMSESSION cookie is sent back to the device.
	Select **HTTPS** to send cookies securely if other Web agents are set up to accept secure cookies, or **HTTP** to send cookies nonsecurely.
Cookie Domain and Protocol When the Cookie is Set on the Device	Enter the valid Internet domain for the cookie and where the browser of the user sends cookie contents. This cookie domain should be the same as the host domain. For example, .jnpr.net
	Select HTTPS to send cookies securely if other Web agents are set up to accept secure cookies, or HTTP to send cookies non-securely.
SiteMinder authentication settings	

Automatic Sign In	Select this option to automatically sign in users with a valid SMSESSION cookie. Then, select the authentication realm to which the users are mapped. If you select this option, note that: • If the protection level associated with a user's SMSESSION cookie is different from the protection level of the realm, the protection level associated with the cookie is used. • This option uses SMSESSION cookie, which is already present in the browser to enable single sign-on. • This option provides a single sign-on experience for users. • This option enables users to sign in using a standard Siteminder Web Agent that generates an SMSESSION cookie. When you select this option, you must also configure the following suboptions: • To assign user roles, use this user authentication realm–Select an authentication realm for automatically signed-in users. The users are mapped to a role based on the role mapping rules defined in the selected realm. • If Automatic Sign In fails, redirect to–Enter an alternative URL for users who sign in through the automatic sign-In mechanism. The users are redirected to the specified URL if the authentication fails and if there is no redirect response from the SiteMinder policy server. If you leave this field empty, users are prompted to sign back in.
Authenticate using custom agent	Select this option if you want to authenticate using the custom Web agent. Using this option, the system generates the SMSESSION cookie, just like any other Web agent configured within the organization.
Authenticate using HTML form post	Select this option if you want to post user credentials to a standard Web agent that you have already configured rather than contacting the SiteMinder policy server directly. If you select this option, the Web agent contacts the policy server to determine the appropriate sign-in page to display to the user. To configure the system to "act like a browser" that posts credentials to the standard Web agent, you must enter the following information. • **Target**–Specify the target URL. • **Protocol**–Specify the protocol for communication between the system and the specified Web agent. Select **HTTP** for non-secure communication. Select **HTTPS** for secure communication. • **Webagent**–Specify the name of the Web agent to obtain SMSESSION cookies. An IP address is not allowed for this field. (Specifying the IP address as the Web agent prevents some browsers from accepting cookies.) • **Port**– Specify the port for the protocol. Enter port **80** for HTTP or port **443** for HTTPS. • **Path**–Specify the path of the Web agent's sign-in page. The path must start with a backslash (/) character. In the Web agent sign-in page URL, the path appears after the Web agent. • **Parameters**– Specify the post parameters to be sent when a user signs in. Common SiteMinder variables that you can use include _ _USER_ _, _ _PASS_ _, and _ _TARGET_ _. These variables are replaced by the username and password entered by the user on the Web agent's sign-in page and by the value specified in the Target field. These are the default parameters for login.fcc—if you have made customizations, you may need to change these parameters.

Delegate authentication to a standard agent	Select this option to delegate authentication to a standard agent. When the user accesses the system sign-in page, the FCC URL associated with the protected resource's authentication scheme is determined. The system redirects the user to that URL, setting the system sign-in URL as the target. After successfully authenticating with the standard agent, an SMSESSION cookie is set in the user's browser and the user is redirected back. The system then automatically signs in the user and establishes a session. You must enable the Automatic Sign-In option to use this feature. If you enable this option and a user already has a valid SMSESSION cookie when trying to access a resource, the system tries to automatically sign in using the existing SMSESSION cookie. If the cookie is invalid, the SMSESSION cookie and corresponding system cookies are cleared and a "timeout" page is displayed. The system successfully delegates authentication when the user clicks the **sign back in** option. If you select this option, your authentication scheme must have an associated FCC URL.
SiteMinder authorization settings	This feature is available only on Connect Secure.
Authorize requests against SiteMinder policy server	Use SiteMinder policy server rules to authorize user Web resource requests. If you select this option, make sure that you create the appropriate rules in SiteMinder that start with the server name followed by a forward slash, such as: www.yahoo.com/, www.yahoo.com/*, and www.yahoo.com/r/f1.
If authorization fails, redirect to	Specify an alternative URL that users are redirected to if the device fails to authorize and no redirect response is received from the SiteMinder policy server. If you leave this field empty, users are prompted to sign back into the device. If you are using an authorization-only access policy , you must enter an alternative URL in this field regardless of whether the Authorize requests against SiteMinder policy server option is selected. Users are redirected to this URL when an access denied error occurs.
Resource for insufficient protection level	Specify a resource on the Web agent to which the users are redirected when they do not have the appropriate permissions.
Ignore authorization for files with extensions	Specify the file extensions corresponding to file types that do not require authorization. Enter the extensions of each file type that you want to ignore, separating each with a comma. For example, enter **.gif, .jpeg, .jpg, .bmp** to ignore various image types. You cannot use wildcard characters (such as *, *.*, or .*) to ignore a range of file types.
User Record Synchronization	This feature is available only on Connect Secure.
Enable User Record Synchronization	Select this option to retain the bookmarks and individual preferences regardless of which system you log in to.
Logical Auth Server Name	Specify a logical authentication server name.

Table33: SiteMinder Advanced Configuration Options

Settings	Guidelines
Poll Interval (seconds)	Specify the interval at which the system polls the SiteMinder policy server to check for a new key.
Max. Connections	Control the maximum number of simultaneous connections that the system is allowed to make to the policy server. The default setting is 20.
Max. Requests/ Agent	Control the maximum number of requests that the policy server connection handles before the system ends the connection. If necessary, tune to increase performance. The default setting is 1000.

Idle Timeout (minutes)	Control the maximum number of minutes a connection to the policy server may remain idle (the connection is not handling requests) before the system ends the connection. The default setting of "none" indicates no time limit.
Authorize while Authenticating	Specify that the system should look up user attributes on the policy server immediately after authentication to determine if the user is truly authenticated.
	For example, if your SiteMinder server authenticates users based on an LDAP server setting, you can select this option to indicate that the system should authenticate users through the SiteMinder server and then authorize them through the LDAP server before granting them access. If the user fails authentication or authorization, the user is redirected to the page configured on the policy server.
Enable Session Grace Period	Eliminate the overhead of verifying a user's SMSESSION cookie each time the user requests the same resource by indicating that the system should consider the cookie valid for a certain period of time.
	If you do not select this option, the system checks the user's SMSESSION cookie on each request. Note that the value entered here does not affect session or idle timeout checking.
Validate cookie every N seconds (seconds)	Specify the time period for the system to eliminate the overhead of verifying a user's SMSESSION cookie each time the user requests the same resource by indicating that the system should consider the cookie valid for a certain period of time.
Ignore Query Data	Specify that the system does not cache the query parameter in its URLs. Therefore, if a user requests the same resource as is specified in the cached URL, the request should not fail.
Accounting Port	Specify that the value entered in this field must match the accounting port value entered through the Policy Server Management Console in the Web UI. By default, this field matches the policy server's default setting of 44441.
Authentication Port	Specify that the value entered in this field must match the authentication port value entered through the Policy Server Management Console. By default, this field matches the policy server's default setting of 44442.
Authorization Port	Specifies that the value entered in this field must match the authorization port value entered through the Policy Server Management Console. By default, this field matches the policy server's default setting of 44443.
Agent Configuration Settings	
Overlook Session for Methods	Compare the request method to the methods listed in this parameter. If a match is found, Web Agent does not create a new or update an existing SMSESSION cookie, nor will it make any updates to the cookie provider for that request.
	You can enter multiple methods; use a comma to separate method names.
	If Overlook Session for Methods parameter is set but not Overlook Session for URLs, then all requests that match the methods defined in this parameter are processed (SMSESSION cookie creation/update is blocked).
	If both Overlook Session for Methods and Overlook Session for URLsparameters are set, both the method and the URL of the request are matched before proceeding. Then, all URLs with specified methods are processed (SMSESSION cookie creation/update is blocked).

Overlook Session for URLs	Compare the request URL to the URLs listed in this parameter. If a match is found, Web Agent does not create a new or update an existing SMSESSION cookie, nor will it make any updates to the cookie provider for that request.
	Specify a relative URL. For example: If the URL is http://fqdn.host/MyDocuments/index.html, enter **/MyDocuments/index.html**
	If Overlook Session for URLs is set but not Overlook Session for Methods, then all requests, regardless of the methods, matching the URLs defined in this parameter are processed (SMSESSION cookie creation/update is blocked).
	If both Overlook Session for Methods and Overlook Session for URLsparameters are defined, both the method and the URL of the request are matched before proceeding. Then, all URLs with specified methods are processed (SMSESSION cookie creation/update is blocked).
SiteMinder caching	
Flush Cache	Select this option to delete the resource cache, which caches resource authorization information for 10 minutes.

Displaying the User Accounts Table

To display user accounts:

1. Select **Authentication > Auth. Servers**.

2. Click the link for the authentication server you want to manage.

3. Click the **Users** tab to display the user accounts table.

   The user accounts table includes entries for the accounts that have been created. The Last Sign-in Statistic column shows the last successful sign-in date and time for each user, the user's IP address, and the agent or browser type and version.

4. Use the controls to search for users and manage user accounts:

   - To search for a specific user, enter a username in the Show users named box and click **Update**.

     > TIP: You can use an asterisk (*) as a wildcard, where * represents any number of zero or more characters. For example, to search for all usernames that contain the letters jo, enter ***jo***. The search is case-sensitive. To display the entire list of accounts again, type ***** or delete the field's contents and click **Update**.

   - To limit the number of users displayed on the page, enter a number in the Show *N* users box and click **Update**.

   - To terminate the user session and delete the account, select the check box next to the user account record and click **Delete**.

Related Documentation

- AAA Server Overview

- *Troubleshooting the SiteMinder Server Configuration*

- How To: Configuring SiteMinder

- KB: SiteMinder Protected Resource

- KB: SiteMInder User Access Log

CHAPTER 9 SAML Single Sign-on

This chapter describes how to use Connect Secure in a SAML single sign-on deployment. It includes the following sections:

- Pulse Connect Secure SAML 2.0 SSO Solutions
- SAML 2.0 Configuration Tasks
- Example: Implementing SAML 2.0 Web Browser SSO for Google Apps
- Using SAML AuthnContext Class Variables in Role Mapping and Web ACL Rules
- Investigating a "No valid assertion found in SAML response" Error
- Pulse Connect Secure SAML 1.1 Support

Pulse Connect Secure SAML 2.0 SSO Solutions

This section provides a brief overview of the Security Assertion Markup Language (SAML) standard produced and approved by the Organization for the Advancement of Structured Information Standards (OASIS). It includes the following topics:

- Understanding SAML 2.0
- SAML 2.0 Supported Features Reference

Understanding SAML 2.0

This topic provides a reference to the Security Assertion Markup Language (SAML) standard and an overview of SAML 2.0 use cases. It includes the following information:

- About SAML
- SAML Use Cases

About SAML

SAML is an XML-based framework for communicating user authentication, entitlement, and attribute information. The standard defines the XML-based assertions, protocols, bindings, and profiles used in communication between SAML entities. SAML is used primarily to implement Web browser single sign-on (SSO). SAML enables businesses to leverage an identity-based security system like Connect Secure to enforce secure access to web sites and other resources without prompting the user with more than one authentication challenge.

For complete details on the SAML standard, see the OASIS web site:

http://www.oasis-open.org/committees/tc_home.php?wg_abbrev=security

SAML Use Cases

This section provides a brief summary of the primary SAML use cases. It includes the following information:

- SAML SSO
- SAML ACL

SAML SSO

SAML is primarily used to enable Web browser single sign-on (SSO). Pulse Secure client R5.0 or greater (Win and Mac) also supports SAML SSO. The user experience objective for SSO is to allow a user to authenticate once and gain access to separately secured systems without resubmitting credentials. The security objective is to ensure the authentication requirements are met at each security checkpoint.

In an SSO transaction, the SAML services implemented at each secured system exchange requests and assertions to determine whether to allow access. The SAML assertions used in SSO transactions include authentication statements and attribute statements.

SAML ACL

SAML can also be used to enforce access control list (ACL) permissions. In an ACL transaction, the SAML services implemented for each secured system exchange assertions to determine whether a user can access the resource. The SAML assertions used in ACL transactions include authorization requests and authorization decision statements.

Related Documentation

- SAML 2.0 Supported Features Reference

- Configuring Connect Secure as a SAML 2.0 Service Provider

- Configuring Connect Secure as a SAML 2.0 Identity Provider

- Configuring a SAML 2.0 ACL Web Policy

SAML 2.0 Supported Features Reference

This topic provides an overview of Connect Secure support for Security Assertion Markup Language (SAML) single sign-on (SSO). It includes the following information related to SAML 2.0 support:

- Supported SAML SSO Deployment Modes
- Supported SAML SSO Profiles
- FIPS Support Notes

Supported SAML SSO Deployment Modes

In a SAML deployment, a SAML service provider is a secured resource (an application, web site, or service) that is configured to request authentication from a SAML identity provider. The SAML identity provider responds with assertions regarding the identity, attributes, and entitlements (according to your configuration). The exchange enforces security and enables the SSO user experience.

The system can act as a SAML service provider, a SAML identity provider, or both. The following sections provide illustrations:

- Connect Secure as a SAML Service Provider
- Connect Secure As a SAML Identity Provider (Gateway Mode)
- Connect Secure as a SAML Identity Provider (Peer Mode)

Connect Secure as a SAML Service Provider

If you are working with a partner that has implemented a SAML identity provider, you can deploy the system as a SAML service provider to interoperate with it, thereby enabling SSO for users who should have access to resources in both networks. In this model, the user is authenticated by the SAML identity provider. The system uses the SAML response containing the assertion to make an authentication decision.

The choices the identity provider makes to implement SAML determine the deployment choices, for example whether to use SAML 2.0 or SAML 1.1, whether to reference a published metadata configuration file, and whether to use a POST or artifact profile. When you deploy the system as a SAML service provider, you create a *SAML authentication server* configuration that references the partner SAML identity provider, and a set of access management framework objects (realm, role mapping rules, and sign-in policy) that reference the SAML authentication server.

When you configure the SAML service provider, some particular settings are necessary to support either identity-provider-initiated or service-provider-initiated SSO. The documentation for the configuration steps makes note of these settings. Regardless, you configure the SAML service provider to support both identity-provider-initiated and service-provider-initiated SSO.

Figure 26 illustrates the flow of network communication in a service-provider-initiated SSO scenario with a Web browser client.

Figure26: Connect Secure as a SAML Service Provider in a Service-Provider-Initiated SSO Scenario

1 – The user clicks a link to access a resource.

2a – The service provider sends an HTTP redirect status code (HTTP 302) to the user. The SAML request and all other SAML details are sent as URL parameters in the URL Location header.

2b – The user sends an HTTP GET request to the identity provider. The SAML request and all other SAML details are sent as URL parameters.

If the user already has a session with the identity provider, steps 2.1 and 2.2 are skipped.

2.1 – If the user does not have a session, the identity provider sends an authentication challenge to the user.

2.2 – The user enters sign-in credentials.

3a – The identity provider sends a successful status code (HTTP 200 OK) to the user with a form in the HTML body.

3b – The user sends the form to the service provider.

4 – The external resource is delivered to the user's browser.

Figure 27 illustrates the flow of network communication in an identity-provider-initiated SSO scenario with a Web browser client.

Figure27: Connect Secure as a SAML Service Provider in an Identity-Provider-Initiated SSO Scenario
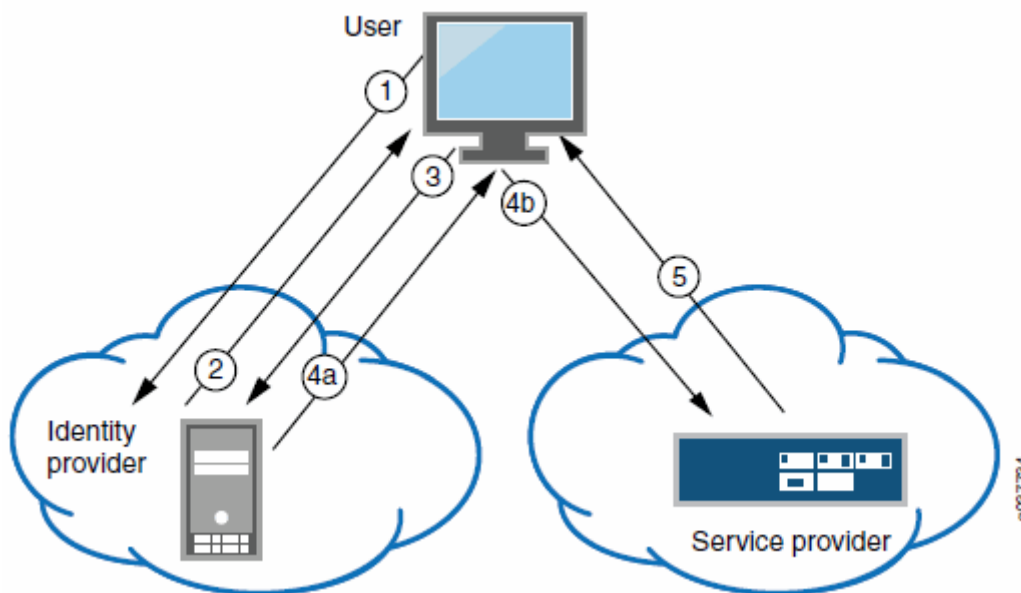
1 – The user authenticates to the identity provider.

2 – The identity provider returns a portal page with links to external resources.

3 – The user clicks a link for an external resource.

4a – The identity provider sends a successful status code (HTTP 200 OK) to the user with a form in the HTML body.

4b – The user sends the form to the service provider.

5 – The external resource is delivered to the user's browser.

Figure 28 illustrates the flow of network communication when a user clicks a Pulse client connection.

Figure28: Connect Secure as a SAML Service Provider in a Pulse-Client-Initiated Connection
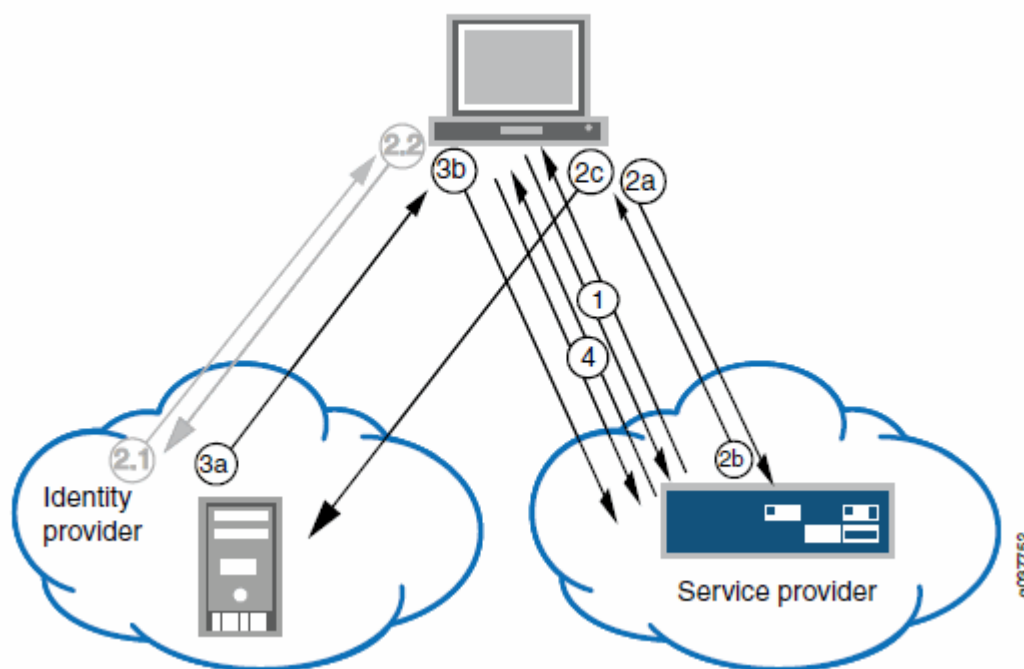
1 – The user clicks the Pulse client connection. The Pulse client and system exchange IF-T/TLS messages. The Pulse client learns that authentication is a SAML exchange, and Pulse launches its embedded client Web browser.

2a – The service provider sends an HTTP redirect status code (HTTP 302) to the user. The SAML request and all other SAML details are sent as URL parameters in the URL Location header.

2b – The user sends an HTTP GET request to the identity provider. The SAML request and all other SAML details are sent as URL parameters.

If the user already has a session with the identity provider, steps 2.1 and 2.2 are skipped.

2.1 – If the user does not have a session, the identity provider sends an authentication challenge to the user.

2.2 – The user enters sign-in credentials.

3a – The identity provider sends a successful status code (HTTP 200 OK) to the user with a form in the HTML body.

3b – The user sends the form to the service provider.

4 – The setup client is run on the endpoint, and the Pulse client and system set up an SSL VPN tunnel.

Connect Secure As a SAML Identity Provider (Gateway Mode)

When you deploy the system in front of enterprise resources that support SAML and have been configured as a SAML service provider, the system acts as a gateway for user access to the secured resource, just as it does with its other resource policies. In the SAML exchange, the system acts as a SAML identity provider. When deployed as a gateway, the SAML SSO communication is always identity-provider-initiated. The system maintains the session and uses its rewriting or pass-through proxy features to render data to the user.

In a gateway mode deployment, you configure the system as a SAML identity provider to correspond with the SAML service provider, and you create a *SAML SSO* resource policy configuration to determine the users and resources to which the SAML SSO experience applies. The SAML SSO resource policy supports two types of behavior that are possible with the HTTP responses sent by SAML service providers:

- The SAML service provider sends HTTP responses that can be handled by HTTP cookies and therefore do not require user interaction. In this case, the SAML SSO resource policy can be configured to use cookies to handle the HTTP transaction.

- The SAML service provider sends HTTP responses that require user interaction. For example, the SAML service provider might send an HTTP 200 OK with an embedded form that requires action from the user, execution of JavaScript, or data to be automatically submitted on load. Or, the resource might send an HTTP 3xx redirect that requires acceptance by the user. In these cases, the SAML SSO resource policy can be configured to forward the HTTP responses through the rewriter, which rewrites the HTTP response and sends it to the end user.

231

Figure 29 illustrates the communication that occurs when the SAML SSO policy is configured to handle the SAML service provider responses using cookies.

Figure29: Connect Secure as a SAML Identity Provider (Gateway Mode) – User/Browser Action Not Required
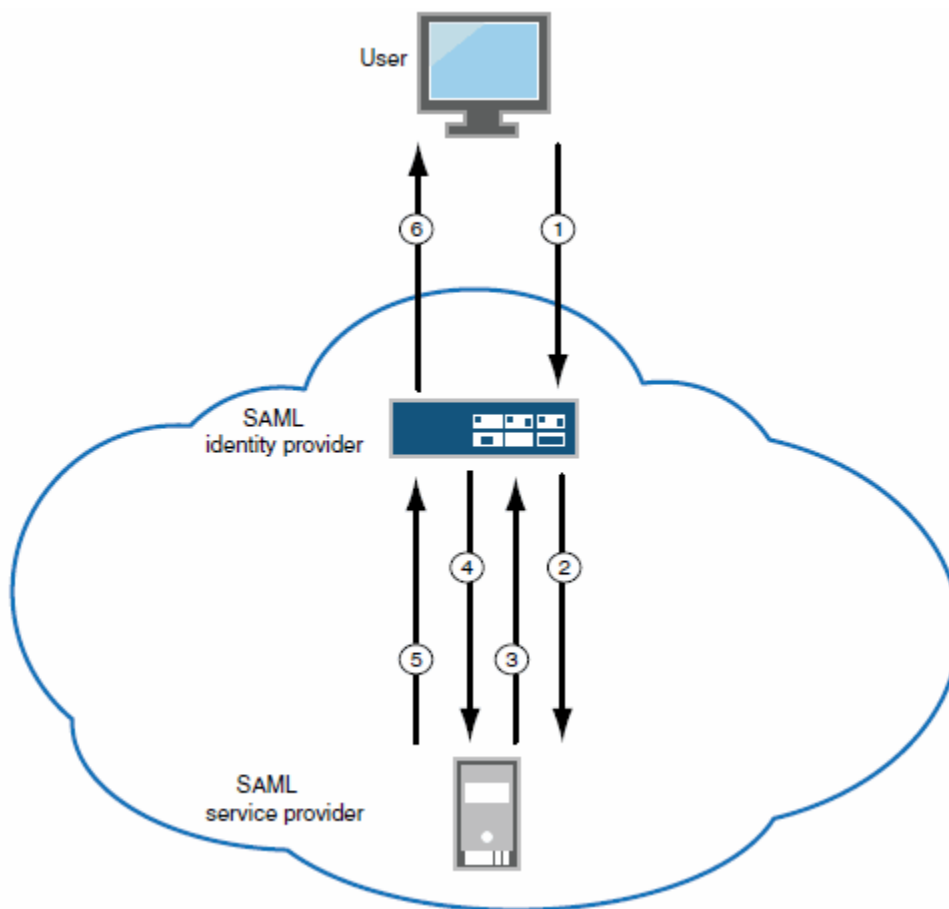
1 – User requests a SAML protected resource.
2 – The system executes the SAML SSO policy and the identity provider sends an HTTP request containing the SAML assertion to the SAML service provider.
3 – The SAML service provider sends an HTTP response. The SAML SSO process extracts the cookies from the response and stores them in the cookie cache.
4 – The system rewriter process sends the request for the resource (sending the cookies received in step 3).
5 – The SAML service provider sends the resource.
6 – The system rewrites the resource and sends it to the user.

Figure 30 illustrates the communication that occurs when the SAML SSO policy is configured to rewrite the SAML service provider responses and send them to the user/browser for action.

Figure30: Connect Secure as a SAML Identity Provider (Gateway Mode) – User/Browser Action Required
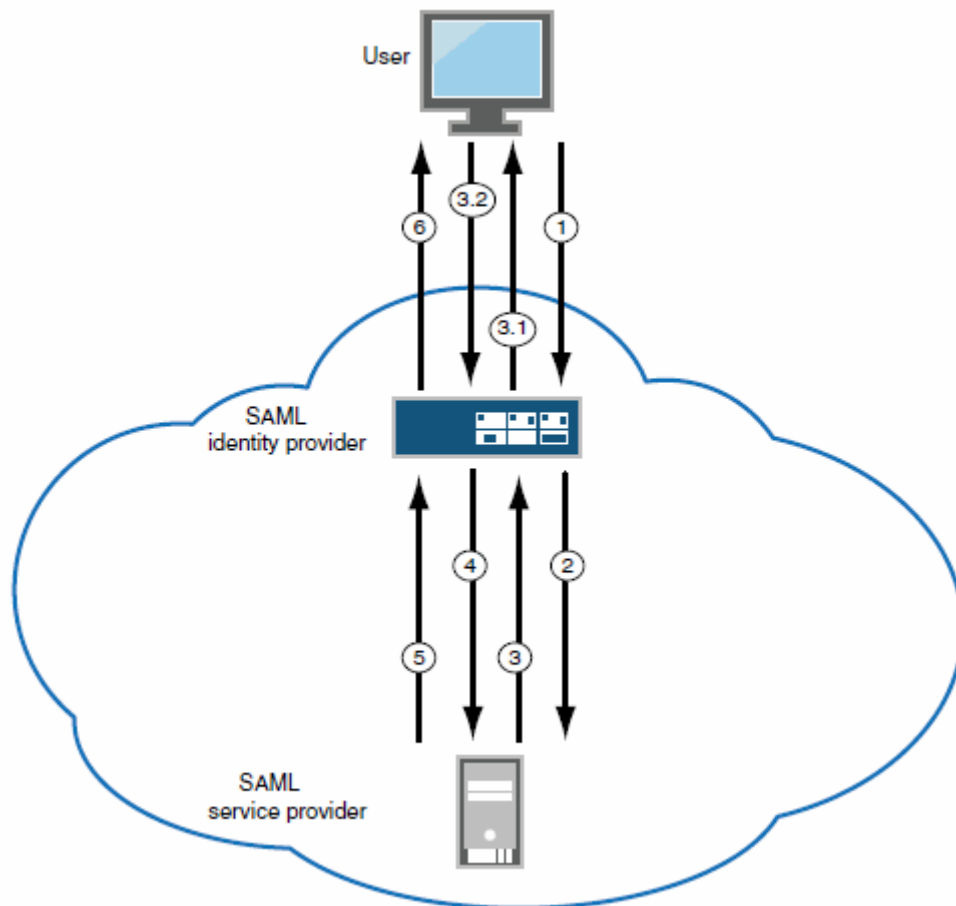
1 – User requests a SAML protected resource.
2 – The system executes the SAML SSO policy and the system identity provider sends an HTTP request containing the SAML assertion to the SAML service provider.
3 – The SAML service provider sends an HTTP response. The system SAML SSO process forwards the entire response to the to the rewriter.
3.1 – The rewriter rewrites the response and sends it to the user.
3.2 – The user/browser completes any action required and sends a response (an HTTP GET/POST request).
4 – The rewriter processes it as any other HTTP web request and forwards to the SAML service provider.
5 – The SAML service provider sends the resource.
6 – The system rewrites the resource and sends it to the user. Steps 5 and 6 can involve many transactions related to Web browsing or use of the resource.

Connect Secure as a SAML Identity Provider (Peer Mode)

When deployed to support access to external resources (for example, public cloud resources), the system does not have to be a gateway to user access. The user can access the external resource directly, and the traffic does not flow through the device. In a peer mode deployment,

you configure the system as a SAML identity provider to correspond with the external SAML service provider, and you create a *SAML External Apps SSO* resource policy configuration to determine the users and resources to which the SAML SSO experience applies.

In a service-provider-initiated SSO scenario, the user requests a resource protected by the SAML service provider. The SAML service provider redirects the user to the sign-in page. The access management framework processes the authentication request, performs host checking rules and role mapping rules. If authentication is successful, the system redirects the user to the protected resource.

In an identity-provider-initiated SSO scenario, the user first creates a session. The access management framework processes are run when the user signs in. The SAML External Apps SSO policy is enforced when the user browses to the SAML protected external application.

When you configure the SAML identity provider, some settings are necessary to support either identity-provider-initiated or service-provider-initiated SSO. The documentation for the configuration steps makes note of these settings. Regardless, you configure the SAML identity provider to support both identity-provider-initiated and service-provider-initiated SSO.

Figure 31 illustrates the flow of network communication in a service-provider-initiated SSO scenario.

Figure31: Connect Secure as a SAML Identity Provider (Peer Mode) in a Service-Provider-Initiated SSO Scenario

1 – The user clicks a link to access a resource.

2a – The service provider sends an HTTP redirect status code (HTTP 302) to the user. The SAML request and all other SAML details are sent as URL parameters in the URL Location header.

2b – The user sends an HTTP GET request to the identity provider. The SAML request and all other SAML details are sent as URL parameters.

If the user already has a session with the identity provider, steps 2.1 and 2.2 are skipped.

2.1 – If the user does not have a session, the identity provider sends an authentication challenge to the user.

2.2 – The user enters sign-in credentials.

3a – The identity provider sends a successful status code (HTTP 200 OK) to the user with a form in the HTML body.

3b – The user sends the form to the service provider.

4 – The external resource is delivered to the user's browser.

235

Figure 32 illustrates the flow of network communication in an identity-provider-initiated SSO scenario.

Figure32: Connect Secure as a SAML Identity Provider (Peer Mode) in an Identity-Provider-Initiated SSO Scenario

1 – The user authenticates to the identity provider.	
2 – The identity provider returns a portal page with links to external resources.	
3 – The user clicks a link for an external resource.	
4a – The identity provider sends a successful status code (HTTP 200 OK) to the user with a form in the HTML body.	
4b – The user sends the form to the service provider.	
5 – The external resource is delivered to the user's browser.	

Supported SAML SSO Profiles

Table 34 summarizes support for SAML 2.0 deployment profiles.

Table34: Supported SAML 2.0 Deployment Profiles

Profile	Message Flows	Binding	Service Provider	Identity Provider (Gateway)	Identity Provider (Peer)
Web browser SSO	<AuthnRequest> from service provider to identity provider	HTTP redirect	Supported	–	Supported
		HTTP POST	Not supported	–	Supported
		HTTP artifact	Not supported	–	Not supported
	Identity provider <Response> to service provider	HTTP POST	Supported	Supported	Supported
		HTTP artifact	Supported	Supported	Supported
	Authentication context classes sent in <RequestedAuthnContext>	HTTP POST	All authentication context classes	–	Password-Protected, TimeSyncToken, and TLSClient only
		HTTP artifact	All authentication context classes	–	Password-Protected, TimeSyncToken, and TLSClient only
Basic attribute profile	Simple <Attribute> statements (name-value pairs) as part of assertions	HTTP POST	Consumes and stores Attribute statements	Sends Attribute statements	Sends Attribute statements
		HTTP artifact			
Assertion query / request	Artifact resolution <ArtifactResolve> <ArtifactResponse>	SOAP	Supported	Supported	Supported
Single logout	Logout request	HTTP redirect	Supported	Not supported	Not supported
		HTTP POST	Not supported	Not supported	Not supported
		HTTP artifact	Not supported	Not supported	Not supported
		SOAP	Not supported	Not supported	Not supported
	Logout response	HTTP redirect	Supported	Not supported	Not supported
		HTTP POST	Not supported	Not supported	Not supported
		HTTP artifact	Not supported	Not supported	Not supported
		SOAP	Not supported	Not supported	Not supported

NOTE: Connect Secure does not act as an Attribute Authority.

FIPS Support Notes

Historically, in FIPS deployments, private keys were managed in a way that can be problematic for SAML functionality that depends on access to the private key. Table 35 summarizes our support for SAML for FIPS deployments.

Table35: SAML Support for FIPS Deployments

	SAML 2.0		SAML 1.1	
Release	Service Provider	Identity Provider	Consumer	Producer
8.0	Same as 7.4.	Same as 7.4	No limitations	Artifact profile only
7.4	Device certificate signing is supported; however, the ECDSA certificates is not supported. There are no other limitations.	Device certificate signing is supported; however, the ECDSA certificates is not supported. There are no other limitations.	No limitations	Artifact profile only
7.3	Same as 7.2.	Same as 7.2.	No limitations	Artifact profile only
7.2	The following configuration restrictions apply to the Authentication Server configuration: • Set Select Device Certificate for Signing to **Not Applicable**. • Set Select Device Certificate for Encryption to **Not Applicable**.	The following configuration restrictions apply to the Sign-In SAML Identity Provider configuration: • Set Protocol Binding to Use for SAML Response to **Artifact**. POST is not supported. • Set Signing Certificate to **No Signing**. • Set Decryption Certificate to **No Encryption**. • Clear the **Enable Artifact Response Signing and Encryption** check box. Support for valid 7.1 configurations (Legacy mode)	No limitations	Artifact profile only
7.1	The following configuration restrictions apply to the authentication server configuration: • Set Select Device Certificate for Signing to **Not Applicable**. • Set Select Device Certificate for Encryption to **Not Applicable**.	Artifact profile only	No limitations	Artifact profile only
7.0	–	–	No limitations	Artifact profile only

Related Documentation

- Understanding SAML 2.0
- Configuring Connect Secure as a SAML 2.0 Service Provider
- Configuring Connect Secure as a SAML 2.0 Identity Provider

SAML 2.0 Configuration Tasks

This section includes the tasks you perform to enable and configure SAML services. It includes the following information:

Configuring System-Wide SAML Settings

This section describes tasks related to configuring system-wide SAML settings. It includes the following topics:

- Configuring Global SAML Settings
- Managing SAML Metadata Files

Configuring Global SAML Settings

The system-wide SAML settings impact all SAML service provider and identity provider instances.

To configure global SAML settings:

1. Select **System > Configuration > SAML**.

2. Click the **Settings** button to display the configuration page.

3. Complete the settings described in Table 36.

4. Click **Save Changes**.

Table36: SAML Global Configuration Guidelines

Settings	Guidelines
Timeout value for metadata fetch request	Specify the number of seconds after which a download request is abandoned. If the peer SAML entity publishes its metadata at a remote location, the system downloads the metadata file from the specified location.
Validity of uploaded/downloaded metadata file	Specify the maximum duration for which the system considers the metadata file of the peer SAML entity to be valid. If the metadata file provided by the peer SAML entity contains validity information, the lower value takes precedence.
Host FQDN for SAML	Specify the fully qualified domain name for the Connect Secure host. The value you specify here is used in the SAML entity ID and the URLs for SAML services, including: • Entity ID for SAML service provider and SAML identity provider instances. The SAML entitiy ID is the URL where the system publishes its SAML metadata file. • Single sign-on service URL • Single logout service URL • Assertion consumer service URL • Artifact resolution service URL BEST PRACTICE: The system uses HTTPS for these services. Therefore, we recommend that you assign a valid certificate to the interface that has the IP address to which this FQDN resolves so that users do not see invalid certificate warnings.

Alternate Host FQDN for SAML	Optional. If you have enabled the **Reuse Existing NC (Pulse) Session** on the SAML Identity Provider Sign-In page, specify the fully qualified domain name used to generate the SSO Service URL. Set up your DNS service to ensure that the alternate hostname resolves to a different IP address when a session is established and when not established. We recommend the following DNS behavior: • If the NC (Pulse) session is not established, the IP address of the alternate hostname should resolve to the public IP address on the device external port. • If the NC (Pulse) session is established, the IP address of the alternate hostname should resolve to the private IP address on the device internal port. BEST PRACTICE: The system uses HTTPS for this service. Therefore, we recommend that you assign a valid certificate to the interface that has the IP address to which this FQDN resolves so that users do not see invalid certificate warnings.
Update Entity IDs	Use this button to regenerate the SAML entity IDs of all configured service providers and identity providers. Typically, you take this action when the Host FQDN for SAML is changed.

Related Documentation

- Managing SAML Metadata Files

Managing SAML Metadata Files

You use the System > Configuration > SAML pages to maintain a table of SAML metadata files for the SAML service providers and identity providers in your network. Using SAML metadata files makes configuration easier and less prone to error.

You can add the metadata files to the system by:

- Uploading a metadata file.

- Retrieving the metadata file from a well-known URL.

To add metadata files:

1. Select **System > Configuration > SAML**.

2. Click **New Metadata Provider** to display the configuration page.

3. Complete the settings described in Table 37.

4. Save the configuration.

Table37: SAML Metadata Provider Configuration Guidelines

Settings	Guidelines
Metadata Provider Location Configuration	Select one of the following methods: • **Local**. Browse and locate the metadata file on your local host or file system. • **Remote**. Enter the URL of the metadata file. Only http and https protocols are supported.
Metadata Provider Verification Configuration	
Accept Untrusted Server Certificate	If you specify a URL for the metadata provider, select this option to allow the system to download the metadata file even if the server certificate is not trusted. This is necessary only for HTTPS URLs.
Accept Unsigned Metadata	If this option is not selected, unsigned metadata is not imported. Signed metadata is imported only after signature verification.

Signing Certificate	Browse and locate the certificate that verifies the signature in the metadata file. This certificate overrides the certificate specified in the signature of the received metadata. If no certificate is uploaded here then the certificate present in the signature of the received metadata is used.
	Select the **Enable Certificate Status Checking** option to verify the certificate before using it. Certificate verification applies both to the certificate specified here and the certificate specified in the signature in the metadata file.

Metadata Provider Filter Configuration

Roles	Select whether the metadata file includes configuration details for a SAML service provider, identity provider, or Policy Decision Point. You may select more than one. If you select a role that is not in the metadata file, it is ignored. If none of the selected roles are present in the metadata file, the system returns an error.
Entity IDs To Import	Enter the SAML Entity IDs to import from the metadata files. Enter only one ID per line. Leave this field blank to import all IDs. This option is available only for uploading local metadata files.

The Refresh button downloads the metadata files from the remote location even if these files have not been modified. This operation applies only to remote locations; local metadata providers are ignored if selected.

To refresh a metadata file:

1. Select **System > Configuration > SAML**.

2. Select the metadata file to refresh and click **Refresh**.

To delete a metadata file:

1. Select **System > Configuration > SAML**.

2. Select the metadata file to delete and click **Delete**.

Related Documentation

- Configuring Connect Secure as a SAML 2.0 Identity Provider

Configuring Connect Secure as a SAML 2.0 Service Provider

This topic describes how to configure the system as a SAML service provider. When the system is a SAML service provider, it relies on the SAML identity provider authentication and attribute assertions when users attempt to sign in to the device. Note that authentication is only part of the security system. The access management framework determines access to the system and protected resources.

The system supports:

- HTTP Redirect binding for sending AuthnRequests

- HTTP Redirect binding for sending/receiving SingleLogout requests/responses

- HTTP POST and HTTP Artifact bindings for receiving SAML responses

- RequestedAuthnContext context class specifications

Before you begin:

- Check to see whether the SAML identity provider uses HTTP POST or HTTP Artifact bindings for SAML assertions.

- Check to see whether the SAML identity provider has published a SAML metadata file that defines its configuration. If the SAML identity provider metadata file is available, configuration is simpler and less prone to error.

- Complete the system-wide SAML settings if you have not already done so. Select **System > Configuration > SAML > Settings**. For details, see "Configuring Global SAML Settings".

- Add metadata for the SAML identity provider to the metadata provider list if you have not already done so. Select **System > Configuration > SAML**. For details, see "Managing SAML Metadata Files".

The sign-in URL for which a session needs to be established for the system as a service provider is identified by the RelayState parameter (HTTP URL parameter for artifact and HTML form parameter for POST.) In a service provider initiated case, the system populates RelayState as an HTTP URL parameter while sending AuthnRequest. In the IdP-Initiated scenario (Connect Secure is a service provider and there is a third-party IdP), the IdP must be configured to set the appropriate Sign-in URL of the system in the RelayState parameter of the HTML form containing the SAML response. For more information, see the SAML 2.0 specification.

To configure the system as a SAML service provider:

1. Select **Authentication > Auth. Servers**.

2. Select **SAML Server** from the New list and then click **New Server** to display the configuration page.

3. Complete the settings as described in Table 38.

4. Save the configuration.

After you save changes for the first time, the page is redisplayed and now has two tabs. Use the Settings tab to modify any of the settings pertaining to the SAML server configuration. Use the Users tab to monitor user sessions.

Next steps:

- Configure the access management framework to use the SAML authentication server. Start with realm and role mapping rules. For details, see "Creating an Authentication Realm" and "Specifying Role Mapping Rules for an Authentication Realm".

- Configure a sign-in policy. When using a SAML authentication server, the sign-in policy can map to a single realm only. For details, see "Defining a Sign-In Policy".

Table38: SAML Service Provider Profile

Settings	Guidellines
Name	Specify a name to identify the server instance.
Settings	
SAML Version	Select **2.0**.
SA Entity Id	This value is prepopulated. It is generated by the system, based on the value for the Host FQDN for SAML setting on the System > Configuration > SAML > Settings page.
Configuration Mode	Select **Manual** or **Metadata**. If a metadata file or location is available from the SAML identity provider, use the metadata option to make configuration simpler and less prone to error. To upload or set the location for the published metadata file, select **System > Configuration > SAML** and click the **New Metadata Provider** button.

Identity Provider Entity ID	The identity provider entity ID is sent as the Issuer value in the assertion generated by the SAML identity provider.
	If you use the metadata option, this setting can be completed by selecting the identity provider entity ID from the list. The list is populated by the identity provider entities defined in metadata files added to the System > Configuration > SAML page.
	If you complete this setting manually, specify the Issuer value in assertions generated by the SAML identity provider. Typically, you ask the SAML identity provider administrator for this setting.
Identity Provider Single Sign On Service URL	The identity provider SSO service URL is a URL provisioned by the SAML identity provider. The setting is required to support service-provider-initiated SSO. If missing, the system cannot successfully redirect the user request.
	If you use the metadata option, this setting can be completed by selecting the SSO service URL from the list. The list is populated by the identity provider entities defined in metadata files added to the System > Configuration > SAML page.
	If you complete this setting manually, ask the SAML identity provider administrator for this setting.
User Name Template	Specify how the system is to derive the username from the assertion. If the field is left blank, it uses the string received in the NameID field of the incoming assertion as the username.
	If you choose a certificate attribute with more than one value, the system uses the first matched value. For example, if you enter <certDN.OU> and the user has two values for the attribute (ou=management, ou=sales), the system uses "management". To use all values, add the SEP attribute to the variable. For example, if you enter <certDN.OUT SEP=":">, the system uses "management:sales". The attributes received in the attribute statement in the incoming assertion are saved under userAttr. These variables can also be used with angle brackets and plain text. If the username cannot be generated using the specified template, the login fails. If the NameID filed of the incoming assertion is of type X509Nameformat, then the individual fields can be extracted using system variable "assertionNameDN".
Allowed Clock Skew (minutes)	Specify the maximum allowed difference in time between the system clock and the SAML identity provider server clock.
	NOTE: SAML is a time sensitive protocol. The time-based validity of a SAML assertion is determined by the SAML identity provider. If the SAML identity provider and SAML service provider clocks are askew, the assertion can be determined invalid, and you will receive the following error:
	"SAML Transferred failed. Please contact your system administrator. Detail: Failure: No valid assertion found in SAML response."
	We recommend you use NTP to ensure the clocks are synchronized and that you set an Allowed Clock Skew value that accommodates any expected or permissible skew.
Support Single Logout	Single logout is a mechanism provided by SAML for logging out a particular user from all the sessions created by the identity provider. Select this option if the system must receive and send a single logout request for the peer SAML identity provider.
	If you use the metadata option, the Single Logout Service URL setting can be completed by selecting the SLO service URL from the list. The list is populated by the identity provider entities defined in metadata files added to the System > Configuration > SAML page. The system sends Single Logout requests to this URL.
	In addition, if you use the metadata option, the Single Logout Response URL setting is completed based on your selection for Single Logout Service URL. If the identity provider has left this setting empty in its metadata file, the system sends the Single Logout response to the SLO service URL.
	If you complete these settings manually, ask the SAML identity provider administrator for guidance.
	The Support Single Logout service for the identity provider must present a valid certificate. For example, the hostname in a single logout request URL must be the same as the Common Name of the certificate presented by the identity provider of that hostname. If an invalid certificate is presented, the single logout feature may not work as intended.

SSO Method

Artifact	When configured to use the Artifact binding, the system contacts the Artifact Resolution Service (ARS) to fetch the assertion using SOAP protocol. If the ARS is hosted on a HTTPS URL, then the certificate presented by the ARS is verified by the system. For this verification to pass successfully, the CA of the server certificate issued to the identity provider ARS must be added to the trusted server CA on the system.

Complete the following settings to configure SAML using the HTTP Artifact binding:

- Source ID. Enter the source ID for the identity provider ARS. Source ID is Base64-encoded, 20-byte identifier for the identity provider ARS. If left blank, this value is generated by the system.
- Source Artifact Resolution Service URL. For metadata-based configuration, this field is completed automatically from the metadata file and is not configurable. For manual configurations, enter the URL of the service to which the SP ACS is to send ArtifactResolve requests. ArtifactResolve requests are used to fetch the assertion from the artifact received by it.
- SOAP Client Authentication. Select **HTTP Basic** or **SSL Client Certificate** and complete the related settings. If you use an SSL client certificate, select a certificate from the device certificate list.
- Select Device Certificate for Signing. Select the device certificate the system uses to sign the AuthnRequest sent to the identity provider SSO service. If you do not select a certificate, the system does not sign AuthnRequest.
- Select Device Certificate for Encryption. Select the device certificate the system uses to decrypt encrypted data received in the SAML response. The public key associated with the device certificate is used by the identity provider for encryption. |
| POST | When configured to use the POST binding, the system uses a response signing certificate to verify the signature in the incoming response or assertion. The certificate file must be in PEM or DER format. The certificate you select should be the same certificate used by the identity provider to sign SAML responses.

Complete the following settings to configure SAML using the HTTP POST binding:

- Response Signing Certificate. If you use the metadata-based configuration option, select a certificate from the list. The list is populated by the identity provider entities defined in metadata files added to the System > Configuration > SAML page.

  If you configure these settings manually, browse to and upload the certificate to be used to validate the signature in the incoming response or assertion.

  If no certificate is specified, the certificate embedded in the response is used.
- Enable Signing Certificate status checking. Select this option to check the validity of the signing certificate before verifying the signature. This setting applies to any certificate used for signature verification. If this option is enabled, the response will be rejected if the certificate is revoked, expired, or untrusted. If this option is selected, the certificate CA must be added to the system Trusted Client CA store.

  If this option is not enabled then the certificate is used without any checks.
- Select Device Certificate for Signing. Select the device certificate the system uses to sign the AuthnRequest sent to the identity provider SSO service. If you do not select a certificate, the system does not sign AuthnRequest.
- Select Device Certificate for Encryption. Select the device certificate the system uses to decrypt encrypted data received in the SAML response. The public key associated with the device certificate is used by the identity provider for encryption. |
| Authentication Context Classes | Use the **Add** and **Remove** buttons to select authentication context classes to be sent in the authentication requests to the SAML identity provider. These are included in the RequestedAuthnContext element.

In the OASIS standard, an authentication context is defined as "the information, additional to the authentication assertion itself, that the relying party may require before it makes an entitlements decision with respect to an authentication assertion."

This feature supports all authentication context classes specified in the SAML 2.0 OASIS Authn Context specification.

For example, if you select X509, the system sends the following context:

<samlp:RequestedAuthnContext>
    <saml:AuthnContextClassRef xmlns:saml="urn:oasis:names:tc:SAML:2.0:assertion">
        urn:oasis:names:tc:SAML:2.0:ac:classes:X509</saml:AuthnContextClassRef>
</samlp:RequestedAuthnContext>

In response, the SAML IdP sends the context data along with the authentication results. The system stores the context data in the session cache, and it can be specified in user attribute role mapping rules. |

Specify a comparison attribute within the RequestedAuthnContext element. The comparison attribute specifies the relative strengths of the authentication context classes specified in the request and the authentication methods offered by a SAML IdP. The following values specified in the SAML 2.0 OASIS core specification can be selected:

- **exact**——Requires the resulting authentication context in the authentication statement to be the exact match of at least one of the authentication contexts specified.

- **minimum**—Requires the resulting authentication context in the authentication statement to be at least as strong as one of the authentication contexts specified.

- **maximum**—Requires the resulting authentication context in the authentication statement to be stronger than any one of the authentication contexts specified.

- **better**—Requires the resulting authentication context in the authentication statement to be as strong as possible without exceeding the strength of at least one of the authentication contexts specified.

Select the same value that is configured on the SAML IdP. If none is specified in the SAML IdP configuration, the implicit default is exact.

Service Provider Metadata Settings	
Metadata Validity	Enter the number of days the system metadata is valid. Valid values are 0 to 9999. 0 specifies the metadata does not expire.
Do Not Publish SA Metadata	Select this option if you do not want the system to publish the metadata at the location specified by the system Service Entity ID field.
Download Metadata	This button appears only after you have saved the authentication server configuration. Use this button to download the metadata of the current SAML service provider.
User Record Synchronization	
Enable User Record Synchronization	Allow users to retain their bookmarks and individual preferences regardless of which device they log in to.
Logical Auth Server Name	Specify the server name if you have enabled user record synchronization.

Related Documentation

- Understanding SAML 2.0

- SAML 2.0 Supported Features Reference

- Investigating a "No valid assertion found in SAML response" Error

Configuring Connect Secure as a SAML 2.0 Identity Provider

This topic describes how to configure the system as a SAML identity provider. It includes the following sections:

- Configuration Overview

- Configuring Sign-in SAML Metadata Provider Settings

- Configuring Sign-in SAML Identity Provider Settings

- Configuring Peer SAML Service Provider Settings

- Configuring a SAML SSO Resource Policy for Gateway Mode Deployments

- Configuring a SAML External Applications SSO Policy

Configuration Overview

Implementing the system as a SAML identity provider includes the following basic steps.

1. Configure system-wide SAML settings. Select **System > Configuration > SAML > Settings**. See "Configuring Global SAML Settings".

2. Add SAML metadata provider files. Select **System > Configuration > SAML**. See "Managing SAML Metadata Files".

3. Configure Sign-In SAML metadata provider settings. See "Configuring Sign-in SAML Metadata Provider Settings".

4. Configure Sign-In SAML identity provider settings. See "Configuring Sign-in SAML Identity Provider Settings".

5. Configure peer service provider settings. See "Configuring Peer SAML Service Provider Settings".

6. Configure a resource policy:

   - For gateway mode deployments, configure a SAML SSO resource policy. See "Configuring a SAML SSO Resource Policy for Gateway Mode Deployments".

   - For peer mode deployments, configure a SAML SSO external applications policy. See "Configuring a SAML External Applications SSO Policy".

Configuring Sign-in SAML Metadata Provider Settings

Sign-in SAML metadata provider settings determine how the system identity provider metadata is published.

To configure the identity provider metadata publication settings:

1. Select **Authentication > Signing In > Sign-In SAML > Metadata Provider** to display the configuration page.

2. Complete the settings described in Table 39.

3. Click **Save Metadata Provider** to save your changes.

Table39: Sign-in SAML Identity Provider Metadata Provider Configuration Guidelines

Settings	Guidelines
Entity ID	This value is prepopulated. It is generated by the system, based on the value for the Host FQDN for SAML setting on the System > Configuration > SAML > Settings page.
Metadata Validity	Specify the maximum duration for which a peer SAML entity can cache the system SAML metadata file. Valid values are 1 to 9999. The default is 365 days.
Do Not Publish SA Metadata	Select this option if you do not want the system to publish the metadata at the location specified by the system Entity ID field. You can use this option to toggle off publication without deleting your settings.
Download Metadata	Use this button to download the system SAML identity provider metadata.

Configuring Sign-in SAML Identity Provider Settings

The settings defined in this procedure are the default settings for the system SAML identity provider communication with all SAML service providers. If necessary, you can use the peer service provider configuration to override these settings for particular service providers.

To configure sign-in SAML identity provider settings:

1. Select **Authentication > Signing In > Sign-In SAML > Identity Provider** to display the configuration page.

2. Complete the settings described in Table 40.

3. Save the configuration.

Table40: Sign-in SAML Identity Provider Configuration Guidelines

Settings	Guidelines
Basic Identity Provider (IdP) Configuration (Published in Metadata)	
Protocol Binding to use for SAML Response	Select **POST**, **Artifact**, or both, depending on your total requirements.
Signing Certificate	Select the certificate used to sign the SAML messages sent by the system. The certificates listed here are configured on the System > Configuration > Certificate > Device Certificates page.
Decryption Certificate	Select the certificate used to decrypt the SAML messages sent by peer service providers. The public key associated with this certificate is used by the peer service provider to encrypt SAML messages exchanged with this identity provider. The decryption certificate must be configured if the peer service provider encrypts the SAML messages sent to the system. The certificates listed here are configured on the System > Configuration > Certificate > Device Certificates page.
Other Configurations	Reuse Existing NC (Pulse) Session. This feature applies to an service-provider-initiated SSO scenario— that is, when a user clicks a link to log into the service provider site. The service provider redirects the user to the identity provider SSO Service URL.
	If this option is selected, a user with an active NC/Pulse session is not prompted to authenticate. The system uses information from the existing session to form the SAML response.
	Accept unsigned AuthnRequest. In an service-provider-initiated SSO scenario, the SP sends an AuthnRequest to the identity provider. This AuthnRequest could be either signed or unsigned. If this option is unchecked, the system rejects unsigned AuthnRequests. Note that the system also rejects signed AuthnRequests if signature verification fails.
Service-Provider-Related IdP Configuration	
Relay State	SAML RelayState attribute sent to the service provider in an identity-provider-initiated SSO scenario. If left blank, the RelayState value is the URL identifier of the resource being accessed.
Session Lifetime	Suggest a maximum duration of the session at the service provider created as a result of the SAML SSO. Select one of the following options:
	• **None**. The identity provider does not suggest a session duration.
	• **Role Based**. Suggest the value of the session lifetime configured for the user role.
	• **Customized**. If you select this option, the user interface displays a text box in which you specify a maximum in minutes.
Sign-In Policy	Select the sign-in URL to which the user is redirected in an service-provider-initiated scenario. The list is populated by the sign-in pages configured on the **Authentication > Signing In > Sign-in Policies** page.
	Note: The user is not redirected if he or she already has a session with the system and had authenticated through this sign-in policy.
Force Authentication Behavior	In an service-provider-initiated scenario, the service provider sends an AuthnRequest to the identity provider. If the service provider AuthnRequest includes the ForceAuthn attribute set to true and the user has a valid session, this setting determines how the identity provider responds. Select one of the following options:
	• **Reject AuthnRequest**. Do not honor the SAML SSO request.
	• **Re-Authenticate User**. Invalidate the user session and prompt for reauthentication.
	Note: This setting prevails over the Pulse session reuse setting.
User Identity	
Subject Name Format	Format of the NameIdentifier field in the generated assertion. Select one of the following options:
	• **DN**. Username in the format of DN(distinguished name).
	• **Email address**. Username in the format of an e-mail address.
	• **Windows**. Username in the format of a Windows domain qualified username.
	• **Other**. Username in an unspecified format.

Subject Name	Template for generating the username that is sent as the value of the NameIdentifier field in the assertion. You may use any combination of available system or custom variables contained in angle brackets and plain text.
Web Service Authentication	These settings apply when the HTTP Artifact binding is used.
Authentication Type	Method used to authenticate the service provider assertion consumer service to the identity provider on the system. Select one of the following options: • **None**. Do not authenticate the assertion consumer service. • **Username/Password**. If you select this option, use the controls to specify username and password settings. • **Certificate**. For certificate-based authentication, the Client CA of the service provider should be present in the system Trusted Client CA list (located on the System > Configuration > Certificates > Trusted Client CAs page).
Artifact Configuration	These settings apply when the HTTP Artifact binding is used.
Source ID	This is the Base64-encoded, 20-byte identifier of the Artifact Resolution Service on the identity provider.
Enable Artifact Response Signing and Encryption	If checked, the identity provider signs and encrypts the artifact response.
Attribute Statement Configuration	Attributes to be sent in SAML Attribute statements can be specified manually as name-value pairs, or you can configure an option to fetch name-value pairs from an LDAP server (or you can specify both manual entries and LDAP entries).
Attribute Name	An ASCII string.
Friendly Name	A more readable friendly name for the attribute. This is optional (an option included in the SAML standard).
Attribute Value	The attribute value can be specified as a hard-coded string, a custom variable, or a user attribute variable. System conventions for specifying user and custom tokens and variables apply. The value can be a combination of a string and a user or custom variable. For example: **Email::<customVar.email>**. The value can also be a combination of user and custom variables and hardcoded text. For example: **mydata=<USER><REALM><customVar.email>**.

Value Type	Select **Single-Valued** or **Multivalued**.
	
	A single-valued attribute can be a combination of a string and a user or custom variable.
	
	If there are multiple single-valued attributes configured with the same attribute names, they are combined and sent as a multivalued attribute.
	
	Select **Multivalue** if you want every individual token defined in the Attribute Value column to be sent as a separate AttributeValue. For example:
	
	    <element ref="saml:AttributeValue" minOccurs="0" maxOccurs="unbounded"/>
	
	If the Attribute value is given as <USER>mars<REALM>juniper<ROLE> and the value type is marked as Multivalued, then the values sent as part of attribute statement are sent as follows:
	
	    Username
	    Realmname
	    Role
	
	Note that only the tokens ['<>'] will be considered when processing a Multivalued attribute marked. The remaining data (for example mars, jupiter) is discarded.
	
	Specifying the token <ROLE> will send only one role. To send all roles, specify the Attribute value with the syntax <ROLE SEP=",">. If you specify <ROLE SEP=","> as a single-valued attribute, it is sent as a single string with "," separated roles. If you specify <ROLE SEP=","> as a multi-valued attribute, each role is sent in a separate <AttributeValue> element.
	
	NOTE: Encryption is set at the assertion level. You cannot encrypt individual attributes.

Directory Server	To fetch attribute name-value pairs from an LDAP server, complete the following settings:

- **Directory Server**—Select the LDAP server from the list. You must add the LDAP server to the Authentication > Auth. Servers list before it can be selected.

- **Username for lookup**—Enter a username template for LDAP lookup. The default is the variable <USERNAME>. The <USERNAME> variable stands for the login credential the user entered when logging in. The value can contain contextual characters as well as variables for substitution.

- **Attribute Name**—Type an LDAP attribute name, such as **cn**. The attribute name is fetched from the LDAP server and sent as SAML Attribute statements as part of a SAML assertion.

- **Friendly Name**—A more readable friendly name for the attribute. This is optional (an option included in the SAML standard).

NOTE: With the LDAP option, the SAML IdP sends attributes in the form configured on the backend LDAP server. If the LDAP server returns an attribute value in multi-valued form, then the SAML attribute statement will also be in multi-valued form.

Configuring Peer SAML Service Provider Settings

The peer service provider list defines the set of service providers configured to communicate with the system SAML identity provider. When you add a peer service provider to the list, you can customize the SAML identity provider settings used to communicate with the individual service provider. If the service provider provides a SAML metadata file, you can use it to simplify configuration, or you can complete more detailed manual steps. If available, we recommend you use metadata so that configuration is simpler and less prone to error.

To configure peer SAML service provider settings:

1. Select **Authentication > Signing In > Sign-In SAML > Identity Provider**.

2. Under Peer Service Provider Configuration, create a list of service providers that are SAML peers to the system SAML identity provider. To add a service provider to the list, click **Add SP** to display the configuration page.

3. Complete the settings described in Table 41.

4. Save the configuration.

Table41: Peer Service Provider Configuration Guidelines

Settings	Guidelines
Configuration Mode	Select **Manual** or **Metadata**.
Service Provider Configuration - Metadata	
Entity Id	If you use metadata, select the SAML entity ID of the service provider. This list contains all the service providers specified in all the metadata files added to the System > Configuration > SAML page.
Select certificates manually	When you use the metadata configuration, the system SAML identity provider iterates through all the signature verification certificates specified when verifying the incoming SAML messages coming from the service provider. Similarly, when encrypting the SAML messages going out, the system SAML identity provider encrypts the messages with the first valid encryption certificate encountered in the metadata. Select this option to override this default behavior and select certificates manually.
Signature Verification Certificate	If you select the **Select certificates manually** option, select the certificate to be used by the identity provider to verify the signature of incoming SAML messages.
Encryption Certificate	If you select the **Select certificates manually** option, select the certificate to be used if the assertions sent by the identity provider must be encrypted.
Service Provider Configuration - Manual	

Entity Id	If you are completing a manual configuration, ask the SAML service provider administrator for this setting.
Assertion Consumer Service URL	SAML service provider URL that receives the assertion or artifact sent by the identity provider.
Protocol Binding supported by the Assertion Consumer Service at the SP	Select **POST**, **Artifact**, or both. This setting must be consistent with the SAML identity provider configuration.
Default Binding	If both POST and Artifact bindings are supported, which is the default? • **Post** • **Artifact** This setting must be consistent with the SAML identity provider configuration.
Signature Verification Certificate	Upload the certificate to be used by the identity provider to verify the signature of incoming SAML messages. If no certificate is specified, the certificate embedded in the incoming SAML message is used for signature verification.
Encryption Certificate	Upload the certificate to be used if the assertions sent by the identity provider must be encrypted. If not certificate is specified, the assertions sent by the identity provider are not encrypted.
Certificate Attribute Configuration for Artifact Resolution Service	Optional. Specify attributes that must be present in the certificate presented to the Artifact Resolution Service (ARS) at the identity provider by the service provider assertion consumer service. This option appears only if the SAML service provider supports the HTTP Artifact binding, the system SAML identity provider has been configured to support the HTTP Artifact binding, and the Web service authentication type specified for the service provider is Certificate.

Certificate Status Checking Configuration

Enable signature verification certificate status checking	Select this option to enable revocation checks for the signing certificate. Uses the configuration on the System > Configuration > Certificates > Trusted Client CAs page.
Enable encryption certificate status checking	Select this option to enable revocation checks for the encryption certificate. Uses the configuration on the System > Configuration > Certificates > Trusted Client CAs page.

Customize identity provider Behavior

Override Default Configuration	Select this option to set custom behavior of the system SAML identity provider for this SP instance. If you select this option, the user interface displays the additional options listed next.
Reuse Existing NC (Pulse) Session	This option cannot be enabled here if it is not selected for the sign-in SAML identity provider default settings.
Accept unsigned AuthnRequest	Individual service providers can choose to accept unsigned AuthnRequest.
Relay State	SAML RelayState attribute sent to the service provider in an identity-provider-initiated SSO scenario. If left blank, the RelayState value is the URL identifier of the resource being accessed.
Session Lifetime	Suggest a maximum duration of the session at the service provider created as a result of the SAML SSO. Select one of the following options: • **None**. The identity provider does not suggest a session duration. • **Role Based**. Suggest the value of the session lifetime configured for the user role. • **Customized**. If you select this option, the user interface displays a text box in which you specify a maximum in minutes.
Sign-In Policy	Select the Sign-In URL to which the user is redirected in an service-provider-initiated scenario. The list is populated by the sign-in pages configured in Authentication > Signing In > Sign-in Policies. Note: The user is not redirected if he or she already has an active session and had authenticated through this sign-in policy.

Force Authentication Behavior	In an service-provider-initiated scenario, the service provider sends an AuthnRequest to the identity provider. If the service provider AuthnRequest includes the ForceAuthn attribute set to true and the user has a valid session, this setting determines how the identity provider responds. Select one of the following options:
	• **Reject AuthnRequest**. Do not honor the SAML SSO request.
	• **Re-Authenticate User**. Invalidate the user session and prompt for reauthentication.
	Note: This setting prevails over the Pulse session reuse setting.

User Identity

Subject Name Format	Format of NameIdentifier field in generated Assertion. Select one of the following options:
	• **DN**. Username in the format of DN (distinguished name).
	• **Email address**. Username in the format of an e-mail address.
	• **Windows**. Username in the format of a Windows domain qualified username.
	• **Other**. Username in an unspecified format.
Subject Name	Template for generating the username that is sent as the value of the NameIdentifier field in the assertion.
	You may use any combination of available system or custom variables contained in angle brackets and plain text.

Web Service Authentication: These settings apply when the HTTP Artifact binding is used.

Authentication Type	Method used to authenticate the service provider assertion consumer service to the identity provider on the system. Select one of the following options:
	• **None**. Do not authenticate the assertion consumer service.
	• **Username/Password**. Use the controls to specify username and password settings.
	• **Certificate**. For certificate-based authentication, the client CA of the service providers should be present in the system trusted client CA list (located on the System > Configuration > Certificates > Trusted Client CAs page).

Artifact Configuration: These settings apply when the HTTP Artifact binding is used.

Source ID	This is the Base64-encoded, 20-byte identifier of the Artifact Resolution Service on the identity provider.
Enable Artifact Response Signing and Encryption	If checked, the identity provider signs and encrypts the Artifact response.

Attribute Statement Configuration

Send Attribute Statements	Select this option if the SAML SP requires additional attributes to be sent with SAML assertions.
	If you enable attribute statements, select one of the following configuration options:
	• **Use IdP Defined Attributes**—Send attributes based on the default settings for the system SAML identity provider communication with all SAML service providers.
	• **Customize IdP Defined Attributes**—Selectively configure the attributes that are sent for this particular peer SAML SP. Attributes to be sent in SAML Attribute statements can be specified manually as name-value pairs, or you can configure an option to fetch name-value pairs from an LDAP server (or you can specify both manual entries and LDAP entries). If you select this option, configure the settings described next.
Attribute Name	An ASCII string.
Friendly Name	A more readable friendly name for the attribute. This is optional (an option included in the SAML standard).

Attribute Value	The attribute value can be specified as a hard-coded string, a custom variable, or a user attribute variable. System conventions for specifying user and custom tokens and variables apply.
	The value can be a combination of a string and a user or custom variable. For example: **Email::<customVar.email>**. The value can also be a combination of user and custom variables and hardcoded text. For example: **mydata=<USER><REALM><customVar.email>**.
Value Type	Select **Single-Valued** or **Multivalued**.
	A single-valued attribute can be a combination of a string and a user or custom variable.
	If there are multiple single-valued attributes configured with the same attribute names, they are combined and sent as a multivalued attribute.
	Select **Multivalue** if you want every individual token defined in the Attribute Value column to be sent as a separate AttributeValue. For example:
	`<element ref="saml:AttributeValue" minOccurs="0" maxOccurs="unbounded"/>`
	If the Attribute value is given as <USER>mars<REALM>juniper<ROLE> and the value type is marked as multivalued, then the values sent as part of attribute statement are sent as follows:
	Username Realmname Role
	Note that only the tokens ['<>'] will be considered when processing a multivalued attribute marked. The remaining data (for example mars, jupiter) is discarded.
	Specifying the token <ROLE> will send only one role. To send all roles, specify the Attribute value with the syntax <ROLE SEP=",">. If you specify <ROLE SEP=","> as a single-valued attribute, it is sent as a single string with "," separated roles. If you specify <ROLE SEP=","> as a multivalued attribute, each role is sent in a separate <AttributeValue> element.
	NOTE: Encryption is set at the assertion level. You cannot encrypt individual attributes.
Directory Server	To fetch attribute name-value pairs from an LDAP server, complete the following settings:
	• **Directory Server**—Select the LDAP server from the list. You must add the LDAP server to the Authentication > Auth. Servers list before it can be selected.
	• **Username for lookup**—Enter a username template for LDAP lookup. The default is the variable <USERNAME>. The <USERNAME> variable stands for the login credential the user entered when logging in. The value can contain contextual characters as well as variables for substitution.
	• **Attribute Name**—Type an LDAP attribute name, such as **cn**. The attribute name is fetched from the LDAP server and sent as SAML Attribute statements as part of a SAML assertion.
	• **Friendly Name**—A more readable friendly name for the attribute. This is optional (an option included in the SAML standard).
	NOTE: With the LDAP option, the SAML IdP sends attributes in the form configured on the backend LDAP server. If the LDAP server returns an attribute value in multivalued form, then the SAML attribute statement will also be in multivalued form.

Configuring a SAML SSO Resource Policy for Gateway Mode Deployments

When deployed as a gateway in front of enterprise resources, the SAML SSO policy acts like other resource policies. When deployed as a gateway, the SAML SSO communication is always identity-provider-initiated. The system maintains the session and uses its rewriting or pass-through proxy features to render data to the user. You use a SAML SSO resource policy when the protected resource supports SAML SSO and has been configured as a SAML service provider.

To configure a SAML SSO resource policy:

1. Select **Users > Resource Policies > Web**.

2. Use the tabs to display the SSO > SAML page.

If your administrator view is not configured to show SAML policies, click the Customize button in the upper-right corner of the page and select the **SSO** and **SAML** check boxes.

3. Click **New Policy** to display the configuration page.

4. Complete the settings described in Table 42.

5. Save the configuration.

Table42: SAML SSO Resource Policy Configuration Guidelines

Settings	Guidelines
Name	Type a name for the policy.
Description	Type a description that would be meaningful to other administrators.
Resources	Specify the fully qualified domain name for the resources for which this policy applies. These are the resources protected at the SAML service provider.
Roles	Select one of the following options: • **Policy applies to ALL roles**. To apply this policy to all users • **Policy applies to SELECTED roles**. To apply this policy only to users who are mapped to roles in the Selected roles list. Make sure to add roles to this list from the Available roles list. • **Policy applies to all roles OTHER THAN those selected below**. To apply this policy to all users except for those who map to the roles in the Selected roles list. Make sure to add roles to this list from the Available roles list.
Action	Select one of the following options: • **Use the SAML SSO defined below**. Typically, this is the setting you use for a SAML SSO resource policy. The system SAML identity provider makes the SSO request when a user tries to access a SAML resource specified in the Resources list. • **Do NOT use SAML**. The system does not perform an SSO request. Use this if there is a problem with the SAML service provider and you want to allow access. • **Use Detailed Rules**. Use this option to configure advanced rules.
SAML SSO Details	
SAML Version	Select **2.0**.
Service Provider Entity ID	Select the service provider entity ID. The service provider entity IDs listed here are configured on the Authentication > Signing In > Sign-in SAML > Identity Provider > Peer Service Provider pages.
Cookie Domain	Enter a comma-separated list of domains to which the system sends the SSO cookie.
Rewrite Response from SP	Select this option if the SAML service provider generates HTTP responses that require user/browser action, such as submission of a form, JavaScript execution, redirection to a different location, and other similar behavior. If you select this option, the system rewrites the HTTP responses sent by the SAML service provider and sends them to the user.

NOTE: The SAML SSO resource policy settings changed from Release 7.1 to 7.2. Policies you created with Release 7.1 are preserved in edit-only mode for legacy use.

Configuring a SAML External Applications SSO Policy

When deployed to support access to external resources (for example, public cloud resources), the system does not have to be a gateway to user access. The user can access the external resource directly, and the traffic does not flow through the device. To enable SAML SSO in these deployments, you configure the system as a SAML identity provider to correspond with

the external SAML service provider, and you configure a SAML external applications SSO policy to determine the users and resources to which the SAML SSO experience applies.

To configure a SAML External Apps SSO resource policy:

1. Select **Users > Resource Policies > Web**.

2. Use the tabs to display the SSO > SAML External Apps SSO page.

   If your administrator view is not configured to show SAML policies, click the Customize button in the upper-right corner of the page and select the **SSO** and **SAML** check boxes.

3. Click **New Policy** to display the configuration page.

4. Complete the settings described in Table 43.

5. Click **Save Changes**.

Table43: SAML SSO External Applications Policy Configuration Guidelines

Settings	Guidelines
Name	Type a name for the policy.
Description	Type a description that would be meaningful to other administrators.
Resources	Specify the fully qualified domain name for the resources for which this policy applies. These are the resources protected at the SAML service provider.
Roles	Select one of the following options: • **Policy applies to ALL roles**. To apply this policy to all users. • **Policy applies to SELECTED roles**. To apply this policy only to users who are mapped to roles in the Selected roles list. Make sure to add roles to this list from the Available roles list. • **Policy applies to all roles OTHER THAN those selected below**. To apply this policy to all users except for those who map to the roles in the Selected roles list. Make sure to add roles to this list from the Available roles list.
Action	Select one of the following actions: • **Use the SAML SSO defined below**. Typically, this is the setting you use for a SAML SSO resource policy. The system SAML identity provider makes the SSO request when a user tries to access to a SAML resource specified in the Resources list. • **Do NOT use SAML**. The system does not perform an SSO request. Use this if there is a problem with the SAML service provider and you want to allow access. • **Use Detailed Rules**. Use this option to configure advanced rules.
SAML SSO Details	
Service Provider Entity ID	Select the service provider entity ID. The service provider entity IDs listed here are configured on the Authentication > Signing In > Sign-in SAML > Identity Provider > Peer Service Provider pages.

Related Documentation

- Example: Implementing SAML 2.0 Web Browser SSO for Google Apps

- SAML 2.0 Supported Features Reference

Configuring a SAML 2.0 ACL Web Policy

To configure the system as a policy enforcement point, you must create a SAML ACL web policy.

To configure a SAML ACL web policy:

1. In the admin console, select **Users > Resource Policies > Web**.

2. Use the tabs to display the Access > SAML ACL page.

   If your administrator view is not configured to show SAML policies, click the Customize button in the upper-right corner of the page and select the **SAML ACL** check box.

3. On the SAML Access Control Policies page, click **New Policy**.

4. Complete the settings described in Table 44.

5. Click **Save Changes**.

6. On the SAML Access Control Policies page, order the policies according to how you want the system to evaluate them. Keep in mind that once the system matches the resource requested by the user to a resource in a policy's (or a detailed rule's) Resource list, it performs the specified action and stops processing policies.

Table44: SAML ACL Web Policy Settings

Setting	Description
Name	Type a name for the policy.
Description	Type a description that would be meaningful to other administrators.
Resources	Specify the fully qualified domain name for the resources for which this policy applies. These are the resources protected at the SAML service provider.
Roles	Select one of the following options: • **Policy applies to ALL roles**. To apply this policy to all users. • **Policy applies to SELECTED roles**. To apply this policy only to users who are mapped to roles in the Selected roles list. Make sure to add roles to this list from the Available roles list. • **Policy applies to all roles OTHER THAN those selected below**. To apply this policy to all users except for those who map to the roles in the Selected roles list. Make sure to add roles to this list from the Available roles list.
Action	Select one of the following options: • **Use the SAML Access Control checks defined below**. The system performs an access control check to the specified URL using the data specified in the SAML Access Control Details section. • **Do not use SAML Access**. The system does not perform an access control check. • **Use Detailed Rules**. Use this option to configure advanced rules.
SAML Access Control Details	SAML Version. Select **2.0**. Configuration Mode. If you select manual, complete the SAML Access Control details. If you select **Metadata**, select the policy decision point to use. If the metadata option is disabled, you have not defined or uploaded a metadata file on the System > Configuration > SAML page. SAML Web Service URL. Completed automatically if using metadata. If you configure manually, enter the URL of the access management system SAML server. For example, enter **https://hostname/ws**. SAML Web Service Issuer. Enter the hostname of the issuer, typically the hostname of the access management system. NOTE: You must enter unique string that the SAML Web service uses to identify itself in authorization assertions.

Authentication Type	Select one of the following options: • **None**—Do not authenticate the system. • **Username**—Authenticate using a username and password. Enter the username and password that the system must send the Web service. • **Certificate Attribute**—Authenticate using a certificate signed by a trusted certificate authority. If you have more than one certificate installed on the system, use the drop-down list to select which certificate to send to the Web service.
User Identity	Subject Name Type—Specify which method the system and SAML Web service should use to identify the user. Select one the following options: • **DN**—Send the username in the format of a DN (distinguished name) attribute. • **Email Address**—Send the username in the format of an e-mail address. • **Windows**—Send the username in the format of a Windows domain qualified username. • **Other**—Send the username in another format agreed upon by the system and the SAML Web service. Subject Name—Use variables to specify the username to the SAML Web service. Or, enter static text. NOTE: You must send a username or attribute that the SAML Web service will recognize. Device Issuer—Enter a name that uniquely identifies the SAML authority, such as the device hostname.
Maximum Cache Time	You can eliminate the overhead of generating an authorization decision each time the user requests the same URL by indicating that the system must cache the access management system's authorization responses. Enter the amount of time the system should cache the responses (in seconds).
Ignore Query Data	By default, when a user requests a resource, the system sends the entire URL for that resource (including the query parameter) to the SAML Web service and caches the URL. You can specify that the system should remove the query string from the URL before requesting authorization or caching the authorization response.

Related Documentation

- Understanding SAML 2.0

- SAML 2.0 Supported Features Reference

Example: Implementing SAML 2.0 Web Browser SSO for Google Apps

This example shows how to implement SAML 2.0 Web browser SSO for Google Apps. It includes the following sections:

- Topology

- Configuring the Google Apps SAML Service Provider

- Configuring the Connect Secure SAML Identity Provider

- Verifying the Google Apps SAML SSO Deployment

Topology

When deployed to support access to external resources (for example, public cloud resources), the system does not have to be a gateway to user access. The user can access the external resource directly, and the traffic does not flow through the device. You configure the system as a SAML identity provider to correspond with the external SAML service provider, and you configure a SAML SSO external applications policy to determine the users and resources to which the SAML SSO experience applies.

When you configure the SAML identity provider, some settings are necessary to support either identity-provider-initiated or service-provider-initiated SSO. The documentation for the configuration steps makes note of these settings. Regardless, you configure the SAML identity provider to support both identity-provider-initiated and service-provider-initiated SSO.

Figure 33 illustrates the flow of network communication in a service-provider-initiated SSO scenario.

Figure33: Connect Secure as a SAML Identity Provider (Peer Mode) in a Service-Provider-Initiated SSO Scenario

1 – The user clicks a link to access a resource.

2a – The service provider sends an HTTP redirect status code (HTTP 302) to the user. The SAML request and all other SAML details are sent as URL parameters in the URL Location header.

2b – The user sends an HTTP GET request to the identity provider. The SAML request and all other SAML details are sent as URL parameters.

If the user already has a session with the identity provider, steps 2.1 and 2.2 are skipped.

2.1 – If the user does not have a session, the identity provider sends an authentication challenge to the user.

2.2 – The user enters sign-in credentials.

3a – The identity provider sends a successful status code (HTTP 200 OK) to the user with a form in the HTML body.

3b – The user sends the form to the service provider.

4 – The external resource is delivered to the user's browser.

Figure 34 illustrates the flow of network communication in an identity-provider-initiated SSO scenario.

Figure34: Connect Secure as a SAML Identity Provider (Peer Mode) in an Identity-Provider-Initiated SSO Scenario

1 – The user authenticates to the identity provider.
2 – The identity provider returns a portal page with links to external resources.
3 – The user clicks a link for an external resource.
4a – The identity provider sends a successful status code (HTTP 200 OK) to the user with a form in the HTML body.
4b – The user sends the form to the service provider.
5 – The external resource is delivered to the user's browser.

Configuring the Google Apps SAML Service Provider

To configure the Google Apps SAML service provider:

1. Log into the Google Apps control panel. The URL is similar to the following: **https://www.google.com/a/cpanel/acmegizmo.com**.

2. Click **Advanced Tools** in the menu bar.

3. Click the **Set up single sign-on (SSO)** link to display its configuration page, as shown in Figure 35.

4. Configure the SAML service provider settings as described in Table 45.

5. Click **Save Changes**.

Figure35: Google Apps Advanced Tools: SSO

Table45: Google Apps SSO Configuration

Settings	Guidelines
Enable Single Sign-on	Select this option.
Sign-in page URL	Type the URL of the system SAML SSO service. The URL formed with the primary host FQDN for SAML has the following form: **https://*SAMLHostName*/dana-na/auth/saml-sso.cgi** For example: **https://i5.lab.juniper.net/dana-na/auth/saml-sso.cgi** The URL formed with the alternate host FQDN for SAML (to support Pulse/NC session detection) has the following form: **https://i5pulse.lab.juniper.net/dana-na/auth/saml-sso.cgi**

Sign-out page URL	We recommend using the URL for the sign-in page for the realm associated with the system SAML identity provider. Users who already have a session will be directed to the sign-in page and can decide whether to log out from the system or not. The default sign-in URL has the form: **https://**_FQDN_ For example: **https://i5.lab.juniper.net/**
Change password URL	We recommend using the URL for the sign-in page for the realm associated with the system SAML identity provider. The system provides password management capabilities for some back-end auth servers (such as AD, LDAP, or Local Auth). When implemented, the password management capabilities are accessed from the sign-in page. The default sign-in URL has the form: **https://**_FQDN_ For example: **https://i5.lab.juniper.net/**
Verification certificate	Click **Browse** and select the device certificate. Then click **Upload** and ensure that the certificate is saved.
Use a domain specific issuer	Select this option.
Network masks	Do not select.
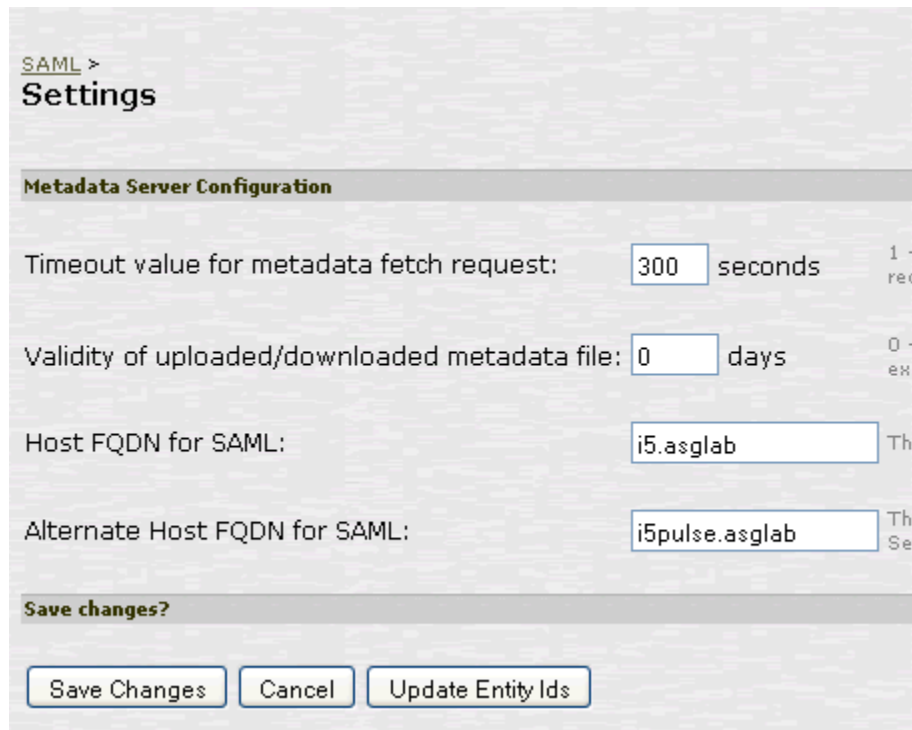
Configuring the Connect Secure SAML Identity Provider

You configure the system SAML identity provider settings to match the Google Apps SAML service provider settings.

To configure the SAML identity provider settings:

1. Select **System > Configuration > SAML > Settings** to complete the global SAML settings. These settings apply to all of your SAML deployments. Figure 36 shows an example of SAML global settings.

Figure36: SAML Global Settings

2. Select **Authentication > Signing In > Sign-In SAML > Identity Provider** to configure SAML identity provider settings. These settings apply to all of your deployments where the device is a SAML identity provider. Figure 37 shows an example of SAML identity provider settings.

Figure37: SAML Identity Provider Settings

3. On the SAML Sign-In Identity Provider page, click **Add SP** and complete the settings for communication with Google Apps. Google Apps does not publish metadata, so the configuration is manually configured. The Google SAML service provider uses the HTTP

POST binding and takes usernames in e-mail address format. Figure 38 shows an example of SP settings for Google Apps.

Figure38: Peer SP Settings

4. Select **Users > Resource Policies > Web > SAML External Apps SSO** and complete settings for the external applications policy that controls the users and the resources that can use the

SSO implementation. Figure 39 shows an example of an SAML external applications SSO policy for Google Apps.

Figure39: SAML External Apps SSO Policy Settings

Verifying the Google Apps SAML SSO Deployment

Access a Google Docs or Google Apps resource as a non-admin user to verify the solution works as expected.

> **TIP:** Use a browser plugin such as HTTP Watch if you want to trace the SAML
> communication between the SAML service provider and SAML identity provider.

To verify service-provider-initiated SSO:

1. Make sure you are not logged into the device or Google.

2. Open a Web browser and open a location on Google Docs or Google Apps. Google Apps
   redirects you to the sign-in page to authenticate.

3. Log in.

   The access management framework processes the authentication request, performs host
   checking rules and role mapping rules. If authentication is successful, the system redirects
   you to the Google Docs or Google Apps location you had requested.

To verify Pulse/NC session detection for service-provider-initiated SSO:

1. Make sure you are not logged into the device or Google.

2. Use Pulse or VPN tunneling client to create an SSL VPN connection.

3. Open a Web browser and open a location on Google Docs or Google Apps.

   You should not have to authenticate to access the Google Docs or Google Apps location.

To verify identity-provider-initiated SSO:

1. Use the system admin console to create a bookmark to a location on Google Docs or
   Google Apps.

2. As a user, log in to the device.

3. Click the bookmark link to the Google Docs or Google Apps location.

   You should not have to authenticate to access the Google Docs or Google Apps location.

Related Documentation

- Google Code: SAML Single Sign-on (SSO) Service for Google Apps

- Google Support: SSO (Single Sign-on)

- Configuring Connect Secure as a SAML 2.0 Identity Provider

Using SAML AuthnContext Class Variables in Role Mapping and Web ACL Rules

This topic describes how to use Security Assertion Markup Language (SAML) AuthnContext
class variables in access management framework rules. For information about SAML
AuthnContext class variables, refer to the SAML 2.0 OASIS Authn Context specification. This topic
includes the following information:

- Configuring SAML AuthnContext Class Variables in the Authentication Server Configuration

- Configuring a Role Mapping Rule Based on a SAML AuthnContext Class Variable

- Configuring a Web ACL Policy Rule Based on a SAML AuthnContext Class Variable

- Using Policy Tracing Logs to Verify the SAML AuthnContext Class Variable Is Used in Rules

Configuring SAML AuthnContext Class Variables in the Authentication Server Configuration

In deployments where the system is a SAML service provider (SAML SP), you can configure
the SAML SP to request authentication context classes from the SAML identity provider (SAML
IdP). The SAML SP includes these in the RequestedAuthnContext element. In response, the
SAML IdP sends the context data along with the authentication results.

The system stores the authnContext data in the session cache. You can use the system variable named samlAuthnContextClass to create rules based on AuthnContext in role mapping and resource policies.

To specify the SAML AuthnContext class variables in the SAML SP configuration:

1. Select **Authentication > Auth. Servers**.

2. Create a new SAML server configuration or edit one you have already created.

   Figure 40 shows the SAML server configuration page. Red boxes highlight the configuration elements for AuthnContext classes.

3. Select the AuthnContext classes that you want to request from the SAML IdP, and select a comparison method.

   This feature supports all authentication context classes described in the SAML 2.0 OASIS Authn Context specification.

   The comparison method values are defined in the SAML 2.0 OASIS core specification. You should specify the same values that have been configured on the SAML IdP. If none is specified in the SAML IdP configuration, the implicit default is exact.

4. Save the configuration.

Figure40: Authentication Server Configuration Page
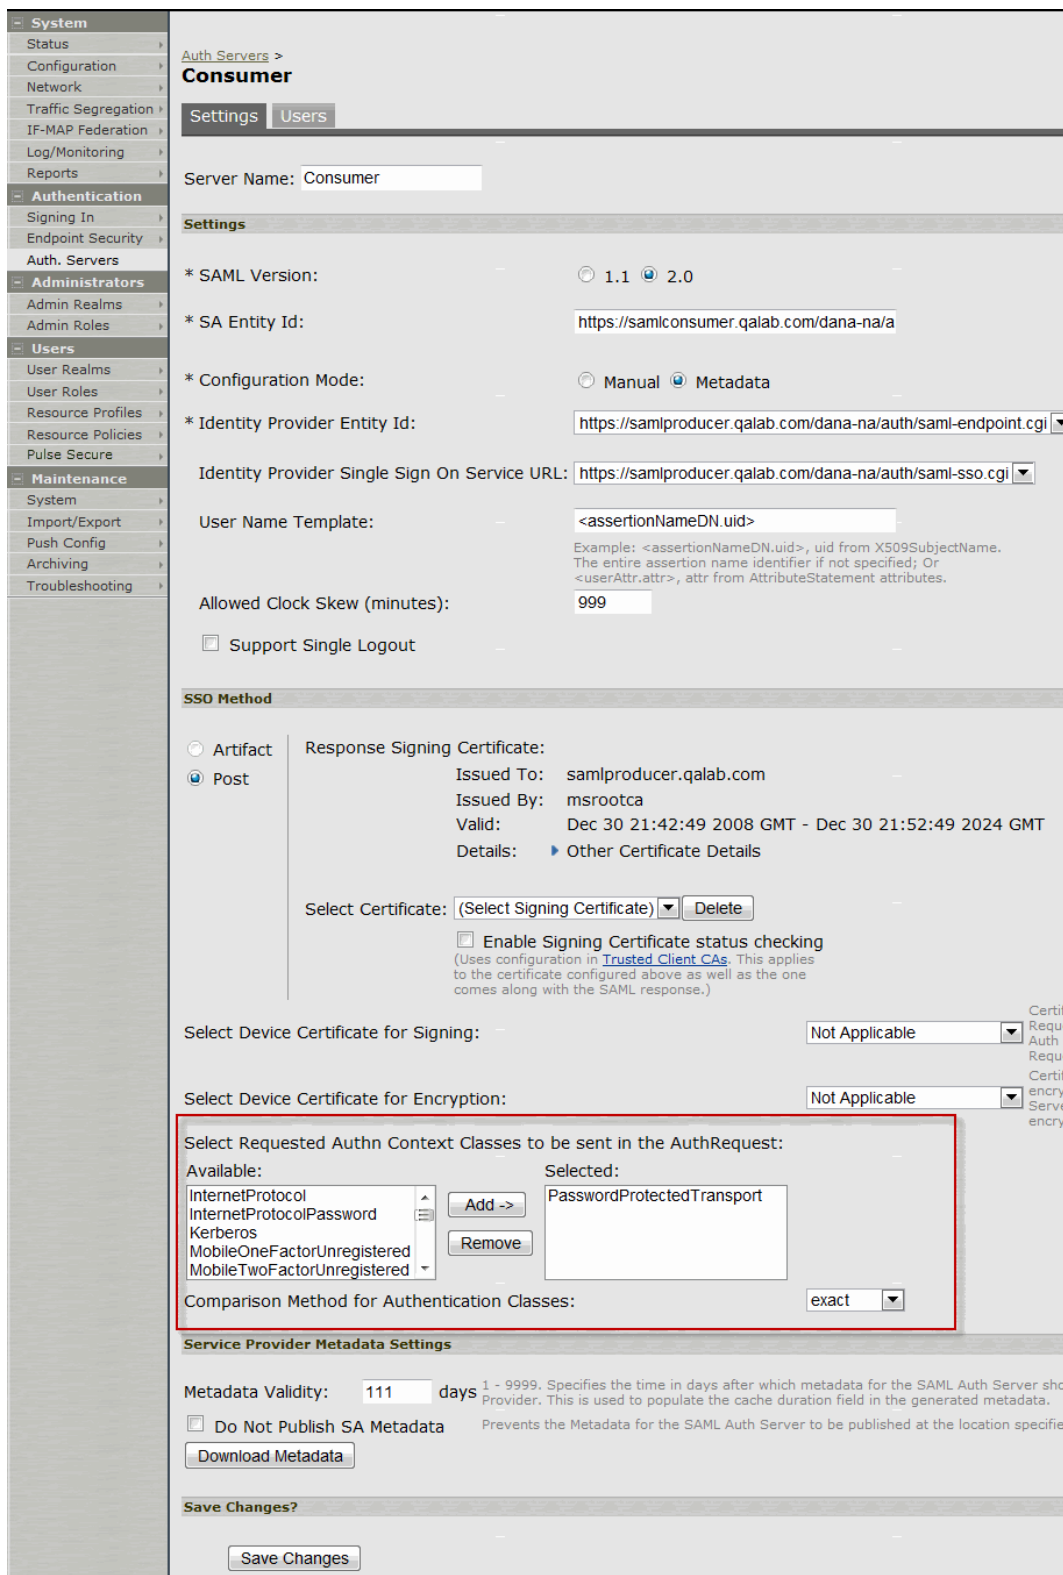
Configuring a Role Mapping Rule Based on a SAML AuthnContext Class Variable

You can use role mapping rule custom expressions to include AuthnContext class data as a factor in role determination.

To configure role mapping rules:

1. Select **Users > User Realms** .

2. Create a new realm or edit a realm you have already created.

3. Click **New Rule** to display the configuration page.

4. Select **Custom Expression** and click **Update** to redisplay the configuration page with the controls related to custom expressions.

   Figure 41 shows the configuration page.

5. Click **Expressions** to display the server catalog dialog box.

   Figure 42 shows the dialog box.

6. Select **samlAuthnContextClass**, select an operator, and click **Insert Expression**.

7. Edit the expression template to match the AuthnContextClassRef data expected from the SAML IdP.

8. Save your changes to the variable expression and return to the rule configuration page.

9. Select the expression, roles for the rule, and the stop option (if desired).

10.   Save your changes to the rule configuration and return to the realm configuration page.

11.   Reorder the rules if necessary.

12.   Save the realm configuration.

Figure41: Authentication Server Configuration Page

Figure42: Server Catalog Expressions and Variables

Configuring a Web ACL Policy Rule Based on a SAML AuthnContext Class Variable

You can use the resource policy detailed rules configuration to include AuthnContext class data as a factor in resource access determinations. This example shows how to use a SAML AuthnContext class variable in Web ACL detailed rules. In the same manner, you can use the AuthnContext class variable in detailed rules for other resource policies.

To configure a resource policy:

1. Select **Resource Policies > Web > Web ACL**.

2. Create a new policy or edit a policy you have already created.

3. Click the **Detailed Rules** tab for the policy.

4. Click **New Rule** to display the detailed rules configuration page.

   Figure 43 shows the detailed rule configuration page.

5. Under Conditions, select **samlAuthnContextClass**, select an operator, and click **Insert Expression**.

6. Edit the condition expression template to match the AuthnContextClassRef data expected from the SAML IdP.

7. Select a rule action and resources to which the rule applies, and save your changes to return to the policy configuration page.

8. Reorder the rules if necessary.

9. Save the configuration.

Figure43: Detailed Rule Configuration Page

Using Policy Tracing Logs to Verify the SAML AuthnContext Class Variable Is Used in Rules

You can use policy tracing logs to verify your configuration.

To create a policy trace log:

1. Select **Troubleshooting > User Sessions > Policy Tracing** to display the configuration page.

2. Specify the username, realm, and source IP address if you know it. If you provide the source IP address, the policy trace log can include events that occur before the user ID is entered into the system.

3. Select the events to trace.

4. Click **Start Recording**.

5. Initiate the action you want to trace, such as a user sign in.

6. Click **View Log** to display the policy trace results log.

7. Click **Stop Recording** when you have enough information.

Figure 44 shows policy trace results. The highlighted entries show the data populated to the samlAuthnContextClass system variable, as well as the role mapping rule that was matched.

Figure44: Policy Tracing Results

Related Documentation

- SAML 2.0 Supported Features Reference

- Configuring Connect Secure as a SAML 2.0 Service Provider

- Using the SAML Server

- Using Custom Expressions in Rule Configuration

Investigating a "No valid assertion found in SAML response" Error

Problem

Description: SAML is a time sensitive protocol. The time-based validity of a SAML assertion is determined by the SAML identity provider. If the SAML identity provider and SAML service provider clocks are askew, the assertion can be determined invalid, and authentication fails.

Environment: In the scenario described here, the system is deployed as a SAML service provider in a SAML 2.0 deployment.

Symptoms: In this scenario, the following error is returned to the user after the user has submitted credentials to the SAML identity provider:

"`SAML Transferred failed. Please contact your system administrator.`

`Detail: Failure: No valid assertion found in SAML response.`"

Cause

To investigate the error:

1. Select **Maintenance > Troubleshooting > Monitoring > Debug Logs** to display the Debug Log configuration page, shown in Figure 45.

Figure45: Debug Log Page

2. Turn debug logging on, set Debug Log Detail Level to **10**, and Event Codes to **saml**.

3. Reproduce the action that results in the error—in this case, user access to the resource associated with the SAML service provider that prompts the user to submit credentials to the SAML identity provider.

4. Click **Save Debug Log**.

   The console displays the Save As dialog box.

5. Save the file to a location your local host or a location that you can access when sending mail. The file is an encrypted file, so do not try to open it and analyze it yourself.

6. E-mail the debug log to Pulse Customer Support.

Pulse Customer Support will use the file to diagnose the issue. In the debug log, the following log lines indicate issues with the time-based validity of the assertion:

`verifySubjectConfirmationData: assertion has expired`

`processConditions: assertion has expired [NotOnOrAfter condition failed]`

`processConditions: assertion is not yet Valid [NotBefore condition failed]`

These log lines indicate a clock sync issue only if failure of the time-based validity check is unexpected. The same log lines might appear in the debug log to indicate an assertion has expired as expected.

Solution

We recommend you use NTP to ensure the clocks are synchronized and that you set an Allowed Clock Skew value that accommodates any expected or permissible skew. Properly synchronized clocks avoid unexpected failure.

To configure NTP:

1. Select **System > Status** to display the System Status page.

2. Next to System Date & Time, click **Edit** to display the Date and Time page.

3. Specify the settings for the same NTP server used by the SAML identity provider.

4. Save your configuration.

To set the Allowed Clock Skew value:

1. Select **Authentication > Auth. Servers**.

2. Select the SAML authentication server you want to configure to display its configuration page.

3. Specify a number of minutes in the **Allowed Clock Skew** to accommodate any expected or permissible skew.

4. Save your configuration.

Related Documentation

- Configuring Connect Secure as a SAML 2.0 Service Provider

Pulse Connect Secure SAML 1.1 Support

The trend in SAML deployments is converging on the SAML 2.0 specification. Connect Secure Release 7.2 and later continues to support SAML 1.1. The following sections reprint previous information we have provided about SAML 1.1 deployments:

- About SAML Version 1.1
- SAML Version 1.1 Configuration Tasks

About SAML Version 1.1

The following topics provide background information about SAML version 1.1:

- Understanding SAML 1.1
- Understanding SAML 1.1 Profiles
- Understanding SAML 1.1 Assertions

- Creating a Trust Relationship Between SAML 1.1 Systems

Understanding SAML 1.1

The system enables you to pass user and session state information between the device and another trusted access management system that supports the Security Assertion Markup Language (SAML). SAML provides a mechanism for two disparate systems to create and exchange authentication and authorization information using an XML framework, minimizing the need for users to re-enter their credentials when accessing multiple applications or domains.

SAML exchanges are dependent upon a trusted relationship between two systems or domains. In the exchanges, one system acts as a SAML authority (also called an asserting party or SAML responder) that asserts information about the user. The other system acts as a relying party (also called a SAML receiver) that relies on the statement (also called an assertion) provided by the SAML authority. If it chooses to trust the SAML authority, the relying party authenticates or authorizes the user based on the information provided by the SAML authority.

The system supports two SAML use case scenarios:

- The system as the SAML authority—The user signs into a resource by way of the device first, and all other systems are SAML receivers, relying on the system for authentication and authorization of the user. Under this scenario, the system can use either an artifact profile or a POST profile.

- The system as the SAML receiver—The user signs into another system on the network first, and the system is the SAML receiver, relying on the other system for authentication and authorization of the user.

For example, in the first scenario, an authenticated user named John Smith may try to access a resource protected by an access management system. When he does, the system acts as a SAML authority and declares "This user is John Smith. He was authenticated using a password mechanism." The access management system (the relying party) receives this statement and chooses to trust the system (and therefore trust that the system has properly identified the user). The access management system may still choose to deny the user access to the requested resource (for instance, because John Smith has insufficient access privileges on the system), while trusting the information sent by the system.

In the second scenario, John Smith signs in to his company portal and is authenticated using an LDAP server sitting behind the company's firewall. On the company's secure portal, John Smith clicks a link to a resource protected by the system. The following process occurs:

1. The link redirects John Smith to an intersite transfer service on the company portal, which constructs an artifact URL. The artifact URL contains a reference to a SAML assertion stored in the company portal's cache.

2. The portal sends the URL to the system, which can decide whether or not to link to the reference.

3. If the system links to the reference, the portal sends a SOAP message containing the SAML assertion (an XML message containing the user's credentials) to the system, which can then decide whether or not to allow the user access to the requested resource.

   > *i* NOTE: SOAP requests generated by the system (when configured as a SAML 1.1 consumer) are not signed.

4. If the system allows the user access, the system presents to the user the requested resource.

5. If the system rejects the SAML assertion, or the user credentials, the system responds to the user with an error message.

When configuring the system, you can use SAML for:

- Single sign-on (SSO) authentication—In a SAML SSO transaction, an authenticated user is seamlessly signed into another system without resubmitting his credentials. In this type of transaction, the system can be either the SAML authority or the SAML receiver. When acting as the SAML authority, the system makes an authentication statement, which declares the user's username and how he was authenticated. If the relying party (called an assertion consumer service in SAML SSO transactions) chooses to trust the system, the user is seamlessly signed into the assertion consumer service using the username contained in the statement.

  When acting as the SAML receiver, the system requests credential confirmation from the SAML authority, which is the other access management system, such as LDAP or another authentication server. The SAML authority sends an assertion by way of a SOAP message. The assertion is a set of XML statements that the system must interpret, based on criteria that the system administrator has specified in a SAML server instance definition. If the system chooses to trust the asserting party, the system allows the user to sign in seamlessly using the credentials contained in the SAML assertion.

- Access control authorization—In a SAML access control transaction, the system asks an access management system whether the user has access. In this type of transaction, the system is the relying party (also called a policy enforcement point in access control transactions). It consumes and enforces an authorization decision statement provided by the access management system (SAML authority), which declares what the user is allowed to access. If the SAML authority (also called a policy decision point in access control transactions) declares that the user has sufficient access privileges, the user may access the requested resource

  The system does not generate authorization decision statements—it only consumes them.

  In addition to providing users access to a URL based on the authorization decision statement returned by a SAML authority, the system also allows you to define users' access rights to a URL using system-only mechanisms (Users > Resource Profiles > Web Applications/Pages tab). If you define access controls through the system as well as through a SAML authority, both sources must grant access to a URL for a user to access it. For example, you may configure a access policy that denies members of the "Users" role access to www.google.com, but configure another SAML policy that bases a user's access rights on an attribute in an access management system. Even if the access management system permits users access to www.google.com, users are still denied access based on the access policy.

  When asked if a user may access a resource, access management systems that support SAML may return a response of permit, deny, or indeterminate. If the system receives an indeterminate response, it denies the user access.

  The session timeouts on the system and your access management system may not coordinate with one another. If a user's access management system session cookie times out before his destination signaling identifier (DSID) cookie times out, then single sign-on between the two systems is lost. The user is forced to sign in again when he times out of the access management system.

Related Documentation

- Configuring SAML 1.1 SSO Profiles

Understanding SAML 1.1 Profiles

The system accepts authentication assertions generated by a SAML authority using either an artifact profile or a POST profile. This feature allows a user to sign in to a source site or portal without going through the system first, and then to access the system with single sign-on (SSO) through the SAML consumer service.

As a result, the user who authenticates elsewhere can access resources behind the device without signing in again.

Using the Artifact Profile and POST Profile

The two supported profiles provide different methods of accomplishing the same task. The end user's goal is to sign in to all desired resources once, without experiencing multiple sign-in pages for different resources or applications. Although the end user wants transparency, you, the administrator, want to ensure complete security across the resources on your system, regardless of the servers or sites represented.

The artifact profile requires that you construct an automated request-response HTTP message that the browser can retrieve based on an HTTP GET request.

The POST profile requires that you construct an HTML form that can contain the SAML assertion, and which can be submitted by an end user action or a script action, using an HTTP POST method.

Using the Artifact Profile Scenario

The SAML server generally supports the following artifact profile scenario:

1. The user accesses a source site through a browser. The source site might be a corporate portal using a non-Connect Secure authentication access management system.

2. The source site challenges the user for username and password.

3. The user provides username and password, which the source site authenticates through a call to an LDAP directory or other authentication server.

4. The user then clicks a link on the source site, which points to a resource on a server that is protected behind the device.

5. The link redirects the user to the intersite transfer service URL on the source site. The source site pulls an authentication assertion message from its cache and encloses it in a SOAP message. The source site constructs a SAML artifact (a Base64 string) that it returns to the browser in a URI along with the destination and assertion address.

6. The destination site queries the authenticated assertion from the source site, based on the artifact it receives from the source site.

7. The system accepts the assertion as a valid authentication if the elapsed time falls within the allowable clock skew time. If the user also meets the other policy restrictions, the system grants the user access to the requested resource.

The main tasks you are required to fulfill to support the system as the relying party with the artifact profile include:

- Implement the assertion consumer service, which:

  - Receives the redirect URL containing the artifact.

  - Generates and sends the SAML request.

  - Receives and processes the SAML response.

- Integrate the assertion consumer service with the existing system process, which:

  - Maps the SAML assertion to a local user.

- Creates a user session.

- Performs local authorization.

- Serves the resource or denies access.

Using the POST Profile Scenario

The SAML server generally supports the POST profile scenario, as follows:

1. The end user accesses the source web site, hereafter known as the source site.

2. The source site verifies whether or not the user has a current session.

3. If not, the source site prompts the user to enter user credentials.

4. The user supplies credentials, for example, username and password.

5. If the authentication is successful, the source site authentication server creates a session for the user and displays the appropriate welcome page of the portal application.

6. The user then selects a menu option or link that points to a resource or application on a destination web site.

7. The portal application directs the request to the local intersite transfer service, which can be hosted on the source site. The request contains the URL of the resource on the destination site, in other words, the TARGET URL.

8. The intersite transfer service sends an HTML form back to the browser. The HTML FORM contains a SAML response, within which is a SAML assertion. The response must be digitally signed. Typically the HTML FORM will contain an input or submit action that will result in an HTTP POST. This can be a user-clickable Submit button or a script that initiates the HTTP POST programmatically.

9. The browser, either due to a user action or by way of an auto-submit action, sends an HTTP POST containing the SAML response to the destination web site's assertion consumer service.

10. The replying party's assertion consumer (in this case, on the destination web site) validates the digital signature on the SAML response.

11. If valid, the assertion consumer sends a redirect to the browser, causing the browser to access the TARGET resource.

12. The system, on the destination site, verifies that the user is authorized to access the destination site and the TARGET resource.

13. If the user is authorized to access the destination site and the TARGET resource, the system returns the TARGET resource to the browser.

The main tasks you are required to fulfill to support the system as the relying party with the POST profile include:

- Implement the assertion consumer service, which receives and processes the POST form

- Integrate the assertion consumer service with the existing process, which:

  - Maps the SAML assertion to a local user.

  - Creates a user session.

  - Performs local authorization.

  - Serves the resource or denies access.

Related Documentation

- Understanding SAML 1.1 Assertions

Understanding SAML 1.1 Assertions

Each party in the request-response communication must adhere to certain requirements. The requirements provide a predictable infrastructure so that the assertions and artifacts can be processed correctly.

- The artifact is a Base64-encoded string of 40 bytes. An artifact acts as a token that references an assertion on the source site, so the artifact holder—the Connect Secure device—can authenticate a user who has signed in to the source site and who now wants to access a resource protected by the system. The source site sends the artifact to the device in a redirect, after the user attempts to access a resource protected by the system. The artifact contains:

  - TypeCode—A 2-byte hex code of 0x0001 that identifies the artifact type.

  - SourceID—A Base64-encoded string of 20 bytes that determines the source site identity and location. You can use OpenSSL or similar Base64 encoding tool to generate the encoded string. The system maintains a table of SourceID values and the URL for the corresponding SAML responder. The system and the source site communicate this information in a back channel. On receiving the SAML artifact, the system decodes it and ensures that it is 20 bytes. It determines whether or not the SourceID belongs to a known source site, and, if it does, obtains the site location before sending a SAML request. The source site generates the SourceID by computing the SHA-1 hash of the source site's own URL.

  - AssertionHandle—A 20-byte random value that identifies an assertion stored or generated by the source site. At least 8 bytes of this value should be obtained from a cryptographically secure RNG or PRNG.

- The intersite transfer service is the identity provider URL on the source site (not the Connect Secure device). Your specification of this URL in the admin console enables the system to construct an authentication request to the source site, which holds the user's credentials in cache. The request is similar to the following example:

GET http://<intersite transfer hostname and path>?TARGET=<Target>...<HTTP-Version><other HTTP 1.0 or 1.1 components>

In the preceding sample, <intersite transfer hostname and path> consists of the hostname, port number, and path components of the intersite transfer URL at the source and where Target=<Target> specifies the requested target resource at the destination (Connect Secure protected) site. This request might look like:

GET http://10.56.1.123:8002/xferSvc?TARGET=http://www.dest.com/sales.htm

- The intersite transfer service redirects the user's browser to the assertion consumer service at the destination site—in this case, the Connect Secure device. The HTTP response from the source site intersite transfer service must be in the following format:

<HTTP-Version> 302 <Reason Phrase>
<other headers>
Location: http://<assertion consumer hostname and path>?<SAML searchpart><other HTTP 1.0 or 1.1 components>

In the preceding sample, <assertion consumer hostname and path> provides the hostname, port number, and path components of an assertion consumer URL at the destination site and where <SAML searchpart>= …TARGET=<Target> …SAMLart=<SAML artifact>… consists of one target description, which must be included in the <SAML searchpart> component. At least one SAML artifact must be included in the SAML <SAML searchpart> component. The asserting party can include multiple SAML artifacts.

NOTE: You can use status code 302 to indicate that the requested resource resides temporarily under a different URI.

If <SAML searchpart> contains more than one artifact, all of the artifacts must share the same SourceID.

The redirect might look like:

HTTP/1.1 302 Found
Location: http://www.ive.com:5802/artifact?TARGET=/www.ive.com/&SAMLart=artifact

- The user's browser accesses the assertion consumer service, with a SAML artifact representing the user's authentication information attached to the URL.

The HTTP request must appear as follows:

GET http://<assertion consumer hostname and path>?<SAML searchpart> <HTTP-Version><other HTTP 1.0 or 1.1 request components>

In the preceding sample, <assertion consumer hostname and path> provides the hostname, port number, and path components of an assertion consumer URL at the destination site.

<SAML searchpart>= …TARGET=<Taret>…SAMLart=<SAML artifact> …

A single target description MUST be included in the <SAML searchpart> component. At least one SAML artifact MUST be included in the <SAML searchpart> component; multiple SAML artifacts MAY be included. If more than one artifact is carried within <SAML searchpart>, all the artifacts MUST have the same SourceID.

You should not expose the assertion consumer URL unless over SSL 3.0 or TLS 1.0. Otherwise, transmitted artifacts might be available in plain text to an attacker.

- The issuer value is typically the URL of the source site. You can specify the <ISSUER> variable, which will return the issuer value from the assertion.

- The username template is a reference to the SAML name identifier element, which allows the asserting party to provide a format for the username. The SAML specification allows for values in the following formats:

  - Unspecified—Indicates that interpretation of the content is left up to the individual implementations. In this case, you can use the variable assertionName.

  - E-mail Address—Indicates that the content is in the form of an e-mail address. In this case, you can use the variable assertionName.

  - X.509 Subject Name—Indicates that the content is in the form of an X.509 subject name. In this case, you can use the variable assertionNameDN.<RDN>.

  - Windows Domain Qualified Name—Indicates that the content is a string in the form of DomainName\Username.

  You should define the username template to accept the type of username your SAML assertion contains.

- You can prevent eavesdropping on the SAML artifact by synchronizing the clocks on the source and destination sites. The system provides an Allowed Clock Skew attribute that dictates the maximum time difference allowed between the system and the source site. The system rejects any assertions whose timing exceeds the allowed clock skew.

Related Documentation

- Understanding SAML 1.1 Profiles

- Creating a SAML 1.1 Server Instance

Creating a Trust Relationship Between SAML 1.1 Systems

In order to ensure that SAML-enabled systems are only passing information between trusted sources, you must create a trust relationship between the applications that are sending and receiving information.

Configuring Trusted Application URLs

In a trust relationship, you must provide the SAML-enabled systems with the URLs they need to contact each other. In some transactions, only the system that initiates the transaction (the Connect Secure device) needs to know the URL of the other system. (The system uses the URL to initiate the transaction.) In other transactions (SSO transactions using artifact profiles), you need to configure each system with the URL of the other.

The following list shows the different transaction types and the URLs you must configure for each:

- SSO transactions: Artifact profile—On Connect Secure, you must enter the URL of the assertion consumer service. For example, use **https://hostname/acs**.

  You must also enter the following URL for the system on the assertion consumer service. For example, use **https://<SecureAccessHostname>/dana-ws/saml.ws**.

- SSO transactions: POST profile—On Connect Secure, you must enter the URL of the assertion consumer service. For example, use **https://hostname/acs**.

- Access control transactions—On Connect Secure, you must enter the URL of the SAML Web service. For example, use **https://hostname/ws**.

Configuring an Issuer

Before accepting a statement from another system, a SAML-enabled entity must trust the issuer of the statement. You can control which issuers a system trusts by specifying the unique strings of the trusted issuers during the system's configuration. (When sending a statement, an issuer identifies itself by including its unique string in the statement. SAML-enabled applications generally use hostnames to identify issuers, but the SAML standard allows applications to use any string.) If you do not configure a system to recognize an issuer's unique string, the system will not accept that issuer's statements.

The following list shows the different transaction types and the issuers you must configure for each:

- SSO transactions—You must specify a unique string on the system (typically its hostname) that it can use to identify itself and then configure the access management system to recognize that string.

- Access control transactions—You must specify a unique string on the access management system (typically its hostname) that it can use to identify itself and then configure the system to recognize that string.

Configuring Certificates

Within SSL transactions, the server must present a certificate to the client, and then the client must verify (at minimum) that it trusts the certificate authority who issued the server's certificate before accepting the information. You can configure all of the system SAML transactions to use SSL (HTTPS).

Configuring SSO Transactions: Artifact Profile

Artifact profile transactions involve numerous communications back and forth between the system and the access management system. The methods you use to pass data and authenticate the two systems affect which certificates you must install and configure.

The following list shows the different artifact profile configuration options that require special certificate configurations:

- All artifact profile transactions—Regardless of your artifact profile configuration, you must install the certificate of the CA that signed the system Web server certificate on the access management system. (The system requires the access management system to use an SSL connection when requesting an authentication statement. In an SSL connection, the initiator must trust the system to which it is connecting. By installing the CA certificate on the access management system, you ensure that the access management system will trust the CA that issued the system certificate.)

- Sending artifacts over an SSL connection (HTTPS GET requests)—If you choose to send artifacts to the access management system using an SSL connection, you must install the access management system's root CA certificate on the system. (In an SSL connection, the initiator must trust the system to which it is connecting. By installing the access management system's CA certificate on the system, you ensure that the system will trust the CA that issued the access management system's certificate.) You can install the root CA from the System > Configuration > Certificates > Trusted Client CAs page in the admin console. If you do not want to send artifacts over an SSL connection, you do not need to install any additional certificates.

  To enable SSL-based communications from the system to the access management system, enter a URL that begins with HTTPS in the SAML Assertion Consumer Service URL field during the system configuration. You may also need to enable SSL on the access management system.

- Transactions using certificate authentication—If you choose to authenticate the access management system using a certificate, you must:

  - Install the access management system's root CA certificate on the system. You can install the root CA from the System > Configuration > Certificates > Trusted Client CAs page in the admin console.

  - Specify which certificate values the system should use to validate the access management system. You must use values that match the values contained in the access management server's certificate.

If you do not choose to authenticate the access management system, or if you choose to use username/password authentication, you do not need to install any additional certificates.

Configuring SSO Transactions: POST Profile

In a POST profile transaction, the system sends signed authentication statements to the access management system. Generally, it sends them over an SSL connection (recommended), but in some configurations, the system may send statements through a standard HTTP connection.

The following list shows the different POST profile configuration options that require special certificate configurations:

- All POST profile transactions—Regardless of your POST profile configuration, you must specify which certificate the system should use to sign its statements. You can choose a certificate in the Users > Resource Policies > Web > SSO SAML > [Policy] > General page in the admin console. Then, you must install the device certificate on the access management system. You can download the certificate from the System > Configuration > Certificates > Device Certificates > [Certificate] > Certificate Details page.

- Sending POST data over an SSL connection (HTTPS)—If you choose to send statements to the access management system using an SSL connection, you must install the access management system's root CA certificate on the system. (In an SSL connection, the initiator must trust the system to which it is connecting. By installing the access management

system's certificate on the system, you ensure that the system will trust the CA that issued the access management system's certificate.) You can install the root CA from the System > Configuration > Certificates > Trusted Client CAs page in the admin console. If you do not want to post statements over an SSL connection, you do not need to install any additional certificates.

To enable SSL-based communications from the system to the access management system, enter a URL that begins with HTTPS in the SAML assertion consumer service URL field during the system configuration. You may also need to enable SSL on the access management system.

Configuring Access Control Transactions

In an access control transaction, the system posts an authorization decision query to the access management system. To ensure that the access management system responds to the query, you must determine which certificate options are required by your configuration.

The following list shows the different access control configuration options that require special certificate configurations:

- Sending authorization data over an SSL connection—If you choose to connect to the access management system using an SSL connection, you must install the access management system's root CA on the system. (In an SSL connection, the initiator must trust the system to which it is connecting. By installing the access management system's certificate on the system, you ensure that the system will trust the CA that issued the access management system's certificate.) You can install the root CA from the System > Configuration > Certificates > Trusted Client CAs page in the admin console.

- Transactions using certificate authentication—If you choose to use certificate authentication, you must configure the access management system to trust the CA that issued the certificate. Optionally, you may also choose to accept the certificate based on the following additional options:

  - Upload the certificate public key to the access management system.

  - Validate the system using specific certificate attributes.

  These options require that you specify which certificate the system should pass to the access management system. You can choose a certificate in the Users > Resource Policies > Web > SAML ACL > [Policy] > General page in the admin console.

  To determine how to configure your access management system to validate the certificate, see your access management system's documentation. If your access management system does not require certificate authentication, or if it uses username/password authentication, you do not need to configure the system to pass the access management server a certificate. If you do not specify a trust method, your access management system may accept authorization requests from any system.

Configuring User Identity

In a trust relationship, the two entities must agree on a way to identify users. You may choose to share a username across systems, select an LDAP or certificate user attribute to share across systems, or hardcode a user ID. (For example, you may choose to set the Subject Name field to "guest" to easily allow access across systems.)

To ensure that the two systems are passing common information about users, you must specify which information the system should pass using options in the User Identity section of the Users > Resource Policies > Web > SAML SSO > [Policy] > General page and the Users > Resource Policies > Web > SAML ACL > [Policy] > General page. Choose a username or attribute that the access management system will recognize.

Related Documentation

- Using Trusted Client CAs

- Understanding SAML 1.1

- Configuring SAML 1.1 SSO Profiles

- Configuring General Role Options

SAML Version 1.1 Configuration Tasks

The following topics describe how to configure the features that support SAML version 1.1:

- Creating a SAML 1.1 Server Instance
- Configuring SAML 1.1 SSO Profiles
- Creating a SAML 1.1 SSO POST Profile
- Creating a SAML 1.1 ACL Resource Policy

Creating a SAML 1.1 Server Instance

To create a new SAML server instance:

1. In the admin console, choose **Authentication > Auth. Servers**.

2. Select **SAML Server** from the New list, and then click **New Server**.

3. Complete the settings as described in Table 46.

4. Click **Save Changes**.

After you save changes for the first time, the page is redisplayed and now has two tabs. The Settings tab allows you to modify any of the settings pertaining to the SAML Server instance. The Users tab lists valid users of the server.

Table46: SAML Authentication Server (SAML 1.1)

Setting	Guideline
Name	Specify a name to identify the server instance.
Settings	
SAML Version	Select **1.1**.
Source Site Inter-Site Transfer Service URL	User is redirected to this URL in destination first scenario.
Issuer Value for Source Site	Typically the URI or hostname of the issuer of the assertion.
User Name Template	Enter the mapping string from the SAML assertion to a user realm. For example, enter **<assertionNameDN.CN>** to derive the username from the CN value in the assertion.

Allowed Clock Skew	The maximum allowed difference in time between the system clock and the source site clock.
	NOTE: SAML is a time sensitive protocol. The time-based validity of a SAML assertion is determined by the SAML identity provider. If the SAML identity provider and SAML service provider clocks are askew, the assertion can be determined invalid, and you will receive the following error:
	"SAML Transferred failed. Please contact your system administrator. Detail: Failure: No valid assertion found in SAML response."
	We recommend you use NTP to ensure the clocks are synchronized and that you set an Allowed Clock Skew value that accommodates any expected or permissible skew.
SSO Method	
Artifact	• Source ID. A Base64-encoded string of 20 bytes that the system uses to recognize an assertion from a given source site.
	• Source SOAP Responder Service URL
	• SOAP Client Authentication. Select **HTTP Basic** or **SSL Client Certificate** and complete the related settings.
	NOTE: SOAP requests generated by the system (when configured as a SAML 1.1 consumer) are not signed.
POST	• Response Signing Certificate. Enter the name of, or browse to locate, the PEM-formatted signing certificate, which is loaded for the SAML response signature verification.
	The certificate you select should be the same certificate used for signing the SAML response at the source site. The source site may send this certificate along with the SAML response, depending on the source site configuration. By default, the system performs signature verification of the SAML response first on the locally configured certificate. If a certificate is not configured locally in the SAML authentication server, then the system performs the signature verification on the certificate included in the SAML response from the source site.
	• **Enable Signing Certificate status checking**. Select this option to check the validity of the signing certificate configured in the SAML authentication server POST profile. It is possible that the certificate has already expired or has been revoked.
User Record Synchronization	
Enable User Record Synchronization	Allow users to retain their bookmarks and individual preferences regardless of which device they log in to.
Logical Auth Server Name	Logical name of the authentication server.

Related Documentation

Configuring SAML 1.1 SSO Profiles

When enabling SSO transactions to a trusted access management system, you must indicate whether the access management system should "pull" user information from Connect Secure or whether Connect Secure should "push" it to the access management system. You indicate which communication method the two systems should use by selecting a profile during configuration. A *profile* is a method that two trusted sites use to transfer a SAML statement. When configuring the system, you may choose to use an artifact or POST profile.

When you choose to communicate using the *artifact profile* (also called browser/artifact profile), the trusted access management server "pulls" authentication information from the system.

Figure 46 shows the SAML communication process when the implementation uses the artifact profile.

Figure46: Artifact Profile
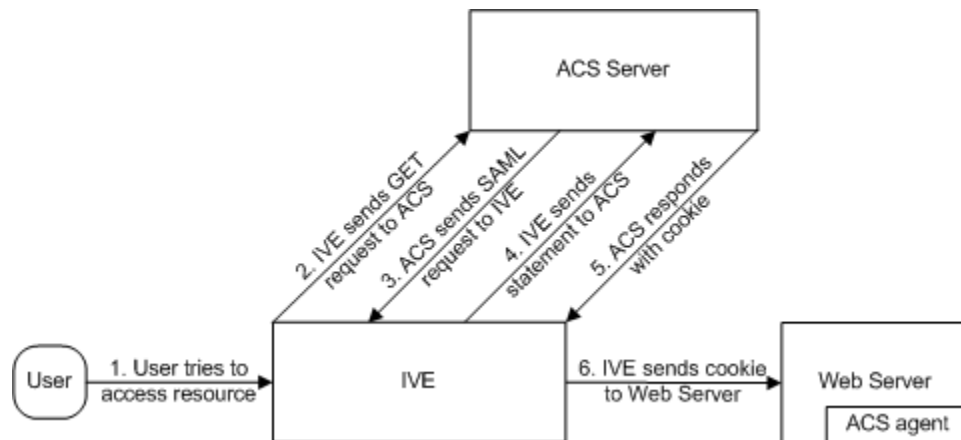
The system and an assertion consumer service (ACS) use the following process to pass information:

1. The user tries to access a resource—A user is signed into the device and tries to access a protected resource on a Web server.

2. The system sends an HTTP or HTTPS GET request to the ACS—the system intercepts the request and checks whether it has already performed the necessary SSO operation to honor the request. If not, the system creates an authentication statement and passes an HTTP query variable called an artifact to the assertion consumer service.

   An artifact profile is a Base64-encoded string that contains the source ID of the source site (that is, a 20-byte string that references the system) and a randomly generated string that acts as a handle to the authentication statement. (Note that a handle expires 5 minutes after the artifact is sent, so if the assertion consumer service responds after 5 minutes, the system does not send a statement. Also note that the system discards a handle after its first use to prevent the handle from being used twice.)

3. The ACS sends a SAML request to the system—The assertion consumer service uses the source ID sent in the previous step to determine the location of the device. Then the assertion consumer service sends a statement request wrapped in a SOAP message to the following address on the system:

   https://<<ivehostname>/danaws/saml.ws

   The request includes the statement handle passed in the previous step.

   > NOTE: The system only supports type 0x0001 artifacts. This type of artifact passes a reference to the source site's location (that is, the source ID of the system), rather than sending the location itself. To handle type 0x0001 artifacts, the assertion consumer service must maintain a table that maps source IDs to the locations of partner source sites.

4. The system sends an authentication statement to the ACS—the system uses the statement handle in the request to find the correct statement in the system cache and then sends the appropriate authentication statement back to the assertion consumer service. The unsigned statement contains the user's identity and the mechanism he used to sign into the device.

5. The ACS sends a cookie to the system—The assertion consumer service accepts the statement and then it sends a cookie back to the system that enables the user's session.

6. The system sends the cookie to the Web server—the system caches the cookie to handle future requests. Then the system sends the cookie in an HTTP request to the Web server

whose domain name matches the domain in the cookie. The Web server honors the session without prompting the user for credentials.

i NOTE: If you configure the system to use artifact profiles, you must install the Web server certificate on the assertion consumer service.

To write a SAML SSO artifact profile resource policy:

1. In the admin console, select **Users > Resource Policies > Web**.

2. If your administrator view is not already configured to show SAML policies, make the following modifications:

   1. Click the **Customize** button in the upper right corner of the page.

   2. Select the **SSO** check box.

   3. Select the **SAML** check box below the SSO check box.

   4. Click **OK**.

5. Use the tabs to display the SSO > SAML page.

6. Click **New Policy**.

7. On the New Policy page, enter:

   1. A name to label this policy.

   2. A description of the policy (optional).

3. In the Resources section, specify the resources to which this policy applies.

4. In the Roles section, specify:

   - **Policy applies to ALL roles**—To apply this policy to all users.

   - **Policy applies to SELECTED roles**—To apply this policy only to users who are mapped to roles in the Selected roles list. Make sure to add roles to this list from the Available roles list.

   - **Policy applies to all roles OTHER THAN those selected below**—To apply this policy to all users except for those who map to the roles in the Selected roles list. Make sure to add roles to this list from the Available roles list.

5. In the Action section, specify:

   - **Use the SAML SSO defined below**—The system performs a single-sign on (SSO) request to the specified URL using the data specified in the SAML SSO details section. The system makes the SSO request when a user tries to access a SAML resource specified in the Resources list.

   - **Do NOT use SAML**—The system does not perform a SSO request.

   - **Use Detailed Rules**—To specify one or more detailed rules for this policy.

6. In the SAML SSO Details section, specify:

   - **SAML Assertion Consumer Service URL**—Enter the URL that the system should use to contact the assertion consumer service (that is, the access management server). For example, https://<hostname>:<port>/dana-na/auth/saml-consumer.cgi. (Note that the system also uses this field to determine the SAML recipient for its assertions.)

   i NOTE: If you enter a URL that begins with HTTPS, you must install the assertion consumer service's root CA on the system.

- **Profile**—Select **Artifact** to indicate that the assertion consumer service should "pull" information from the system during SSO transactions.

- **Source ID**—Enter the source ID for the system. It must be a Base64-encoded string. The system decodes it and ensures that it is 20 bytes. You can use OpenSSL or other Base64 tool to generate the Base64-encoded string.

  > *i* NOTE: The system identifier (that is, the source ID) must map to the following URL on the assertion consumer service: https://<ivehostname>/dana-ws/saml.ws

- **Issuer**—Enter a unique string that the system can use to identify itself when it generates assertions (typically its hostname).

  > *i* NOTE: You must configure the assertion consumer service to recognize the unique string.

7. In the User Identity section, specify how the system and the assertion consumer service should identify the user:

   - **Subject Name Type**—Specify which method the system and assertion consumer service should use to identify the user:

     - **DN**—Send the username in the format of a DN (distinguished name) attribute.

     - **Email Address**—Send the username in the format of an e-mail address.

     - **Windows**—Send the username in the format of a Windows domain qualified username.

     - **Other**—Send the username in another format agreed upon by the system and the assertion consumer service.

   - **Subject Name**—Use variables to specify the username that the system should pass to the assertion consumer service. Or, enter static text.

     > *i* NOTE: You must send a username or attribute that the assertion consumer service will recognize.

8. In the Web Service Authentication section, specify the authentication method that the system should use to authenticate the assertion consumer service:

   - **None**—Do not authenticate the assertion consumer service.

   - **Username**—Authenticate the assertion consumer service using a username and password. Enter the username and password that the assertion consumer service must send.

   - **Certificate Attribute**—Authenticate the assertion consumer service using certificate attributes. Enter the attributes that the assertion consumer service must send (one attribute per line). For example, use cn=sales. You must use values that match the values contained in the assertion consumer service certificate.

     > *i* NOTE: If you select this option, you must install the assertion consumer service root CA on the system.

9. **Cookie Domain**—Enter a comma-separated list of domains to which we send the SSO cookie.

10. Click **Save Changes**.

11. On the SAML SSO Policies page, order the policies according to how you want the system to evaluate them. Keep in mind that once the system matches the resource

requested by the user to a resource in a policy's (or a detailed rule's) Resource list, it performs the specified action and stops processing policies.

Related Documentation

- Understanding Digital Certificate Security

Creating a SAML 1.1 SSO POST Profile

When you choose to communicate using a POST profile (also called browser/POST profile), the system "pushes" authentication data to the access management system using an HTTP POST command over an SSL 3.0 connection.

Figure 47 shows the SAML communication process when the implementation uses the POST profile.

Figure47: POST Profile
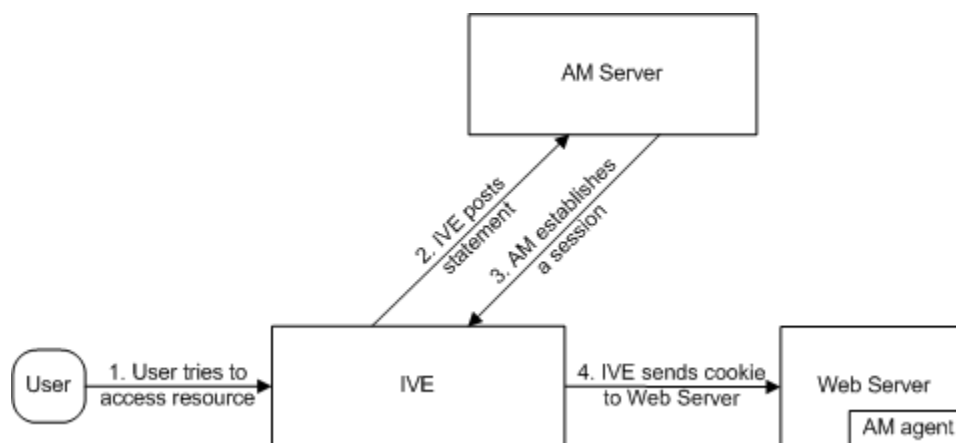
The system and an access management system use the following process to pass information:

1. The user tries to access a resource—A user is signed into the device and tries to access a protected resource on a Web server.

2. The system posts a statement—the system intercepts the request and checks whether it has already performed the necessary SSO operation to honor the request. If not, the system creates an authentication statement, digitally signs it, and posts it directly to the access management server. Since the statement is signed, the access management server must trust the certificate authority that was used to issue the certificate. Note that you must configure which certificate the system uses to sign the statement.

3. The AM establishes a session—If the user has the proper permissions, the access management server sends a cookie back to the system that enables the user's session.

4. The system sends the cookie to the Web server—the system caches the cookie to handle future requests. Then the system sends the cookie in an HTTP request to the Web server whose domain name matches the domain in the cookie. The Web server honors the session without prompting the user for credentials.

   > 🛈 NOTE: If you configure the system to use POST profiles, you must install the assertion consumer service's root CA on the system and determine which method the assertion consumer service uses to trust the certificate.

To write a SAML SSO POST profile resource policy:

1. In the admin console, select **Users > Resource Policies > Web**.

2. If your administrator view is not already configured to show SAML policies, make the following modifications:

    1. Click the **Customize** button in the upper right corner of the page.

    2. Select the **SSO** check box.

    3. Select the **SAML** check box below the SSO check box.

    4. Click **OK**.

5. Use the tabs to display the SSO > SAML page.

6. Click **New Policy**.

7. On the SAML SSO Policy page, enter:

- A name to label this policy.

- A description of the policy (optional).

8. In the Resources section, specify the resources to which this policy applies.

9. In the Roles section, specify:

- **Policy applies to ALL roles**—To apply this policy to all users.

- **Policy applies to SELECTED roles**—To apply this policy only to users who are mapped to roles in the Selected roles list. Make sure to add roles to this list from the Available roles list.

- **Policy applies to all roles OTHER THAN those selected below**—To apply this policy to all users except for those who map to the roles in the Selected roles list. Make sure to add roles to this list from the Available roles list.

10.   In the Action section, specify:

- **Use the SAML SSO defined below**—The system performs a single-sign on (SSO) request to the specified URL using the data specified in the SAML SSO details section. The system makes the SSO request when a user tries to access a SAML resource specified in the Resources list.

- **Do NOT use SAML**—The system does not perform an SSO request.

- **Use Detailed Rules**—To specify one or more detailed rules for this policy.

11.   In the SAML SSO Details section, specify:

- **SAML Assertion Consumer Service URL**—Enter the URL that the system should use to contact the assertion consumer service (that is, the access management server). For example, use **https://hostname/acs**.

- **Profile**—Select POST to indicate that the system should "push" information to the assertion consumer service during SSO transactions.

- **Issuer**—Enter a unique string that the system can use to identify itself when it generates assertions. Typically, the issuer string is a hostname.

  > *i*   NOTE: You must configure the assertion consumer service to recognize the unique string.

- **Signing Certificate**—Specify which certificate the system should use to sign its assertions.

12.   In the User Identity section, specify how the system and the assertion consumer service should identify the user:

- **Subject Name Type**—Specify which method the system and assertion consumer service should use to identify the user:

  - **DN**DN—Send the username in the format of a DN (distinguished name) attribute.

  - **Email Address**—Send the username in the format of an e-mail address.

  - **Windows**—Send the username in the format of a Windows domain qualified username.

  - **Other**—Send the username in another format agreed upon by the system and the assertion consumer service.

- **Subject Name**—Use variables to specify the username that the system should pass to the assertion consumer service. Or, enter static text.

  > *i*   NOTE: You must send a username or attribute that the assertion consumer service will recognize.

- **Cookie Domain**—Enter a comma-separated list of domains to which we send the SSO cookie.

13.   Click **Save Changes**.

14.   On the SAML SSO Policies page, order the policies according to how you want the system to evaluate them. Keep in mind that once the system matches the resource requested by the user to a resource in a policy's (or a detailed rule's) Resource list, it performs the specified action and stops processing policies.

Related Documentation

- Understanding Digital Certificate Security

- Writing a Web Proxy Resource Policy

Creating a SAML 1.1 ACL Resource Policy

When enabling access control transactions to a trusted access management system, the system and trusted access management system exchange information using the method shown in Figure 48.

Figure48: Access Control Policies
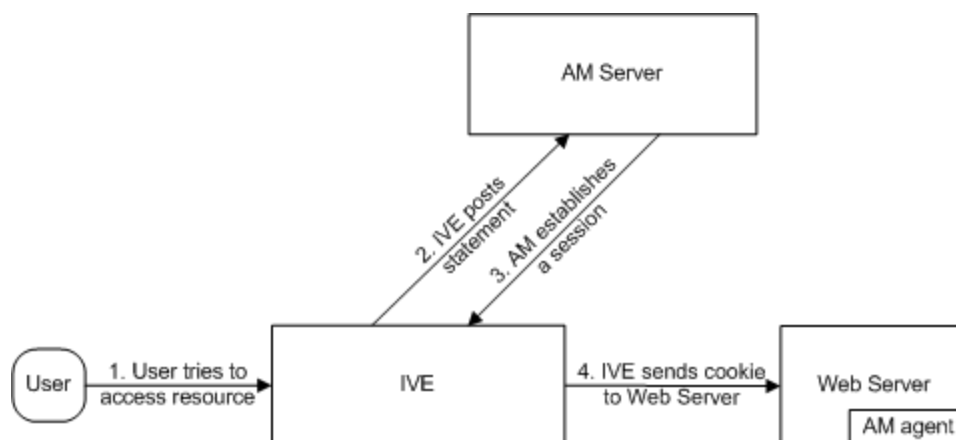
The system and an access management system use the following process to pass information:

1. The user tries to access a resource—A user is signed into the system and tries to access a protected resource on a Web server.

2. The system posts an authorization decision query—If the system has already made an authorization request and it is still valid, the system uses that request. (The authorization

request is valid for the period of time specified in the admin console.) If it does not have a valid authorization request, the system posts an authorization decision query to the access management system. The query contains the user's identity and the resource that the access management system needs to authorize.

3. The access management system posts an authorization decision statement—The access management system sends an HTTPS POST containing a SOAP message that contains the authorization decision statement. The authorization decision statement contains a result of permit, deny, or indeterminate.

4. The system sends the request to the Web browser—If the authorization decision statement returns a result of permit, the system allows the user access. If not, the system presents an error page to the user telling him that he does not have the proper access permissions.

> **NOTE:** If you configure the system to use access control transactions, you must install the SAML Web service root CA on the system.

To create a SAML access control resource policy:

1. In the admin console, select **Users > Resource Policies > Web**.

2. If your administrator view is not already configured to show SAML policies, make the following modifications:

   1. Click the **Customize** button in the upper right corner of the page.

   2. Select the **SAML ACL** check box below the Access check box.

   3. Click **OK**.

3. Use the tabs to display the Access > SAML ACL page.

4. On the SAML Access Control Policies page, click **New Policy**.

5. On the New Policy page, enter:

   1. A name to label this policy.

   2. A description of the policy (optional).

3. In the Resources section, specify the resources to which this policy applies.

4. In the Roles section, specify:

   - **Policy applies to ALL roles**—To apply this policy to all users.

   - **Policy applies to SELECTED roles**—To apply this policy only to users who are mapped to roles in the Selected roles list. Make sure to add roles to this list from the Available roles list.

   - **Policy applies to all roles OTHER THAN those selected below**—To apply this policy to all users except for those who map to the roles in the Selected roles list. Make sure to add roles to this list from the Available roles list.

5. In the Action section, specify:

   - **Use the SAML Access Control checks defined below**—The system performs an access control check to the specified URL using the data specified in the SAML Access Control Details section.

   - **Do not use SAML Access**—The system does not perform an access control check.

   - **Use Detailed Rules**—To specify one or more detailed rules for this policy.

6. In the SAML Access Control Details section, specify:

- **SAML Web Service URL**—Enter the URL of the access management system's SAML server. For example, use **https://hostname/ws**.

- **Issuer**—Enter the hostname of the issuer, which in most cases is the hostname of the access management system.

  i NOTE: You must enter a unique string that the SAML Web service uses to identify itself in authorization assertions.

7. In the User Identity section, specify how the system and the SAML Web service should identify the user:

   - **Subject Name Type**—Specify which method the system and SAML Web service should use to identify the user:

     - **DN**—Send the username in the format of a DN (distinguished name) attribute.

     - **Email Address**—Send the username in the format of an e-mail address.

     - **Windows**—Send the username in the format of a Windows domain qualified username.

     - **Other**—Send the username in another format agreed upon by the system and the SAML Web service.

   - **Subject Name**—Use variables to specify the username that the system should pass to the SAML Web service. Or, enter static text.

     i NOTE: You must send a username or attribute that the SAML Web service will recognize.

8. In the Web Service Authentication section, specify the authentication method that the SAML Web service should use to authenticate the system:

   - **None**—Do not authenticate the system.

   - **Username**—Authenticate the system using a username and password. Enter the username and password that the system must send the Web service.

   - **Certificate Attribute**—Authenticate the system using a certificate signed by a trusted certificate authority. If you have more than one certificate installed on the system, use the drop-down list to select which certificate to send to the Web service.

     i NOTE: If you select this option, you must install the Web server certificate on the access management system Web server and determine which method the SAML Web service uses to trust the certificate.

9. In the Options section, specify:

   - **Maximum Cache Time**—You can eliminate the overhead of generating an authorization decision each time the user requests the same URL by indicating that the system must cache the access management system's authorization responses. Enter the amount of time the system should cache the responses (in seconds).

   - **Ignore Query Data**—By default, when a user requests a resource, the system sends the entire URL for that resource (including the query parameter) to the SAML Web service and caches the URL. You can specify that the system should remove the query string from the URL before requesting authorization or caching the authorization response.

10. Click **Save Changes**.

11. On the SAML Access Control Policies page, order the policies according to how you want the system to evaluate them. Keep in mind that once the system matches the resource

requested by the user to a resource in a policy's (or a detailed rule's) Resource list, it performs the specified action and stops processing policies.

Related Documentation

- Understanding Digital Certificate Security

- Writing a Web Proxy Resource Policy

CHAPTER 10 Device Access Management Framework

Understanding the Device Access Management Framework

The *device access management framework* leverages mobile device management (MDM) services so that you can use familiar Pulse Secure client policies to enforce security objectives based on your device classification scheme: device status is MDM enrollment complete or incomplete; device status is MDM-policy compliant or noncompliant; device is employee owned or company owned; device platform is iOS, Android, or neither; and so forth.

In this framework, the MDM is a device authorization server, and MDM record attributes are the basis for granular access policy determinations. For example, you can implement policies that allow devices that have a clean MDM posture assessment and are compliant with MDM policies to access the network, but deny access to servers when you want to prevent downloads to employee-owned devices or to a particular platform that might be vulnerable. To do this, you use the device attributes and status maintained by the MDM in Pulse Secure client role-mapping rules, and specify the device-attribute-based roles in familiar Pulse Secure client policies.

The framework simply extends the user access management framework realm configuration to include use of device attributes as a factor in role mapping rules. Figure 49 illustrates the similarities.

Figure49: User Access Management Framework and Device Access Management Framework
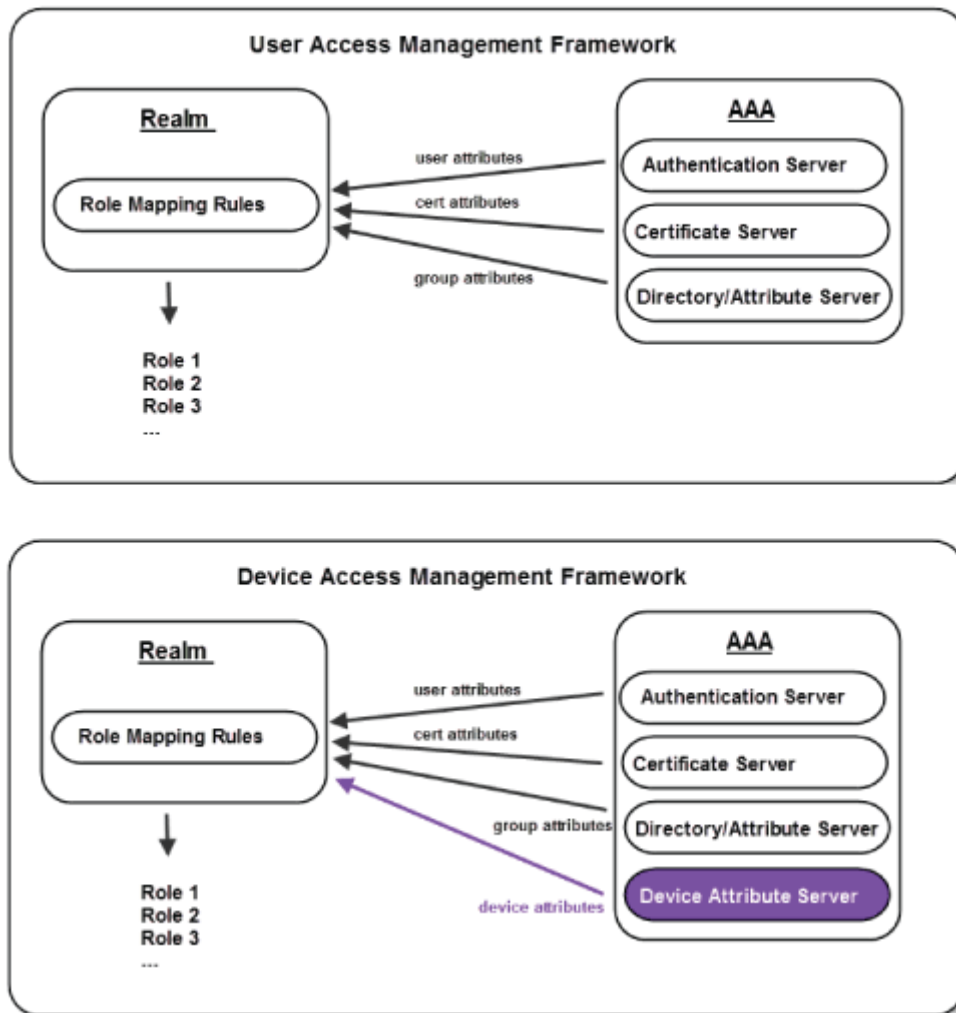
Table 47 summarizes vendor support for this release.

Table47: MDM Vendors

Product	Vendor
Policy Secure	• AirWatch MDM • MobileIron MDM In addition, you can integrate with Pulse Secure Endpoint Profiler as a device attribute server. The Endpoint Profiler catalogues mobile device platform information.
Connect Secure	• AirWatch MDM • MobileIron MDM

Table 48 summarizes supported methods for determining the device identifiers.

Table48: Device Identifiers

Product	Policies
Policy Secure	• Device certificate (recommended) • MAC address from callingStationID RADIUS request attribute in 802.1x communication
Connect Secure	Device certificate (required)

Table 49 summarizes policy reevaluation features.

Table49: Policy Reevaluation

Product	Policy Reevaluation
Policy Secure	Policies are evaluated during sign in and when the MDM data changes. Policy Secure queries the MDM database for updates at an interval you specify. This feature enables you to leverage the MDM posture assessment checks and enforce compliance. For example, the MDM might detect that a device is out of compliance with its security policies, such as a password policy. At the next device check interval, Policy Secure queries the MDM for updated attribute data. In this example, it learns that a formerly compliant device is now noncompliant. It assigns the device the noncompliant role and sends the 802.1x authenticator the corresponding RADIUS attribute to place it in a remediation VLAN.
Connect Secure	The MDM is query and policies evaluated only during sign-in. If desired, you can use the user role session timeout setting to force users to sign in periodically. If you use a certificate server for user authentication, the users are not prompted to sign in again; however, if you have enabled user role notifications, users do receive a notification each time sign-in occurs.

🛈 NOTE: The dynamic policy evaluation feature is not used in the device access management framework.

Table 50 summarizes the policies in which you can specify device-attribute-based roles.

Table50: Policies

Product	Policies
Policy Secure	• 802.1x network access control RADIUS return attribute policies (Layer 2) • Infranet Enforcer resource policies (Layer 3)
Connect Secure	Resource policies or resource profiles

Related Documentation

- Deploying a BYOD Policy for AirWatch Managed Devices

- Deploying a BYOD Policy for MobileIron Managed Devices

- *Deploying a BYOD Policy for Devices Discovered by Pulse Secure Endpoint Profiler*

Deploying a BYOD Policy for AirWatch Managed Devices

This example shows how to use policies to enable security based on device identity, device posture, or user identity in a bring your own device (BYOD) environment for an enterprise that uses AirWatch® for mobile device management (MDM). It includes the following information:

- Solution Overview

- Requirements

- Configuring the AirWatch MDM Service

Solution Overview

In the past, to ensure security and manageability of the corporate network, enterprise information technology (IT) departments had restricted network access to company-issued equipment. For mobile phones, the classic example was the company-issued BlackBerry handset. As powerful mobile smart phones and tablets have become commonly held personal possessions, the trend in enterprise IT has been to stop issuing mobile equipment and instead allow employees to use their personal smart phones and tablets to conduct business activities. This has lowered equipment costs, but BYOD environments pose capacity planning and security challenges: how can an enterprise track network access by non-company-issued devices? Can an enterprise implement policies that can restrict the mobile devices that can access the network and protected resources in the same way that SSL VPN solutions restrict user access?

MDM vendors have emerged to address the first issue. MDMs such as AirWatch provide enrollment and posture assessment services that prompt employees to enter data about their mobile devices. The MDM data records include device attributes and posture assessment status that can be used in the access management framework to enforce security policies.

Figure 50 shows a deployment with Connect Secure and the AirWatch MDM cloud service.

Figure50: Solution Topology

The solution shown in this example leverages the Pulse Secure access management framework to support attribute-based network access control for mobile devices. In the *device access management framework*, the MDM is a device authorization server and MDM record attributes are the basis for access policy determinations. For example, suppose your enterprise wants to enforce a policy that allows access only to mobile devices that have enrolled with the MDM or are compliant with the MDM posture assessment policies. You can use the attributes and status maintained by the MDM in role-mapping rules to implement the policy.

In this framework, a native supplicant is used to authenticate the user of the device. The device itself is identified using a client certificate that contains device identity. The client certificate can be used to identify the device against the MDM records and authenticate the user against a certificate server.

The Pulse Secure solution supports granular, attribute-based resource access policies. For example, you can implement policies that allow devices that have a clean MDM posture assessment and are compliant with MDM policies to access the network, but deny access to servers when you want to prevent downloads to employee-owned devices or to a particular platform that might be vulnerable.

Requirements

Table 51 lists version information for the solution components shown in this example.

Table51: Component Version Information

Component	Version
Connect Secure	Release 8.0r1 or later is required.
AirWatch MDM	Release 6.4.1.2 is used in this example. Any version that supports the device ID and device attributes you plan to query is compatible.

Configuring the AirWatch MDM Service

This solution assumes you know how to configure and use the features of your MDM, and that you can enroll employees and their devices. For more information about the AirWatch MDM, refer to its documentation and support resources. This section focuses on the following elements of the MDM configuration that are important to this solution:

- Device identifier—The primary key for device records. Your MDM configuration determines whether a universal unique identifier (UUID), unique device identifier (UDID), or serial number is used as the device identifier. For AirWatch, UDID is supported and recommended.

- Device attributes—A standard set of data maintained for each device. For AirWatch, see Table 52.

When the user installs the MDM application on the device and completes enrollment, the MDM pushes the device certificate to the device. After enrollment, the MDM maintains a database record that includes information about the enrollee—attributes related to device identity, user identity, and posture assessment against MDM policies. Table 52 describes these attributes. In this solution, these attributes are used in the role mapping that is the basis for network access and resource access policies. When you configure role-mapping rules, you select the normalized Connect Secure attribute name.

Table52: AirWatch Device Attributes

AirWatch Attribute	Normalized Connect SecureName	Description	Data Type
BlockLevelEncryption	BlockLevelEncryption	True if block-level encryption is enabled; false otherwise.	Boolean
ComplianceStatus	complianceReason	Values: Compliant, Non-Compliant.	String
ComplianceStatus	isCompliant	True if the status is compliant with MDM policies; false otherwise.	Boolean
CompromisedStatus	CompromisedStatus	True if the status is compromised; false otherwise.	Boolean

CompromisedStatus	isCompromised	True if the device is compromised; false otherwise.	Boolean
DataProtectionEnabled	DataProtectionEnabled	True if data protection is enabled; false otherwise.	Boolean
DeviceFriendlyName	deviceName	The concatenated name used to identify the device/user combination.	String
EnrollmentStatus	isEnrolled	True if MDM value is Enrolled; false otherwise.	Boolean
FileLevelEncryption	FileLevelEncryption	True if file-level encryption is enabled; false otherwise.	Boolean
Id.Value	deviceId	Device identifier.	String
Imei	IMEI	IMEI number of the device.	String
IsPasscodeCompliant	IsPasscodeCompliant	True if the passcode is compliant with the MDM policy; false otherwise	Boolean
IsPasscodePresent	IsPasscodePresent	True if a passcode has been configured; false otherwise.	Boolean
LastComplianceCheckOn	LastComplianceCheckOn	The refresh date and timestamp of the last status reported.	Timestamp
LastCompromisedCheckOn	LastCompromisedCheckOn	The refresh date and timestamp of the last status reported.	Timestamp
LastSeen	lastSeen	Date and time the device last made successful contact with the MDM.	Timestamp
LocationGroupName	LocationGroupName	MDM location group configuration value.	String
MacAddress	macAdress	The Wi-Fi MAC address.	String
Model	model	Model is automatically reported by the device during registration.	String
OperatingSystem	osVersion	OS version.	String
Ownership	ownership	Values: C, E, or S (Corporate, Employee, or Shared).	String
PhoneNumber	phoneNumber	Phone number entered during registration.	String
Platform	platform	Platform specified during registration.	String
SerialNumber	serialNumber	Serial number.	String
Udid	UDID	Unique device identifier.	String
UserEmailAddress	userEmail	E-mail address of device user.	String
UserName	userName	Name of device user.	String
Uuid	UUID	Universal unique identifier.	String

To configure the MDM:

1. Enroll devices in the MDM using the methods supported by the MDM.

2. Create a profile. The profile determines many MDM management options. The following configurations are key to this solution:

    1. Certificate template. Create a configuration that specifies the field and type of identifier for client device certificates. See Figure 51.

        The MDM configuration templates provide flexibility in how the device identifier can be placed in the device certificate's subject or alternative subject. We recommend you include the user ID in the certificate, so the certificate can identify both the user and the device. For example:

CN=<EnrollmentUser>, serialNumber=<DeviceUid>, o=Company

2. Credential profile. Create a configuration that specifies the certificate authority and certificate template configuration. See Figure 52.

3. VPN profile. Create a configuration that specifies the system VPN, security options, and the credential configuration. See Figure 53.

4. Save and deploy the profile to devices registered with your organization. See Figure 54.

5. Enable API access and generate the AirWatch API key (tenant code). The tenant code is part of the REST API configuration. The tenant code must be included in the system MDM server configuration. It is sent in the API call. See Figure 55.

Figure51: AirWatch Certificate Template Configuration

Figure52: AirWatch Profile Credential Configuration

Figure53: AirWatch Profile VPN Configuration

Figure54: Deploying a Profile to Your Organization's Managed Devices

Figure55: AirWatch API Tenant Code Configuration

Configuring the Device Access Management Framework

This section describes the basic steps for configuring the device access management framework:

1. Configuring the MDM Authentication Server
2. Configuring the Certificate Server
3. Adding the MDM Certificate to the Trusted Client CA Configuration
4. Configuring User Roles
5. Configuring a Realm and Role Mapping Rules
6. Configuring a Sign-In Policy

Configuring the MDM Authentication Server

The MDM authentication server configuration is used by the system to communicate with the MDM. In the device access management framework, the MDM server is used as the device authorization server.

To configure the authentication server:

1. Select **Authentication > Auth Servers** to navigate to the authentication server configuration pages.
2. Select **MDM Server** and click **New Server** to display the configuration page shown in Figure 56.
3. Complete the configuration as described in Table 53.
4. Save the configuration.

Figure56: Authentication Server Configuration Page

Table53: Authentication Server Configuration Guidelines

Settings	Guidelines
Name	Specify a name for the configuration.
Type	Select **AirWatch**.
Server	
Server Url	Specify the URL for your AirWatch server. This is the URL AirWatch has instructed you to use to access its RESTful Web API (also called a RESTful Web service). The URL for the AirWatch MDM server used in this example has the following form: **https://apidev-as.Awmdm.com** NOTE: You must configure your firewalls to allow communication between these two nodes over port 443.
Viewer Url	Specify the URL for the AirWatch report viewer. This URL is used for links from the Active Users page to the AirWatch report viewer. The URL for the AirWatch MDM viewer for this example has the following form: **https://apidev.awmdm.com/AirWatch/Devices/DeviceDetails/<deviceAttr.deviceId>**
Request Timeout	Specify a timeout period (0-60 seconds) for queries to the MDM server. The default is 15 seconds.
Administrator	
Username	Specify the username for an account that has privileges to access the AirWatch RESTful Web API.
Password	Specify the corresponding password.
Tenant Code	Copy and paste the AirWatch API tenant code. See Figure 57.
Device Identifier	
ID Template	Construct a template to derive the device identifier from the certificate attributes. The template can contain textual characters as well as variables for substitution. The variables are the same as those used in role mapping custom expressions and policy conditions. Enclose variables in angle brackets like this <variable>. For example, suppose the certificate DN is: CN=<EnrollmentUser>, serialNumber=<DeviceUid>, o=Company. With this configuration, the certificate could identify both the user and the device. In this example, the device ID template is <certDN.serialNumber>.
ID Type	Select the device identifier type that matches the selection in the MDM certificate configuration: • **UUID**—Not applicable for the AirWatch MDM. • **Serial Number**—The device serial number. • **UDID**—The device unique device identifier. This is supported by the AirWatch MDM.

Configuring the Certificate Server

The certificate server configuration enables device users to authenticate using the certificate pushed to the device by the MDM. The certificates are used for user authentication, and the users do not have to enter user credentials.

To configure authentication with the certificate server:

1. Select **Authentication** > **Auth. Servers**.

2. Select **Certificate Server** and click **New Server** to display the configuration page shown in Figure 57.

3. Complete the configuration as described in Table 54.

4. Save the configuration.

Figure57: Certificate Server Configuration Page

Table54: Certificate Server Settings

Settings	Guidelines
Name	Specify a name to identify the server within the system.
User Name Template	Specify a username template. Specify how the system should construct a username. You may use any combination of certificate variables contained in angle brackets and plain text. The username template you configure must be consistent with the MDM certificate template configuration. Your goal is to identify the values specified in the MDM certificate that are to be used as the username in the system. This value populates the <USER> and <USERNAME> session variables for use throughout the rest of the system configuration. For example, suppose the certificate DN is: CN=<EnrollmentUser>, serialNumber=<DeviceUid>, o=Company. With this configuration, the certificate could identify both the user and the device. In this example, the username template is <certDN.CN>.

Adding the MDM Certificate to the Trusted Client CA Configuration

The system uses the uploaded certificate to verify that the browser-submitted certificate is valid. You must upload the MDM certificate that signed the client certificate that was pushed to the

mobile devices. Typically, you obtain this certificate from the MDM when your company establishes its account with them.

To import a trusted client CA certificate:

1. Select **System > Configuration > Certificates > Trusted Client CAs** to display the page shown in Figure 58.

Figure58: Trusted Client CA Management Page
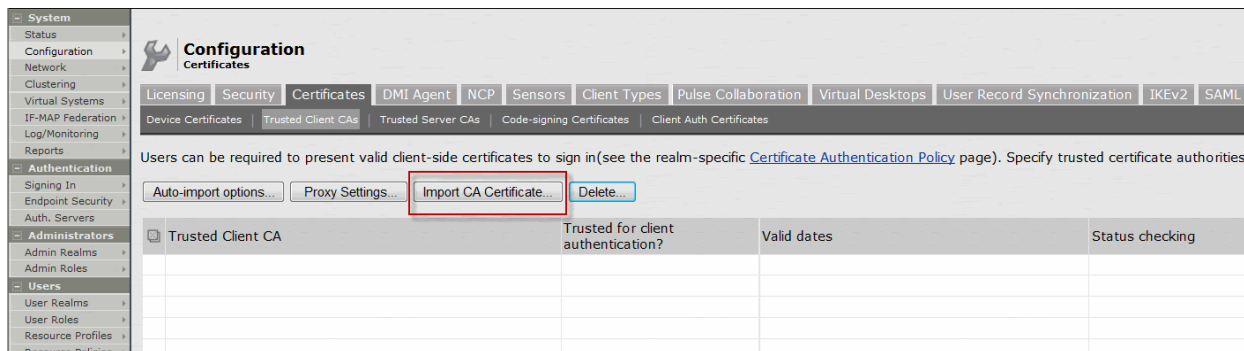
2. Click **Import CA Certificate** to display the page shown in Figure 59.

Figure59: Import Trusted Client CA Page
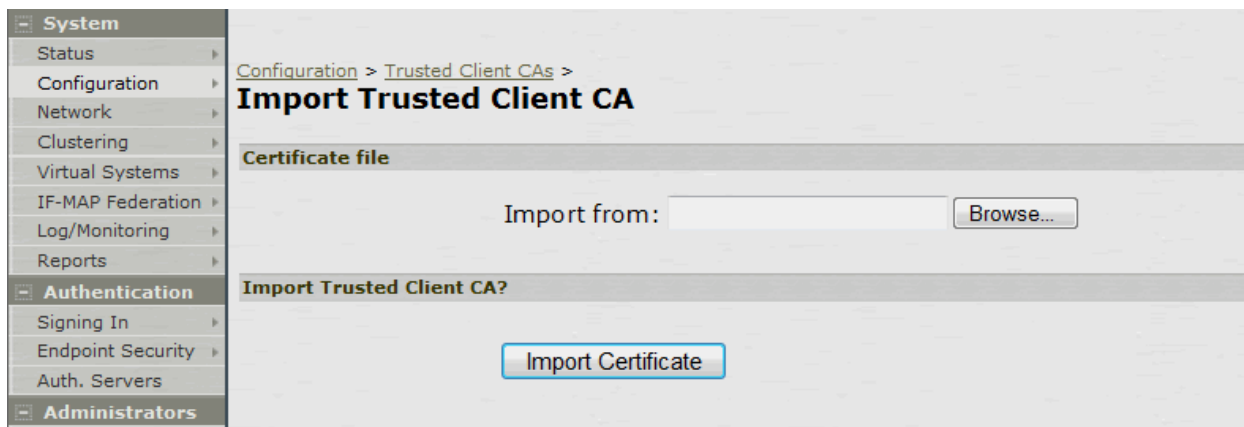
3. Browse to the certificate file, select it, and click **Import Certificate** to complete the import operation.

4. Click the link for the Trusted Client CA to display its details. Figure 60 shows the configuration for this example.

Figure60: Trusted Client CA Configuration for AirWatch

Configuring User Roles

User roles are classifiers for network access control policies. You create a set of roles to use in your classification scheme: device status is MDM enrollment complete or incomplete; device status is MDM-policy compliant or noncompliant; device is employee owned or company owned; device platform is iOS, Android, or neither; and so forth.

The user role configuration also includes options to customize user interface features that are appropriate for a particular role. For MDM deployments, you can use the Personalized Greeting UI option to send a notification message to the device when the role has been applied.

To configure user roles:

1. Select **Users > User Role** to navigate to the role configuration page.

2. Click **New Role** to display the configuration page shown in Figure 61.

3. Complete the configuration for general options as described in Table 55.

4. Save the configuration.

5. Click **UI options** to display the configuration page shown in Figure 62.

6. Complete the configuration for UI options as described in Table 55.

7. Save the configuration.

8. Click **Session Options** to display the configuration page shown in Figure 63.

9. Complete the configuration for session options as described in Table 55.

10. Save the configuration.

Figure61: User Role Configuration Page – General Settings

Figure62: User Role Configuration Page – UI Options

Figure63: User Role Configuration Page – Session Options

Table55: User Role Configuration Guidelines

Settings	Guidelines
Overview tab	
Name	Specify a name for the configuration.
Description	Describe the purpose of the role so that other administrators are aware of it.
Options	Select **UI Options** so that you can customize a message to be sent to the device when the role is applied.
UI Options tab	
Personalized greeting	Select the **Show notification message** option and enter a message to be sent to the device (through the MDM API) after sign-in and this role has been applied.
	In this example, we are using the system to enforce MDM enrollment by flagging compromised devices. The message, therefore, is:
	Your device is compromised. Network access may be limited.
	The message is forwarded to the device using the MDM server Push Notification feature.
	NOTE: When multiple roles are assigned, UI options are not merged. The UI options for the first role that matches are applied.
Session Options	
Session lifetime	Use the session lifetime options to establish the time limits that would require the user to sign in again.

Configuring a Realm and Role Mapping Rules

The user realm configuration associates the authentication server data and MDM server data with user roles.

To configure the realm and role mapping rules:

1. Select **Users > User Realms > New User Realm** to display the configuration page shown in Figure 64.

2. Complete the configuration as described in Table 56.

3. Save the configuration.

   Upon saving the new realm, the system displays the role mapping rules page.

4. Click **New Rule** to display the configuration page shown in Figure 65.

5. Complete the configuration as described in Table 57.

6. Save the configuration.

7. Click the **Authentication Policy** tab and then click the **Certificate** subtab to display the certificate restriction configuration page shown in Figure 66.

8. Complete the configuration as described in Table 59.

9. Save the configuration.

Figure64: Realm Configuration Page

Table56: Realm Configuration Guidelines

Settings	Guidelines
Name	Specify a name for the realm.
	If you enable sign-in using a realm suffix in the sign-in policy configuration, the realm name must match the username realm configured in the VPN profile. See Figure 53.
Description	Describe the purpose of the realm so that other administrators are aware of it.
Servers	
Authentication	Select the user authentication server for this realm's users. This example uses the certificate server configured in the earlier step. When you use a certificate server, users are not prompted for their credentials. You can also select the authentication server used for employees. In that case, users are prompted by the sign-in page to provide their username and password.
User Directory/Attribute	Do not select.
Accounting	Do not select.
Device Attributes	Select the MDM server configured in the earlier step.
Dynamic Policy Evaluation	
Dynamic Policy Evaluation	Do not select this option. A limitation for this release is that role evaluation occurs only when the user signs in. To force role reevaluation, you must force the users to sign in again.
Refresh interval	Do not select.
Refresh roles	Do not select.
Refresh resource policies	Do not select.
Session Migration	
Session Migration	Do not select this option. Session migration is useful for endpoints running Pulse Secure client software, which is not the case for the endpoints in this MDM example.

Figure65: Role Mapping Configuration Page

Table57: Role Mapping Configuration Guidelines

Settings	Guidelines
Rule based on	Select **Device Attribute** and click **Update** to update the configuration page so that it displays settings for role mapping using device attributes.
Name	Specify a name for the configuration.
Rule	Select a device attribute (see Table 58) and a logical operator (is or is not), and type a matching value or value pattern. In this example, select **isCompromised** and the logical operator **is**, and enter the value **1** (true). This means that devices with a compromised status match the rule.
Role assignment	Select the roles to apply if the data matches the rule.

TIP: You likely are to create multiple roles and role-mapping rules to assign roles for different policy purposes. Your realm can have a set of rules based on user attribute, group membership, and device attribute. Be mindful that the user and device can map to multiple roles. Use stop rules and order your rules carefully to implement the policy that you want.

Table 58 describes the AirWatch record attributes that can be used in role mapping rules.

Table58: AirWatch Device Attributes

Role Mapping Attribute Name	AirWatch Attribute Name	Description	Data Type
BlockLevelEncryption	BlockLevelEncryption	True if block-level encryption is enabled; false otherwise.	Boolean
complianceReason	ComplianceStatus	Values: Compliant, Non-Compliant.	String
CompromisedStatus	CompromisedStatus	True if the status is compromised; false otherwise.	Boolean
DataProtectionEnabled	DataProtectionEnabled	True if data protection is enabled; false otherwise.	Boolean
deviceId	Id.Value	Device identifier.	String
deviceName	DeviceFriendlyName	The concatenated name used to identify the device/user combination.	String
FileLevelEncryption	FileLevelEncryption	True if file-level encryption is enabled; false otherwise.	Boolean
IMEI	Imei	IMEI number of the device.	String
isCompliant	ComplianceStatus	Values: Compliant.	String
isCompromised	CompromisedStatus	True if the device is compromised; false otherwise.	Boolean
isEnrolled	EnrollmentStatus	True if MDM value is Enrolled; false otherwise.	Boolean
IsPasscodeCompliant	IsPasscodeCompliant	True if the passcode is compliant with the MDM policy; false otherwise	Boolean
IsPasscodePresent	IsPasscodePresent	True if a passcode has been configured; false otherwise.	Boolean
LastComplianceCheckOn	LastComplianceCheckOn	The refresh date and timestamp of the last status reported.	Timestamp
LastCompromisedCheckOn	LastCompromisedCheckOn	The refresh date and timestamp of the last status reported.	Timestamp
lastSeen	LastSeen	Date and time the device last made successful contact with the MDM.	Timestamp
LocationGroupName	LocationGroupName	MDM location group configuration value.	String
macAdress	MacAddress	The Wi-Fi MAC address.	String
model	Model	Model is automatically reported by the device during registration.	String
osVersion	OperatingSystem	OS version.	String
ownership	Ownership	Values: C, E, or S (Corporate, Employee, or Shared).	String
phoneNumber	PhoneNumber	Phone number entered during registration.	String
platform	Platform	Platform specified during registration.	String
serialNumber	SerialNumber	Serial number.	String
UDID	Udid	Unique device identifier.	String
userEmail	UserEmailAddress	E-mail address of device user.	String
userName	UserName	Name of device user.	String

UUID	Uuid	Universal unique identifier.	String

NOTE: By design, you should be able to specify **true** or **false**, or **1** or **0**, for Boolean data types in your role mapping rules. Due to a issue in this release, you must use **1** for true and **0** for false.

Figure66: Realm Configuration Page – Certificate Restrictions

Table59: Realm Configuration Certificate Restriction Guidelines

Settings	Guidelines
Allow all users	Do not select this option. If you select this option, the system does not request a client certificate during the TLS handshake.
Allow all users and remember certificate	If you select this option, the system requests a client certificate during the TLS handshake. It does allow endpoints to authenticate without a client certificate. For those with a client certificate, the certificate attributes are placed in the session context. TIP: Without a certificate, device attributes cannot be determined, and the session can be mapped only to roles that do not require particular device attributes. You might use this option to grant restricted access or to send a notification that MDM enrollment is required for a greater level of access.
Only allow users with a client-side certificate	If you select this option, the system requests a client certificate during the TLS handshake. It does not allow endpoints to authenticate without a valid client certificate. If the realm is configured with a certificate server, like this example, this option is the only option that can be selected.

Configuring a Sign-In Policy

A sign-in policy associates devices with a realm.

To configure a sign-in policy:

1. Select **Authentication > Signing In > Sign-In Policies** to navigate to the sign-in policies configuration page.

2. Click **New URL** to display the configuration page shown in Figure 67.

---

3. Complete the configuration as described in Table 60.

4. Save the configuration.

Figure67: Sign-In Policy Configuration Page

Table60: Sign-In Policy Configuration Guidelines

Settings	Guidelines
User type	Select **Users**.
Sign-in URL	Enter a URL.
Description	Describe the purpose of the sign-in policy so that other administrators are aware of it.
Sign-In Page	Select a sign-in page.
Authentication Realm	
User experience	Select one of the following options: • **User types the realm name** • **User picks from a list of authentication realms**
Realm	Select the realm you configured in the earlier step.
Configure Sign-in Notifications	
Pre-Auth Sign-in Notification	Not used in this scenario.
Post-Auth Sign-in Notification	Not used in this scenario.

Configuring a Resource Policy

A resource policy enforces role-based access to resources accessed during the SSL VPN session. You use the device access management framework to assign roles to devices, and you use the resource policy to deny access to resources that should not be downloaded onto a specific device platform—in this example, Android devices.

In this scenario, the role configuration and role mapping configuration create a classification for Android devices. Figure 68 shows the user role configuration.

Figure68: User Role Configuration Page – General Settings
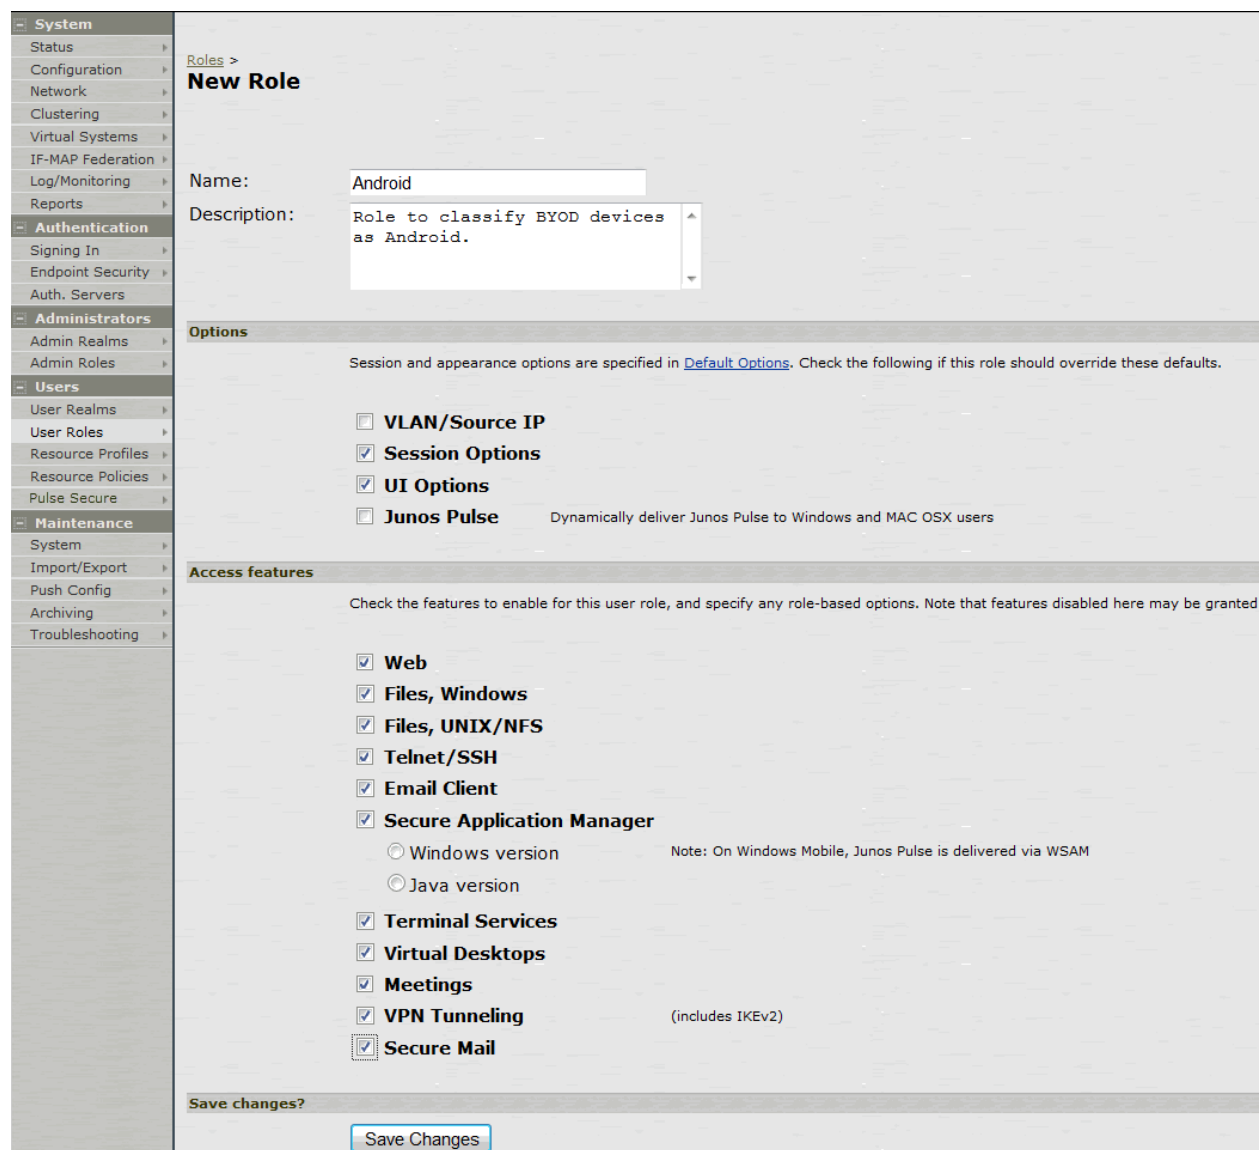
Figure 69 shows the role mapping configuration.

Figure69: Role Mapping Configuration Page

To configure a resource policy:

1. Select **Resource Policies > VPN Tunneling > Access Control** to display the access control policy configuration pages.

2. Click **New Policy** to display the configuration page shown in Figure 70.

3. Complete the configuration as described in Table 61.

4. Save the configuration.

Figure70: Resource Access Policy Configuration Page
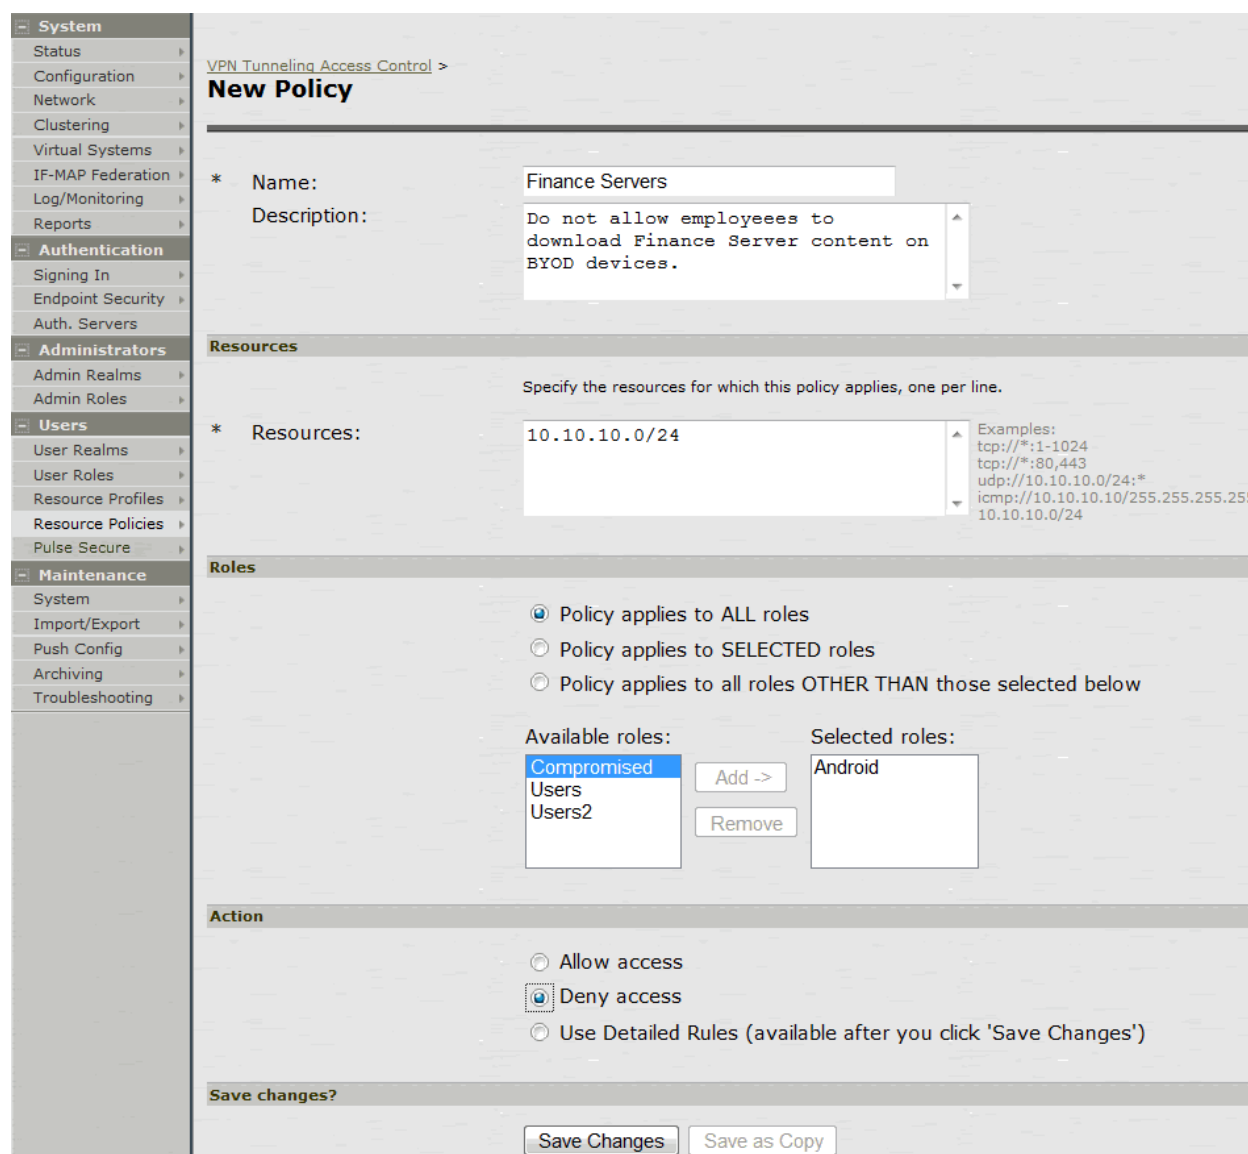
Table61: Resource Access Policy Configuration Guidelines

Settings	Guidelines
Name	Specify a name for the configuration.
Description	Describe the purpose of the configuration so that other administrators are aware of it.
Resources	
Resources	Specify the resources for which this policy applies, one per line.
Roles	
Roles	Select the roles to which the policy applies. In this example, **Android** is selected.
Action	

Action	Select one of the following actions:
	• **Allow Access**
	• **Deny Access**
	• **Use Detailed Rules**
	In this example, we deny access from Android devices.

Related Documentation

- Using Logs to Verify Proper Configuration

- *User and Policy Administration Overview*

- Using Policy Tracing and Debug Logs

- Understanding the Device Access Management Framework

Deploying a BYOD Policy for MobileIron Managed Devices

This example shows how to use policies to enable security based on device identity, device posture, or user identity in a bring your own device (BYOD) environment for an enterprise that uses MobileIron® for mobile device management (MDM). It includes the following information:

- Solution Overview

- Requirements

- Configuring the MobileIron MDM Service

- Configuring the Device Access Management Framework

- Configuring a Resource Policy

Solution Overview

In the past, to ensure security and manageability of the corporate network, enterprise information technology (IT) departments had restricted network access to company-issued equipment. For mobile phones, the classic example was the company-issued BlackBerry handset. As powerful mobile smart phones and tablets have become commonly held personal possessions, the trend in enterprise IT has been to stop issuing mobile equipment and instead allow employees to use their personal smart phones and tablets to conduct business activities. This has lowered equipment costs, but BYOD environments pose capacity planning and security challenges: how can an enterprise track network access by non-company-issued devices?

MDM vendors have emerged to address the first issue. MDMs such as MobileIron provide enrollment and posture assessment services that prompt employees to enter data about their mobile devices.

The integration of MDM device attribute and posture assessment status data into Pulse Secure remote access solutions bolsters security in the BYOD arena by enabling the enterprise IT to enforce network security policies that restrict the mobile devices that can access the network and protected resources in the same way our SSL VPN solution restricts user access. The access management framework uses the MDM data as a factor in making network and resource access policy decisions.

Figure 71 shows a deployment with Connect Secure and the MobileIron MDM cloud service.

Figure71: Solution Topology

The solution shown in this example leverages the Pulse Secure access management framework to support attribute-based network access control for mobile devices. In the *device access management framework*, the MDM is a device authorization server and MDM record attributes are the basis for access policy determinations. For example, suppose your enterprise wants to enforce a policy that allows access only to mobile devices that have enrolled with the MDM or are compliant with the MDM posture assessment policies. You can use the attributes and status maintained by the MDM in role-mapping rules to implement the policy.

In this framework, a native supplicant is used to authenticate the user of the device. The device itself is identified using a client certificate that contains device identity. The client certificate can be used to identify the device against the MDM records and authenticate the user against a certificate server.

The Pulse Secure solution supports granular, attribute-based resource access policies. For example, you can implement policies that allow devices that have a clean MDM posture assessment and are compliant with MDM policies to access the network, but deny access to servers when you want to prevent downloads to employee-owned devices or to a particular platform that might be vulnerable.

Requirements

Table 62 lists version information for the solution components shown in this example.

Table62: Component Version Information

Component	Version
Connect Secure	Release 8.0r1 or later is required.
MobileIron MDM	Release 5.6 is used in this example. Any version that supports the device ID and device attributes you plan to query is compatible.

Configuring the MobileIron MDM Service

This solution assumes you know how to configure and use the features of your MDM, and that you can enroll employees and their devices. For more information about the MobileIron MDM, refer to its documentation and support resources. This section focuses on the following elements of the MDM configuration that are important to this solution:

- Device identifier—The primary key for device records. Your MDM configuration determines whether a universal unique identifier (UUID), unique device identifier (UDID), or serial number is used as the device identifier. For MobileIron, UUID is supported and recommended.

- Device attributes—A standard set of data maintained for each device. For MobileIron, see Table 63.

When the user installs the MDM application on the device and completes enrollment, the MDM pushes the device certificate to the device. After enrollment, the MDM maintains a database record that includes information about the enrollee—attributes related to device identity, user identity, and posture assessment against MDM policies. Table 63 describes these attributes. In this solution, these attributes are used in the role mapping that is the basis for network access and resource access policies. When you configure role-mapping rules, you specify the normalized Connect Secure attribute name.

Table63: MobileIron Device Attributes

MobileIron Attribute	Normalized Connect Secure Name	Description	Data Type
@id	deviceId	Device identifier.	String
blockedReason	blockedReason	Reason MDM has blocked the device. Can be a multivalued string. Values are: • AllowedAppControlPolicyOutOfCompliance • AppControlPolicyOutOfCompliance • DataProtectionNotEnabled • DeviceAdminDeactivated • DeviceComplianceStatusUnknown • DeviceCompliant • DeviceCompromised • DeviceExceedsPerMailboxLimit • DeviceManuallyBlocked • DeviceNotRegistered • DisallowedAppControlPolicyOutOfCompliance • ExchangeReported • HardwareVersionNotAllowed • OsVersionLessThanSupportedOsVersion • PolicyOutOfDate • RequiredAppControlPolicyOutOfCompliance	String

compliance	complianceReason	MDM policy compliance status. Can be a multivalued string. Values are:	String
		• AllowedAppControlPolicyOutOfCompliance	
		• AppControlPolicyOutOfCompliance	
		• DataProtectionNotEnabled	
		• DeviceAdminDeactivated	
		• DeviceComplianceStatusUnknown	
		• DeviceCompliant	
		• DeviceCompromised	
		• DeviceExceedsPerMailboxLimit	
		• DeviceManuallyBlocked	
		• DeviceNotRegistered	
		• DisallowedAppControlPolicyOutOfCompliance	
		• ExchangeReported	
		• HardwareVersionNotAllowed	
		• OsVersionLessThanSupportedOsVersion	
		• PolicyOutOfDate	
		• RequiredAppControlPolicyOutOfCompliance	
compliance	isCompliant	True if the device is in compliance with its MDM security policies; false otherwise.	Boolean
compliance	isCompromised	True if the device is compromised; false otherwise.	Boolean
countryName	countryName	Country name corresponding with the country code of the device.	String
currentPhoneNumber	phoneNumber	Phone number entered during registration.	String
emailAddress	userEmail	E-mail address of device user.	String
employeeOwned	ownership	Values: Employee or Corporate.	String
homeOperator	homeOperator	The service operator for the device when it is not roaming.	String
iPhone IMEI (iOS), imei (Android)	Imei	IMEI number of the device.	String
iPhone UDID	UDID	Unique device identifier.	String
isBlocked	isBlocked	True if the device is blocked from accessing the ActiveSync server; false otherwise.	Boolean
isQuarantined	isQuarantined	True if the device is quarantined by the MDN; false otherwise.	Boolean
lastConnectAt	lastSeen	Date and time the device last made successful contact with the MDM.	Timestamp
manufacturer	manufacturer	Manufacturer is automatically reported by the device during registration.	String
mdmManaged	mdmManaged	True if the MDM profile is enabled on the device; false otherwise. This field applies only to iOS devices. For other devices, the value is always false.	Boolean
ModelName, model, device_model	model	Model is automatically reported by the device during registration.	String
name	deviceName	The concatenated name used to identify the device/user combination.	String
operator	operator	Service provider. The value PDA indicates no operator is associated with the device.	String
OSVersion (iOS), os_version (Android)	osVersion	OS version.	String

platform	platform	Platform specified during registration.	String
principal	userId	User ID.	String
quarantinedReason	quarantinedReason	MDM policy compliance status. Can be a multivalued string. Values are: • AllowedAppControlPolicyOutOfCompliance • AppControlPolicyOutOfCompliance • DataProtectionNotEnabled • DeviceAdminDeactivated • DeviceComplianceStatusUnknown • DeviceCompliant • DeviceCompromised • DeviceExceedsPerMailboxLimit • DeviceManuallyBlocked • DeviceNotRegistered • DisallowedAppControlPolicyOutOfCompliance • ExchangeReported • HardwareVersionNotAllowed • OsVersionLessThanSupportedOsVersion • PolicyOutOfDate • RequiredAppControlPolicyOutOfCompliance	
SerialNumber	serialNumber	Serial number.	String
statusCode	isEnrolled	True if the device has completed enrollment or registration; false otherwise.	Boolean
uuid	UUID	Universal unique device identifier.	String
userDisplayName	userName	Name of device user.	String
wifi_mac (iOS), wifi_mac_addr (Android)	macAdress	The Wi-Fi MAC address.	String

To configure the MDM:

1. Enroll devices in the MDM using the methods supported by the MDM.

2. Create a Simple Certificate Enrollment Protocol (SCEP) configuration that specifies the field and type of identifier for client device certificates. See Figure 72.

   The MDM configuration templates provide flexibility in how the device identifier can be placed in the device certificate's subject or alternative subject. We recommend you include the user ID in the certificate, so the certificate can identify both the user and the device. For example:

   CN=<DEVICE_UUID>, uid=<USER_ID>, o=Company

3. Create a VPN configuration that specifies the Juniper SSL connection type and the URL for the system sign-in page. See Figure 73. During the enrollment process, this profile is provisioned to the device. Select the SCEP configuration completed in Step 1.

4. Select the VPN configuration and apply it to a group label you have provisioned to manage this group of devices. See Figure 74.

5. Apply the group label to the devices when you add them to the MDM. See Figure 75. If they have already been added to the MDM, use the edit configuration utilities in the device inventory page to apply the group label.

Figure72: MobileIron SCEP Configuration

Figure73: MobileIron VPN Configuration

Figure74: Applying the VPN Configuration to a Label

Figure75: Adding a Device to the MDM

Configuring the Device Access Management Framework

This section describes the basic steps for configuring the device access management framework:

1. Configuring the MDM Authentication Server

2. Configuring the Certificate Server

3. Adding the MDM Certificate to the Trusted Client CA Configuration

4. Configuring User Roles

5. Configuring a Realm and Role Mapping Rules

6. Configuring a Sign-In Policy

Configuring the MDM Authentication Server

The MDM authentication server configuration is used by the system to communicate with the MDM. In the device access management framework, the MDM server is used as the device authorization server.

To configure the authentication server:

1. Select **Authentication > Auth Servers** to navigate to the authentication server configuration pages.

2. Select **MDM Server** and click **New Server** to display the configuration page shown in Figure 76.

3. Complete the configuration as described in Table 64.

4. Save the configuration.

Figure76: Authentication Server Configuration Page

Table64: Authentication Server Configuration Guidelines

Settings	Guidelines
Name	Specify a name for the configuration.
Type	Select **MobileIron**.
Server	
Server Url	Specify the URL for your MobileIron server. This is the URL MobileIron has instructed you to use to access its RESTful Web API (also called a RESTful Web service). The URL for the MobileIron server used in this example has the following form: **https://m.mobileiron.net/pulsesecuretest** NOTE: You must configure your firewalls to allow communication between these two nodes over port 443.
Viewer Url	Specify the URL for the MobileIron report viewer. This URL is used for links from the Active Users page to the MobileIron report viewer. The URL for the MobileIron viewer for this example has the following form: **https://m.mobileiron.net/ pulsesecuretest /admin/admin.html#smartphones:all**
Request Timeout	Specify a timeout period (0-60 seconds) for queries to the MDM server. The default is 15 seconds. Calibrate this value based on your observations on how long a query to the MDM server takes over your network. If your network experiences latency when querying the MDM cloud service, increase the timeout to account for the latency. The system queries the MDM when a user attempts to sign in. If a timeout occurs, role mapping proceeds without attributes.
Administrator	
Username	Specify the username for an account that has privileges to access the MobileIron RESTful Web API.
Password	Specify the corresponding password.
Device Identifier	
ID Template	Construct a template to derive the device identifier from the certificate attributes. The template can contain textual characters as well as variables for substitution. The variables are the same as those used in role mapping custom expressions and policy conditions. Enclose variables in angle brackets like this <variable>. For example, suppose the certificate DN is: CN=<DEVICE_UDID>, uid=<USER_ID>, o=Company. With this configuration, the certificate could identify both the user and the device. In this example, the device ID template is <certDN.CN>.
ID Type	Select the device identifier type that matches the selection in the MDM certificate configuration: • **UUID**—Not applicable for the MobileIron MDM. • **Serial Number**—The device serial number. • **UDID**—The device unique device identifier. This is supported by the MobileIron MDM.

Configuring the Certificate Server

The certificate server configuration enables device users to authenticate using the certificate pushed to the device by the MDM. The certificates are used for user authentication, and the users do not have to enter user credentials.

To configure authentication with the certificate server:

1. Select **Authentication** > **Auth. Servers**.

2. Select **Certificate Server** and click **New Server**.

3. Complete the configuration as described in Table 65.

4. Save the configuration.

Table65: Certificate Server Settings

Settings	Guidelines
Name	Specify a name to identify the server within the system.
User Name Template	Specify a username template. Specify how the system should construct a username. You may use any combination of certificate variables contained in angle brackets and plain text. The username template you configure must be consistent with the MDM certificate template configuration. Your goal is to identify the values specified in the MDM certificate that are to be used as the username in the system. This value populates the <USER> and <USERNAME> session variables for use throughout the rest of the system configuration. For example, suppose the certificate DN is: CN=<DEVICE_UDID>, uid=<USER_ID>, o=Company. With this configuration, the certificate could identify both the user and the device. In this example, the username template is <certDN.UID>.

Adding the MDM Certificate to the Trusted Client CA Configuration

The system uses the uploaded certificate to verify that the browser-submitted certificate is valid. You must upload the MDM certificate that signed the client certificate that was pushed to the mobile devices. Typically, you obtain this certificate from the MDM when your company establishes its account with them.

To import a trusted client CA certificate:

1. Select **System > Configuration > Certificates > Trusted Client CAs** to display the page shown in Figure 77.

Figure77: Trusted Client CA Management Page
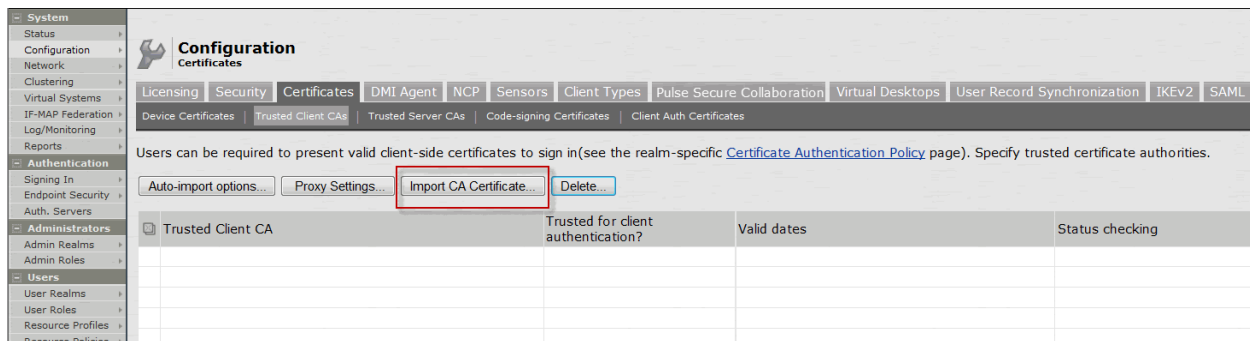
2. Click **Import CA Certificate** to display the page shown in Figure 78.

Figure78: Import Trusted Client CA Page

3. Browse to the certificate file, select it, and click **Import Certificate** to complete the import operation.

---

4. Click the link for the Trusted Client CA to display its details. Figure 79 shows the configuration for this example.

Figure79: Trusted Client CA Configuration for MobileIron

Configuring User Roles

User roles are classifiers for network access control policies. You create a set of roles to use in your classification scheme: device status is MDM enrollment complete or incomplete; device status is MDM-policy compliant or noncompliant; device is employee owned or company owned; device platform is iOS, Android, or neither; and so forth.

The user role configuration also includes options to customize user interface features that are appropriate for a particular role. For MDM deployments, you can use the Personalized Greeting UI option to send a notification message to the device when the role has been applied.

To configure user roles:

1. Select **Users > User Role** to navigate to the role configuration page.

2. Click **New Role** to display the configuration page shown in Figure 80.

3. Complete the configuration for general options as described in Table 66.

4. Save the configuration.

5. Click **UI options** to display the configuration page shown in Figure 81.

6. Complete the configuration for UI options as described in Table 66.

7. Save the configuration.

8. Click **Session Options** to display the configuration page shown in Figure 82.

9. Complete the configuration for session options as described in Table 66.

10. Save the configuration.

Figure80: User Role Configuration Page – General Settings

Figure81: User Role Configuration Page – UI Options

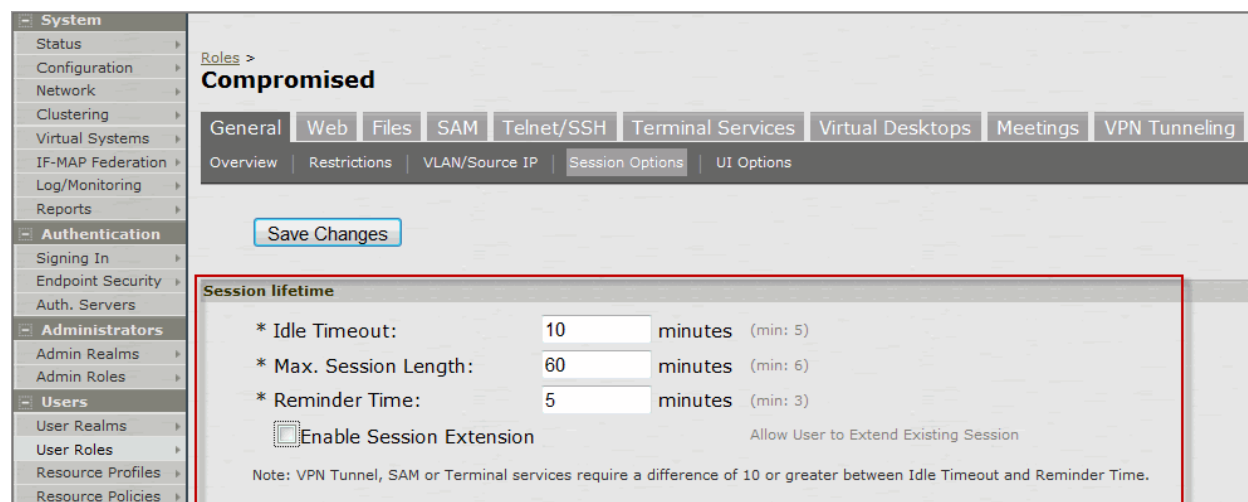Figure82: User Role Configuration Page – Session Options

Table66: User Role Configuration Guidelines

Settings	Guidelines
Overview tab	
Name	Specify a name for the configuration.
Description	Describe the purpose of the role so that other administrators are aware of it.
Options	Select **UI Options** so that you can customize a message to be sent to the device when the role is applied.
UI Options tab	
Personalized greeting	Select the **Show notification message** option and enter a message to be sent to the device (through the MDM API) after sign-in and this role has been applied.
	In this example, we are using the system to enforce MDM enrollment by flagging compromised devices. The message, therefore, is:
	Your device is compromised. Network access may be limited.
	The message is forwarded to the device using the MDM server Push Notification feature.
	NOTE: When multiple roles are assigned, UI options are not merged. The UI options for the first role that matches are applied.
Session Options	
Session lifetime	Use the session lifetime options to establish the time limits that would require the user to sign in again.

Configuring a Realm and Role Mapping Rules

The user realm configuration associates the authentication server data and MDM server data with user roles.

To configure the realm and role mapping rules:

1. Select **Users > User Realms > New User Realm** to display the configuration page shown in Figure 83.

2. Complete the configuration as described in Table 67.

3. Save the configuration.

   Upon saving the new realm, the system displays the role mapping rules page.

4. Click **New Rule** to display the configuration page shown in Figure 84.

5. Complete the configuration as described in Table 68.

6. Save the configuration.

7. Click the **Authentication Policy** tab and then click the **Certificate** subtab to display the certificate restriction configuration page shown in Figure 85.

8. Complete the configuration as described in Table 70.

9. Save the configuration.

Figure83: Realm Configuration Page

Table67: Realm Configuration Guidelines

Settings	Guidelines
Name	Specify a name for the realm.
	If you enable sign-in using a realm suffix in the sign-in policy configuration, the realm name must match the username realm configured in the VPN profile. See Figure 75.
Description	Describe the purpose of the realm so that other administrators are aware of it.
Servers	
Authentication	Select the user authentication server for this realm's users. This example uses the certificate server configured in the earlier step. When you use a certificate server, users are not prompted for their credentials. You can also select the authentication server used for employees. In that case, users are prompted by the sign-in page to provide their username and password.
User Directory/Attribute	Do not select.
Accounting	Do not select.
Device Attributes	Select the MDM server configured in the earlier step.
Dynamic Policy Evaluation	
Dynamic Policy Evaluation	Do not select this option. A limitation for this release is that role evaluation occurs only when the user signs in. To force role reevaluation, you must force the users to sign in again.
Refresh interval	Do not select.
Refresh roles	Do not select.
Refresh resource policies	Do not select.
Session Migration	
Session Migration	Do not select this option. Session migration is useful for endpoints running Pulse Secure client software, which is not the case for the endpoints in this MDM example.

Figure84: Role Mapping Configuration Page
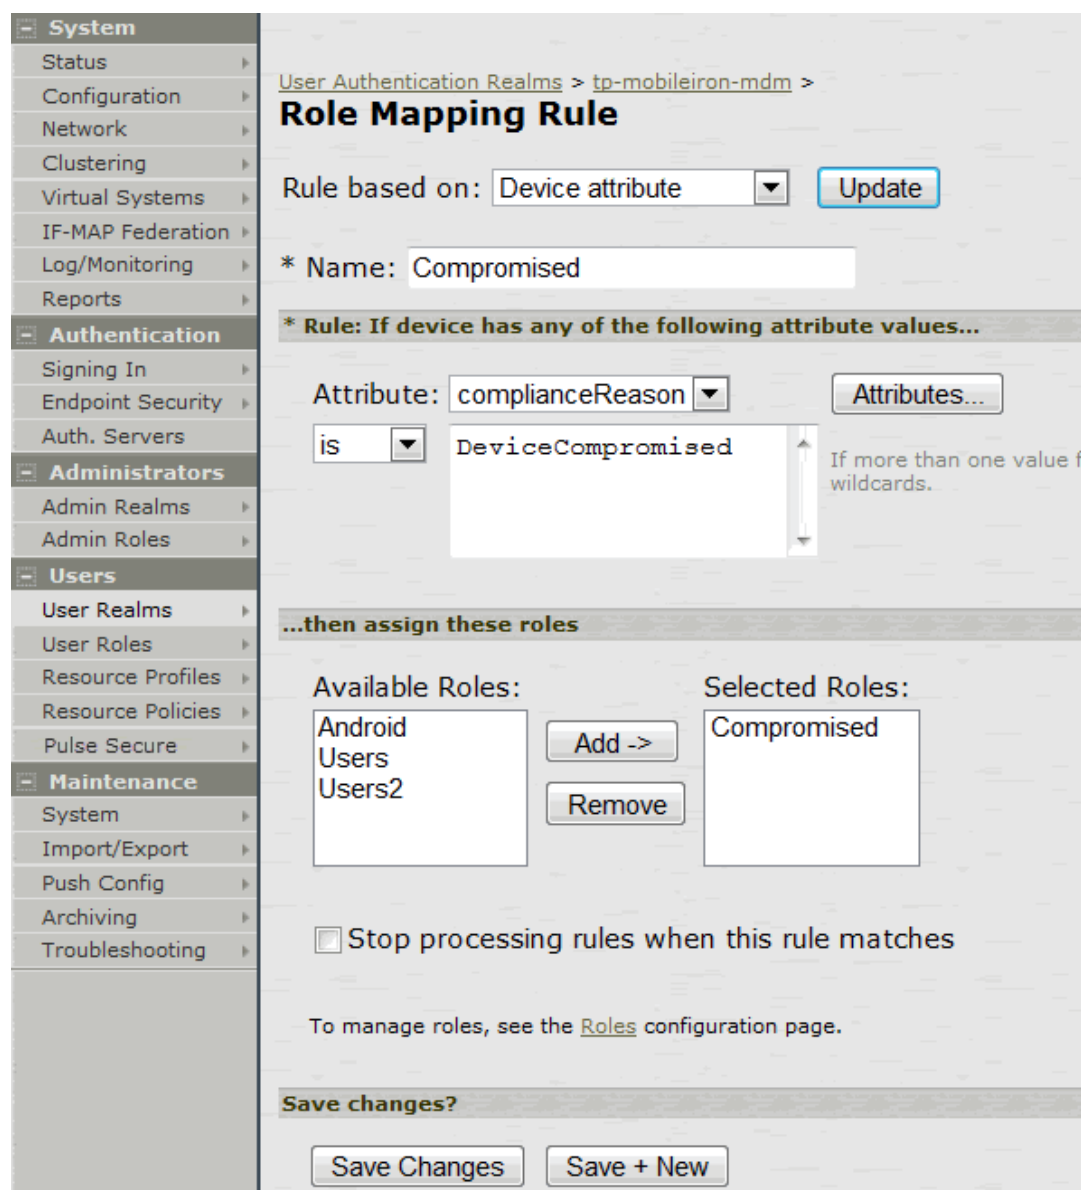
Table68: Role Mapping Configuration Guidelines

Settings	Guidelines
Rule based on	Select **Device Attribute** and click **Update** to update the configuration page so that it displays settings for role mapping using device attributes.
Name	Specify a name for the configuration.
Rule	Select a device attribute (see Table 69) and a logical operator (is or is not), and type a matching value or value pattern. In this example, select **isCompromised** and the logical operator **is**, and enter the value **1** (true). This means that devices with a compromised status match the rule.
Role assignment	Select the roles to apply if the data matches the rule.

TIP: You likely are to create multiple roles and role-mapping rules to assign roles for different policy purposes. Your realm can have a set of rules based on user attribute, group membership, and device attribute. Be mindful that the user and device can map to multiple roles. Use stop rules and order your rules carefully to implement the policy that you want.

Table 69 describes the MobileIron record attributes that can be used in role mapping rules.

Table69: MobileIron Device Attributes

Role Mapping Attribute Name	MobileIron Attribute Name	Description	Data Type
blockedReason	blockedReason	Reason MDM has blocked the device. Can be a multivalued string. Values are: • AllowedAppControlPolicyOutOfCompliance • AppControlPolicyOutOfCompliance • DataProtectionNotEnabled • DeviceAdminDeactivated • DeviceComplianceStatusUnknown • DeviceCompliant • DeviceCompromised • DeviceExceedsPerMailboxLimit • DeviceManuallyBlocked • DeviceNotRegistered • DisallowedAppControlPolicyOutOfCompliance • ExchangeReported • HardwareVersionNotAllowed • OsVersionLessThanSupportedOsVersion • PolicyOutOfDate • RequiredAppControlPolicyOutOfCompliance	String
complianceReason	compliance	MDM policy compliance status. Can be a multivalued string. Values are: • AllowedAppControlPolicyOutOfCompliance • AppControlPolicyOutOfCompliance • DataProtectionNotEnabled • DeviceAdminDeactivated • DeviceComplianceStatusUnknown • DeviceCompliant • DeviceCompromised • DeviceExceedsPerMailboxLimit • DeviceManuallyBlocked • DeviceNotRegistered • DisallowedAppControlPolicyOutOfCompliance • ExchangeReported • HardwareVersionNotAllowed • OsVersionLessThanSupportedOsVersion • PolicyOutOfDate • RequiredAppControlPolicyOutOfCompliance	String

countryName	countryName	Country name corresponding with the country code of the device.	String
deviceId	@id	Device identifier.	String
deviceName	name	The concatenated name used to identify the device/user combination.	String
homeOperator	homeOperator	The service operator for the device when it is not roaming.	String
Imei	iPhone IMEI (iOS), imei (Android)	IMEI number of the device.	String
isBlocked	isBlocked	True if the device is blocked from accessing the ActiveSync server; false otherwise.	Boolean
isCompliant	compliance	True if the device is in compliance with its MDM security policies; false otherwise.	Boolean
isCompromised	compliance	True if the device is compromised; false otherwise.	Boolean
isEnrolled	statusCode	True if the device has completed enrollment or registration; false otherwise.	Boolean
isQuarantined	isQuarantined	True if the device is quarantined by the MDN; false otherwise.	Boolean
lastSeen	lastConnectAt	Date and time the device last made successful contact with the MDM.	Timestamp
manufacturer	manufacturer	Manufacturer is automatically reported by the device during registration.	String
macAdress	wifi_mac (iOS), wifi_mac_addr (Android)	The Wi-Fi MAC address.	String
mdmManaged	mdmManaged	True if the MDM profile is enabled on the device; false otherwise. This field applies only to iOS devices. For other devices, the value is always false.	Boolean
model	ModelName, model, device_model	Model is automatically reported by the device during registration.	String
operator	operator	Service provider. The value PDA indicates no operator is associated with the device.	String
osVersion	OSVersion (iOS), os_version (Android)	OS version.	String
ownership	employeeOwned	Values: Employee or Corporate.	String
phoneNumber	currentPhoneNumber	Phone number entered during registration.	String
platform	platform	Platform specified during registration.	String

quarantinedReason	quarantinedReason	MDM policy compliance status. Can be a multivalued string. Values are:	
		• AllowedAppControlPolicyOutOfCompliance	
		• AppControlPolicyOutOfCompliance	
		• DataProtectionNotEnabled	
		• DeviceAdminDeactivated	
		• DeviceComplianceStatusUnknown	
		• DeviceCompliant	
		• DeviceCompromised	
		• DeviceExceedsPerMailboxLimit	
		• DeviceManuallyBlocked	
		• DeviceNotRegistered	
		• DisallowedAppControlPolicyOutOfCompliance	
		• ExchangeReported	
		• HardwareVersionNotAllowed	
		• OsVersionLessThanSupportedOsVersion	
		• PolicyOutOfDate	
		• RequiredAppControlPolicyOutOfCompliance	
serialNumber	SerialNumber	Serial number.	String
UDID	iPhone UDID	Unique device identifier.	String
userEmail	emailAddress	E-mail address of device user.	String
userId	principal	User ID.	String
userName	userDisplayName	Name of device user.	String
UUID	uuid	Universal unique device identifier.	String

NOTE: By design, you should be able to specify **true** or **false**, or **1** or **0**, for Boolean data types in your role mapping rules. Due to a issue in this release, you must use **1** for true and **0** for false.

Figure85: Realm Configuration Page – Certificate Restrictions

Table70: Realm Configuration Certificate Restriction Guidelines

Settings	Guidelines
Allow all users	Do not select this option. If you select this option, the system does not request a client certificate during the TLS handshake.
Allow all users and remember certificate	If you select this option, the system requests a client certificate during the TLS handshake. If the realm has been configured with a user authentication server, it does allow endpoints to authenticate without a client certificate. For those with a client certificate, the certificate attributes are placed in the session context. TIP: Without a certificate, device attributes cannot be determined, and the session can be mapped only to roles that do not require particular device attributes. You might use this option to grant restricted access or to send a notification that MDM enrollment is required for a greater level of access.
Only allow users with a client-side certificate	If you select this option, the system requests a client certificate during the TLS handshake. It does not allow endpoints to authenticate without a valid client certificate. If the realm is configured with a certificate server, like this example, this option is the only option that can be selected.

Configuring a Sign-In Policy

A sign-in policy associates devices with a realm.

To configure a sign-in policy:

1. Select **Authentication > Signing In > Sign-In Policies** to navigate to the sign-in policies configuration page.

2. Click **New URL** to display the configuration page shown in Figure 86.

3. Complete the configuration as described in Table 71.

4. Save the configuration.

Figure86: Sign-In Policy Configuration Page

Table71: Sign-In Policy Configuration Guidelines

Settings	Guidelines
User type	Select **Users**.
Sign-in URL	Enter a URL.
Description	Describe the purpose of the sign-in policy so that other administrators are aware of it.
Sign-In Page	Select a sign-in page.

345

Authentication Realm	
User experience	Select one of the following options: • **User types the realm name** • **User picks from a list of authentication realms**
Realm	Select the realm you configured in the earlier step.
Configure Sign-in Notifications	
Pre-Auth Sign-in Notification	Not used in this scenario.
Post-Auth Sign-in Notification	Not used in this scenario.

Configuring a Resource Policy

A resource policy enforces role-based access to resources accessed during the SSL VPN session. You use the device access management framework to assign roles to devices, and you use the resource policy to deny access to resources that should not be downloaded onto a specific device platform—in this example, Android devices.

In this scenario, the role configuration and role mapping configuration create a classification for Android devices. Figure 87 shows the user role configuration.

Figure87: User Role Configuration Page – General Settings

Figure 88 shows the role mapping configuration.

Figure88: Role Mapping Configuration Page

To configure a resource policy:

1. Select **Resource Policies > VPN Tunneling > Access Control** to display the access control policy configuration pages.

2. Click **New Policy** to display the configuration page shown in Figure 89.

3. Complete the configuration as described in Table 72.

4. Save the configuration.

Figure89: Resource Access Policy Configuration Page
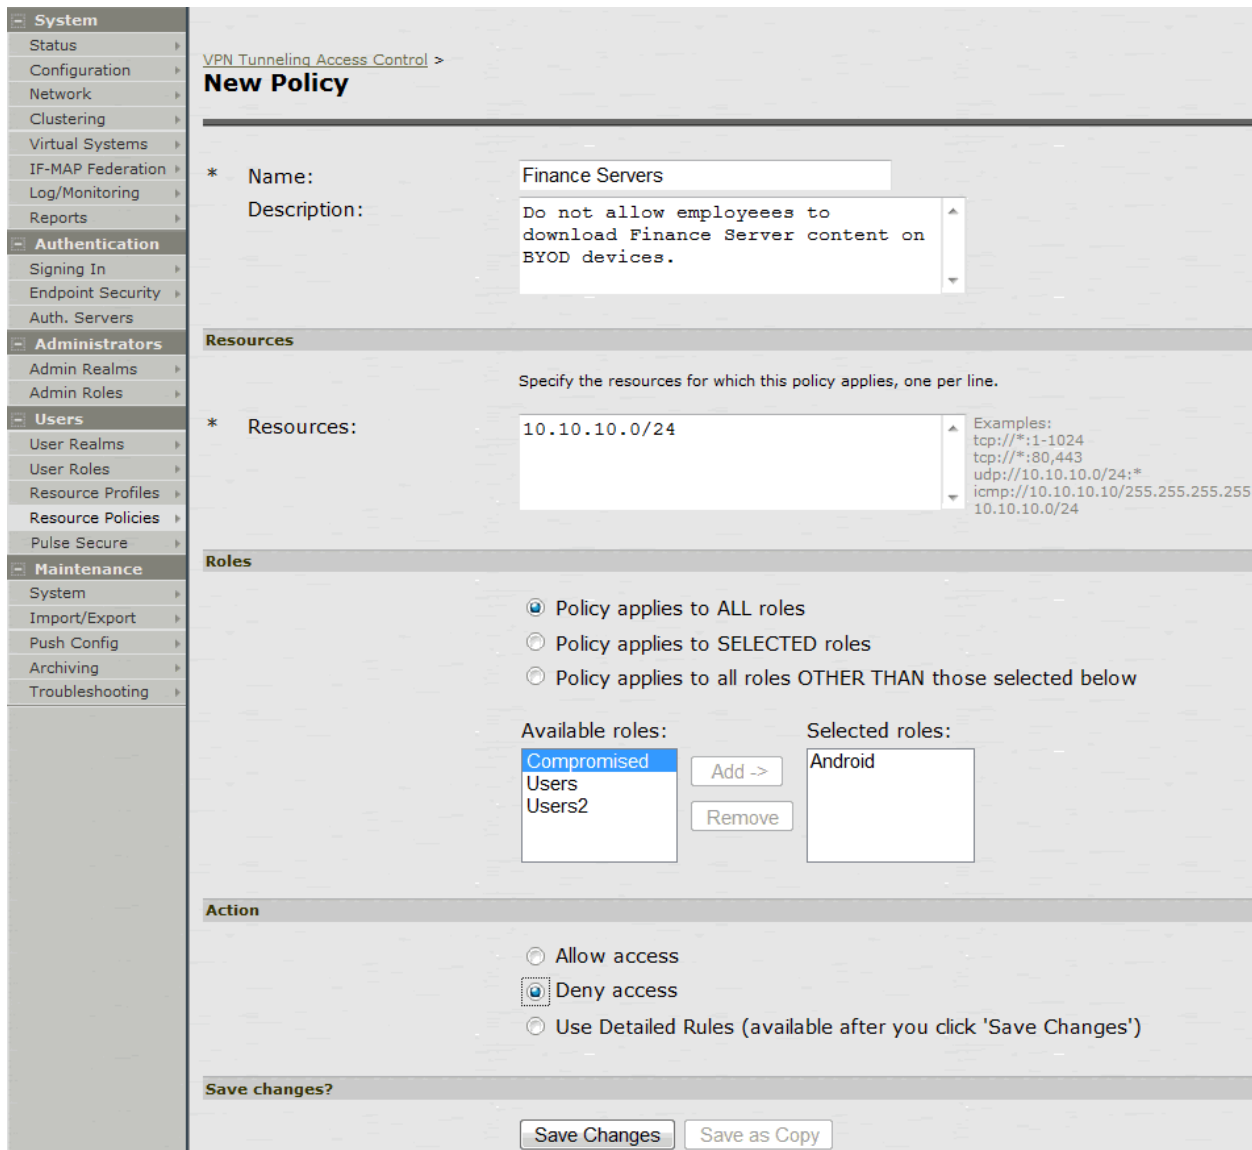
Table72: Resource Access Policy Configuration Guidelines

Settings	Guidelines
Name	Specify a name for the configuration.
Description	Describe the purpose of the configuration so that other administrators are aware of it.
Resources	
Resources	Specify the resources for which this policy applies, one per line.
Roles	
Roles	Select the roles to which the policy applies. In this example, **Android** is selected.
Action	

Action	Select one of the following actions: • **Allow Access** • **Deny Access** • **Use Detailed Rules** In this example, we deny access from Android devices.

Related Documentation

- Using Logs to Verify Proper Configuration

- *User and Policy Administration Overview*

- Using Policy Tracing and Debug Logs

- Understanding the Device Access Management Framework

Using Logs to Verify Proper Configuration

During initial configuration, enable event logs for MDM API calls. You can use these logs to verify proper configuration. After you have verified proper configuration, you can disable logging for these events. Then, enable only for troubleshooting.

To enable logging for MDM API calls:

1. Select **System Log/Monitoring**.

2. Click the **Events** tab.

3. Click the **Settings** tab to display the configuration page.

   Figure 90 shows the configuration page for Policy Secure. The page for Connect Secure is similar.

4. Enable logging for MDM API events and save the configuration.

Figure90: Events Log Settings Configuration Page – Policy Secure

After you have completed the MDM server configuration, you can view system event logs to verify that the polling is occurring.

To display the Events log:

1. Select **System Log/Monitoring**.

2. Click the **Events** tab.

3. Click the **Log** tab.

Figure 91 shows the Events log for Policy Secure.

Figure91: Events Log – Policy Secure

Figure 92 shows the Events log for Connect Secure.

Figure92: Events Log – Connect Secure

Next, to verify user access, you can attempt to connect to a wireless access point with your smart phone, and then view the user access logs.

To display the User Access log:

1. Select **System Log/Monitoring**.

2. Click the **User Access** tab.

3. Click the **Log** tab.

Figure 93 shows the User Access log for Policy Secure. The log for Connect Secure is similar.

Figure93: User Access Log

Related Documentation

- *Deploying a BYOD Policy for AirWatch Managed Devices*

- Deploying a BYOD Policy for AirWatch Managed Devices

- *Deploying a BYOD Policy for MobileIron Managed Devices*

- Deploying a BYOD Policy for MobileIron Managed Devices

Using Policy Tracing and Debug Logs

This topic describes the troubleshooting tools available to diagnose issues. It includes the following information:

- Using Policy Tracing to Troubleshoot Access Issues
- Using the Debug Log

Using Policy Tracing to Troubleshoot Access Issues

It is common to encounter a situation where the system denies a user access to the network or to resources, and the user logs a trouble ticket. You can use the policy tracing utility and log to determine whether the system is working as expected and properly restricting access, or whether the user configuration or policy configuration needs to be updated to enable access in the user's case.

To create a policy trace log:

1. Select **Troubleshooting > User Sessions > Policy Tracing** to display the configuration page.

2. Specify the username, realm, and source IP address if you know it. If you provide the source IP address, the policy trace log can include events that occur before the user ID is entered into the system.

3. Select the events to trace, typically all but Host Enforcer and IF-MAP, unless you have enabled those features.

4. Click **Start Recording**.

5. Initiate the action you want to trace, such as a user sign in.

6. Click **View Log** to display the policy trace results log.

7. Click **Stop Recording** when you have enough information.

Figure 94 shows policy trace results.

Figure94: Policy Tracing Results

Using the Debug Log

The Pulse Secure Global Support Center (PSGSC) might direct you to create a debug log to assist them in helping you debug an issue with the system. The debug log is used only by Pulse Customer Support.

To use debug logging:

1. Select **Troubleshooting > Monitoring > Debug Log** to display the configuration page.

   Figure 95 shows the configuration page for Policy Secure. The page for Connect Secure is similar.

2. Complete the configuration as described in Table 73.

3. Click **Save Changes**. When you save changes with Debug Logging On selected, the system begins generating debug log entries.

4. Initiate the action you want to debug, such as a user sign in. You can reset the debug log file to restart debug logging if it takes you too long to initiate the action.

5. Click **Save Debug Log** to save the debug log to a file that you can send to Pulse Customer Support. You can clear the log after you have saved it to a file.

6. Clear the **Debug Logging On** check box and click **Save Changes** to turn off debug logging.

Figure95: Debug Logging Configuration Page

Table73: Debug Log Configuration Guidelines

Settings	Guidelines
Current Log Size	Displays the size of the current log file. If it is large, use the controls to save, reset, or clear the log file.
Debug Logging On	Select to turn on debug logging.
Debug Log Size	Specify a maximum debug logfile size. The default is 2 MB. The maximum is 250 MB.
Debug Log Detail Level	Specify the debug log detail level. Obtain this from Pulse Customer Support.
Include logs	Select this option to include system logs in the debug log file. Recommended.
Process Names	Specify the process name. Obtain this from Pulse Customer Support.
Event Codes	Specify the event code. Obtain this from Pulse Customer Support. For MDM integration issues, Pulse Customer Support typically likes to collect debugging information for codes **MDM**, **Auth**, **agentman**, and **Realm**. The text is not case sensitive.

Related Documentation

- Using Logs to Verify Proper Configuration

- *Deploying a BYOD Policy for AirWatch Managed Devices*

- Deploying a BYOD Policy for AirWatch Managed Devices

- *Deploying a BYOD Policy for MobileIron Managed Devices*

CHAPTER 11 Authentication Realms

- Understanding Authentication Realms
- Creating an Authentication Realm
- Defining Authentication Access Policies
- Role Mapping Rules
- Specifying Role Mapping Rules for an Authentication Realm
- Machine Authentication for Pulse Secure Connections
- Pulse Secure Connection Realm and Role Preferences for Machine Authentication
- Using the LDAP Server Catalog
- Customizing User Realm UI Views

Understanding Authentication Realms

An authentication realm specifies the conditions that users must meet in order to sign into the system. A realm consists of a grouping of authentication resources, including:

- An authentication server — verifies that the user is who he claims to be. The system forwards credentials that a user submits on a sign-in page to an authentication server.

- A directory server—an LDAP server that provides user and group information to the system that the system uses to map users to one or more user roles.

- An authentication policy—specifies realm security requirements that need to be met before the system submits a user's credentials to an authentication server for verification.

- Role mapping rules—conditions a user must meet in order for the system to map the user to one or more user roles. These conditions are based on either user information returned by the realm's directory server or the user's username.

Authentication realms are an integral part of the access management framework, and therefore are available on all Connect Secure products.

Related Documentation

- About Sign-In Policies
- Defining Authentication Access Policies
- Creating an Authentication Realm

Creating an Authentication Realm

To create an authentication realm:

1. In the admin console, choose **Administrators > Admin Realms** or **Users > User Realms**.

2. On the respective Authentication Realms page, click **New**. Or, select a realm and click Duplicate to base your realm on an existing realm.

3. Enter a name to label this realm and (optionally) a description.

4. If you are copying an existing realm, click **Duplicate**. Then, if you want to modify any of its settings, click the realm's name to enter into edit mode.

5. Select **When editing, start on the Role Mapping page** if you want the Role Mapping tab to be selected when you open the realm for editing.

6. Under Servers, specify:

   - An authentication server to use for authenticating users who sign in to this realm.

   - A directory/attribute server to use for retrieving user attribute and group information for role mapping rules and resource policies. (optional)

   - A RADIUS accounting server to use to track when a user signs in and out of the Infranet Controller (optional).

7. If you want to submit secondary user credentials to an SSO-enabled resource or enable two-factor authentication to access the device, select **Additional authentication server**. Then:

   1. Select the name of the secondary authentication server. Note that you cannot choose an anonymous server, certificate server, or eTrust SiteMinder server.

   2. Select **Username** is specified by user on sign-in page if you want to prompt the user to manually submit his username to the secondary server during the sign-in process. Otherwise, if you want to automatically submit a username to the secondary server, enter static text or a valid variable in the predefined as field. By default, the system submits the <username> session variable, which holds the same username used to sign in to the primary authentication server.

   3. Select **Password** is specified by user on sign-in page if you want to prompt the user to manually submit his password to the secondary server during the sign-in process. Otherwise, if you want to automatically submit a password to the secondary server, enter static text or a valid variable in the predefined as field.

   4. Select **End session if authentication against this server fails** if you want to control access to the system based on the successful authentication of the user's secondary credentials. If selected, authentication fails if the user's secondary credentials fails.

   5. If you want to use dynamic policy evaluation for this realm select **Dynamic policy evaluation** to enable an automatic timer for dynamic policy evaluation of this realm's authentication policy, role mapping rules, and role restrictions. Then:

      1. Use the Refresh interval option to specify how often you want the Infranet Controller to perform an automatic policy evaluation of all currently signedin realm users. Specify the number of minutes (5 to 1440).

      2. Select **Refresh roles** to also refresh the roles of all users in this realm. (This option does not control the scope of the Refresh Now button.)

      3. Select **Refresh resource policies** to also refresh the resource policies (not including Meeting and Email Client) for all users in this realm. (This option does not control the scope of the Refresh Now button.)

      4. Click **Refresh Now** to manually evaluate the realm's authentication policy, role mapping rules, role restrictions, user roles, and resource policies of all currently signed-in realm users. Use this button if you make changes to an authentication policy, role mapping rules, role restrictions, or resource policies and you want to immediately refresh the roles of this realm's users.

      5. Click **Save Changes** to create the realm on the device. The General, Authentication Policy, and Role Mapping tabs for the authentication realm appear.

6. Perform the next configuration steps:

   1. Configure one or more role mapping rules.

   2. Configure an authentication policy for the realm.

Related Documentation

- Defining Authentication Access Policies

- Configuring User Sign In Policies

- Dynamic Policy Evaluation

Defining Authentication Access Policies

An authentication policy is a set of rules that controls one aspect of access management—whether or not to present a realm's sign-in page to a user. An authentication policy is part of an authentication realm's configuration, specifying rules for the system to consider before presenting a sign-in page to a user. If a user meets the requirements specified by the realm's authentication policy, then the system presents the corresponding sign-in page to the user and then forwards the user's credentials to the appropriate authentication server. If this server successfully authenticates the user, then the system moves on to the role evaluation process.

To specify authentication realm access policies:

1. In the admin console, choose **Administrators > Admin Realms or Users > User Realms**.

2. On the respective Authentication Realms page, click Specifying RADIUS Request Attributes a realm and then click the **Authentication Policy** tab.

3. On the Authentication Policy page, configure one or more of the access management options described in the Related Topics section.

Related Documentation

- Specifying Source IP Access Restrictions

- Specifying Password Access Restrictions

- Specifying Certificate Access Restrictions

- Specifying Browser Access Restrictions

- Specifying Session Limits

Role Mapping Rules

Role mapping rules are conditions a user must meet in order for the system to map the user to one or more user roles. These conditions are based on either user information returned by the realm's directory server or the user's username. You must specify role mapping directives in the following format: If the specified condition is|is not true, then map the user to the selected roles.

You create a role mapping rule on Role Mapping tab of an authentication realm. When you click New Rule on this tab, the Role Mapping Rule page appears with an inline editor for defining the rule. This editor leads you through the three steps of creating a rule:

- Specify the type of condition on which to base the rule. Options include:

  - Username

  - User attribute

  - Certificate or certificate attribute

  - Group membership

- Custom expressions

- Specify the condition to evaluate, which consists of:

  - One or more usernames, user attributes, certificate attributes, groups (LDAP), or expressions depending on the type of condition you selected.

  - To what the value(s) should equate, which may include a list of usernames, user attribute values from a RADIUS or LDAP server, client-side certificate values (static or compared to LDAP attributes), LDAP groups, or predefined custom expressions.

- Specify the roles to assign to the authenticated user.

The system compiles a list of eligible roles to which a user may be mapped, which are roles specified by the role mapping rules to which the user conforms. Next, the system evaluates the definition for each role to determine if the user complies with any role restrictions. The system uses this information to compile a list of valid roles, which are roles for which the user meets any additional requirements. Finally, the system either performs a permissive merge of the valid roles or presents a list of valid roles to the user, depending on the configuration specified on the realm's Role Mapping tab.

Related Documentation

- User Roles Overview

Specifying Role Mapping Rules for an Authentication Realm

When creating a new rule that uses LDAP or SiteMinder user attributes, LDAP group information, or custom expressions, you must use the server catalog.

To specify role mapping rules for an authentication realm:

1. In the admin console, choose **Administrators > Admin Realms** or **Users > User Realms**.

2. On the respective Authentication Realms page, select a realm and then click the **Role Mapping** tab.

3. Click **New Rule** to access the Role Mapping Rule page. This page provides an inline editor for defining the rule.

4. In the Rule based on list, choose one of the following:

   - **Username**—Username is the system username entered on the sign-in page. Choose this option if you want to map users to roles based on their system usernames. This type of rule is available for all realms.

   - **User attribute**—User attribute is a user attribute from a RADIUS, LDAP, or SiteMinder server. Choose this option if you want to map users to roles based on an attribute from the corresponding server. This type of rule is available only for realms that use a RADIUS server for the authentication server, or that use an LDAP or SiteMinder server for either the authentication server or directory server. After choosing the User attribute option, click Update to display the Attribute list and the Attributes button. Click the Attributes button to display the server catalog.

     - To add SiteMinder user attributes, enter the SiteMinder user attribute cookie name in the Attribute field in the server catalog, and then click Add Attribute. When you are finished adding cookie names, click OK. The system displays the names of the SiteMinder user attribute cookies in the Attribute list on the Role Mapping Rule page.

     - For information on how to use the server catalog to add LDAP user attributes.

   - **Certificate or Certificate attribute**—Certificate or Certificate attribute is an attribute supported by the users' client-side certificate. Choose this option if you want to map users

to roles based on certificate attributes. The Certificate option is available for all realms; the Certificate attribute option is available only for realms that use LDAP for the authentication or directory server. After choosing this option, click Update to display the Attribute text box.

- **Group membership**—Group membership is group information from an LDAP or native Active Directory server that you add to the server catalog Groups tab. Choose this option if you want to map users to roles based on either LDAP or Active Directory group information. This type of rule is available only for realms that use an LDAP server for either the authentication server or directory server or that use an Active Directory server for authentication. (Note that you cannot specify an Active Directory server as an authorization server for a realm.)

- **Custom Expressions**—Custom Expressions is one or more custom expressions that you define in the server catalog. Choose this option if you want to map users to roles based on custom expressions. This type of rule is available for all realms. After choosing this option, click Update to display the Expressions lists. Click the Expressions button to display the Expressions tab of the server catalog.

  NOTE: If you add more than one custom expression to the same rule, the system creates an "OR" rule for the expressions. For example, you might add the following expressions to a single rule:

  - Expression 1: cacheCleanerStatus = 1

  - Expression 2: loginTime = (8:00AM TO 5:00PM)

  Based on these expressions, a user would match this rule if Cache Cleaner was running on his system OR if he signed into the device between 8:00 and 5:00.

5. Under Rule, specify the condition to evaluate, which corresponds to the type of rule you select and consists of:

   1. Specifying one or more usernames, SiteMinder user attribute cookie names, RADIUS or LDAP user attributes, certificate attributes, LDAP groups, or custom expressions.

   2. Specifying to what the value(s) should equate, which may include a list of usernames, user attribute values from a RADIUS, SiteMinder, or LDAP server, client-side certificate values (static or LDAP attribute values), LDAP groups, or custom expressions.

      For example, you can choose a SiteMinder user attribute cookie named department from the Attribute list, choose is from the operator list, and then enter "sales" and "eng" in the text box.

      Or, you can enter a custom expression rule that references the SiteMinder user attribute cookie named department:

      <userAttr.department = ("sales" and "eng")>

3. Under ...then assign these roles:

   1. Specify the roles to assign to the authenticated user by adding roles to the **Selected Roles list**.

   2. Check **Stop processing rules when this rule matches**if you want to stop evaluating role mapping rules if the user meets the conditions specified for this rule.

3. Click **Save Changes**to create the rule on the Role Mapping tab. When you are finished creating rules:

Make sure to order role mapping rules in the order in which you want to evaluate them. This task is particularly important when you want to stop processing role mapping rules upon a match.

Related Documentation

- Role Mapping Rules

- Policies, Rules & Restrictions, and Conditions Overview

Machine Authentication for Pulse Secure Connections

Pulse Secure client supports machine authentication. Machine authentication uses machine credentials (machine name and password or machine certificate) to authenticate the endpoint. You can enable machine authentication for the system as part of a Pulse Secure Connection and distribute the connection to endpoints through the normal Pulse distribution methods. The following describes the requirements for a machine authentication environment:

- Machine authentication for Connect Secure is available for Pulse layer 3 connections only.

- The authentication server used by the Pulse connection must be Active Directory/Windows NT for machine name/password authentication or a certificate server for machine certificate authentication. You can also use machine credentials when authenticating to RADIUS servers that verify the machine credentials against an Active Directory listing.

- The endpoint must be a member of a Windows domain and the machine credentials must be defined in Active Directory. Typically, during login, the user must enter domain/user in the username box.

- The Pulse connection must be configured so that no prompts are presented during the login process. For example, prompts for realm or role selection or a server certificate trust prompt cause the connection to fail.

- For machine certificate authentication, the domain workstation logon certificate must be issued by the domain certificate authority. The root certificate (CA) must be in the Machine Trusted Certificate store instead of the certificate store for a particular user.

To enable a Pulse connection for machine authentication:

1. Click **Users > Pulse Secure > Connections** and create or select a connection set.

2. Create or edit a connection. Machine authentication is available for connection type **Connect Secure or Policy Secure (L3), Policy Secure (802.1X), or SRX** only.

3. Under Connection is established, select one of the following options:

   - **Automatically when the machine starts. Machine credentials used for authentication**— This option enables machine-only authentication. Machine credentials are used to connect to the system before the user logs on. The user does not need to be logged in. The connection is maintained when a user logs on, logs off, or switches to a different logon.

   - **Automatically when the machine starts. Connection is authenticated again when the user signs in into the desktop**—This option enables user-after-desktop authentication. Machine credentials are used to authenticate the endpoint when no user is logged on. When a user logs on, the machine authentication connection is dropped and the user login is used instead. When the user logs off, the machine connection is reestablished.

Pulse Secure Connection Realm and Role Preferences for Machine Authentication

When a Pulse Secure Connection is configured to use machine authentication, any prompts that occur during the login process cause the connection to fail. For example, if the Pulse server authentication policy allows a user to select a realm or a role during the login process, Pulse presents a dialog box to the user and prompts for the realm or role selection. To avoid failed connections due to prompts during machine authentication you can specify a preferred role and realm for a Pulse connection.

For a Pulse connection that is used for machine authentication, you *do not* need to specify the preferred role if any of the following conditions are true:

- Users are mapped to only one role.

- Users are mapped to more than one role, but the realm's role mapping properties are set to merge settings for all assigned roles.

For a Pulse connection that is used for machine authentication, *you must* specify the preferred realm if the authentication sign-in policy allows the user to select a realm. If that realm maps to only one role, you do not need to specify the role.

For a Pulse connection that is used for machine authentication, *you must* specify the preferred role if any of the following conditions are true:

- The realm that the user connects to maps to more than one role and the realm's role mapping properties are set to require that the user must select a role. The preferred role set must be the name of a role assigned in that realm.

- The realm that the user connects to maps to more than one role and the realm's role mapping properties are defined by role mapping rules. You specify the preferred role by specifying the name of a **rule** that assigns the role set. Figure 96 shows a role mapping rule with the rule name highlighted.

Figure96: Pulse Secure Client Role Mapping Rule

When you create a Pulse connection for machine authentication, you must use the connection type **Connect Secure or Policy Secure (L3), Policy Secure (802.1X), or SRX**. To identify the connection as a machine authentication connection, you specify how the connection is established using one of the following options:

- **Automatically when the machine starts. Machine credentials used for authentication**

  This option uses the machine credentials defined in Active Directory for the machine login process and uses the same credentials for user login. When you select this option, the Realm and Role Set Preferences enable you to specify the following options:

  - **Preferred Machine Realm**—Type the realm name that maps to the role you want to assign.

  - **Preferred Machine Role Set**—Type the name of the role. The role must be one that is identified in the realm's role mapping properties. Or specify the name of a role mapping rule that assigns the role set.

- **Automatically when the machine starts. Connection is authenticated again when the user signs in into the desktop**

  This option uses the Active Directory machine credentials for the machine login process. When machine login is complete, Pulse drops that connection and then uses the user credentials for user login. When you select this option, the Realm and Role Set Preferences enable you to specify the following options:

  - **Preferred Machine Realm**—Type the realm name that maps to the role you want to assign.

  - **Preferred Machine Role Set**—Type the name of the role. The role must be one that is identified in the realm's role mapping properties. Or specify the name of a role mapping rule that assigns the role set.

  - **Preferred User Realm**—Type the realm name that maps to the role you want to assign.

  - **Preferred User Role Set**—Type the name of the role. The role must be one that is identified in the realm's role mapping properties. Or specify the name of a role mapping rule that assigns the role set.

  > NOTE: Realm and role prompts are not the only prompts that are possible during the login process. If the Pulse connection has the Dynamic Certificate Trust option enabled, and there is an issue with the server certificate, Pulse asks the user if it is Ok to proceed. That certificate prompt causes a machine connection to fail. Note that the Pulse prompt for upgrading Pulse software is presented after the user connection is established and it will not affect a machine authentication connection.

Related Documentation

- Role Mapping Rules

- Policies, Rules & Restrictions, and Conditions Overview

Using the LDAP Server Catalog

The LDAP server catalog is a secondary window through which you specify additional LDAP information for the system to use when mapping users to roles, including:

- **Attributes**—The Server Catalog Attributes tab shows a list of common LDAP attributes, such as cn, uid, uniquemember, and memberof. This tab is accessible only when accessing the Server Catalog of an LDAP server. You can use this tab to manage an LDAP server's attributes by adding custom values to and deleting values from its server catalog. Note that the system maintains a local copy of the LDAP server's values; attributes are not added to or

deleted from your LDAP server's dictionary. Figure 97 shows an example of adding the newAccount attribute.

Figure97: Server Catalog > Attributes Tab — Adding an Attribute for LDAP
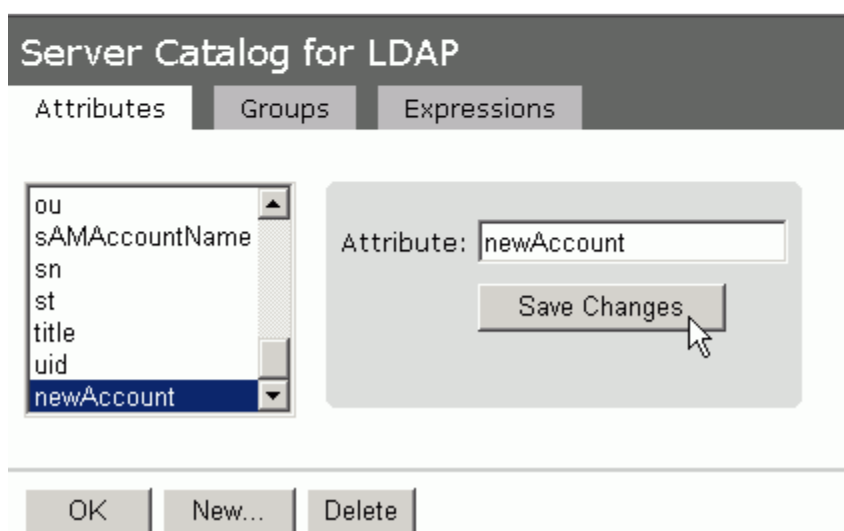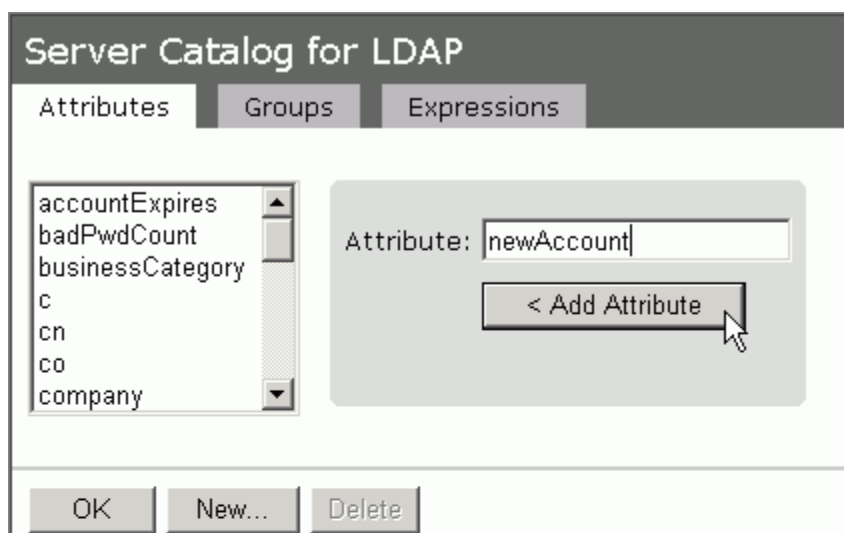
- **Groups**—The Server Catalog Groups tab provides a mechanism to easily retrieve group information from an LDAP server and add it to the server catalog. You specify the BaseDN of your groups and optionally a filter to begin the search. If you do not know the exact container of your groups, you can specify the domain root as the BaseDN, such as dc=juniper, dc=com.The search page returns a list of groups from your server, from which you can choose groups to enter into the Groups list.

  > *i* NOTE: The BaseDN value specified in the LDAP server's configuration page under "Finding user entries" is the default BaseDN value. The Filter value defaults to (cn=*).

  You can also use the Groups tab to specify groups. You must specify the Fully Qualified Distinguished Name (FQDN) of a group, such as cn=GoodManagers, ou=HQ, ou=Juniper, o=com, c=US, but you can assign a label for this group that appears in the Groups list. Note that this tab is accessible only when accessing the Server Catalog of an LDAP server. Figure 98 and Figure 99 show examples of adding LDAP and Active Directory groups.

Figure98: Server Catalog > Groups Tab — Adding LDAP Groups
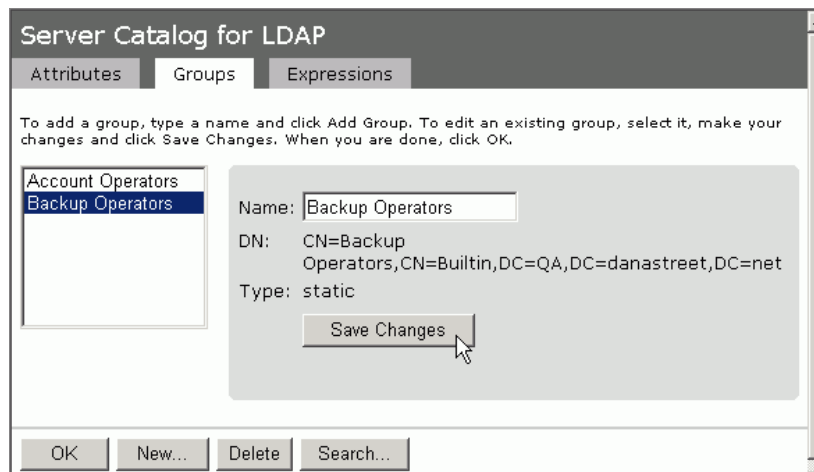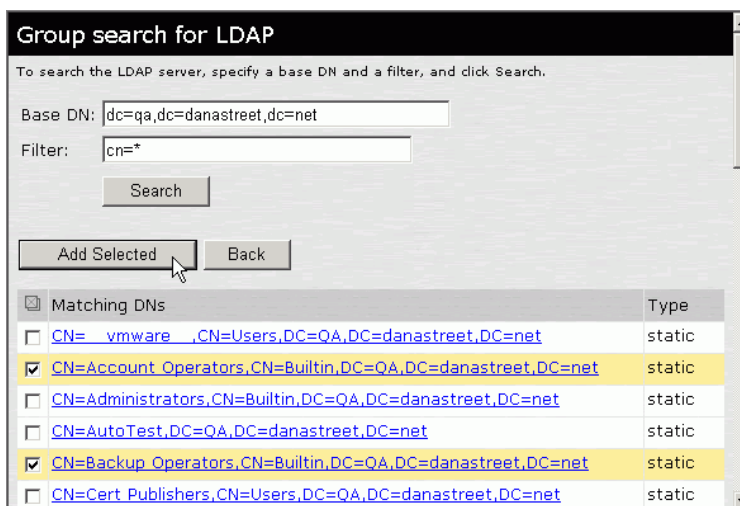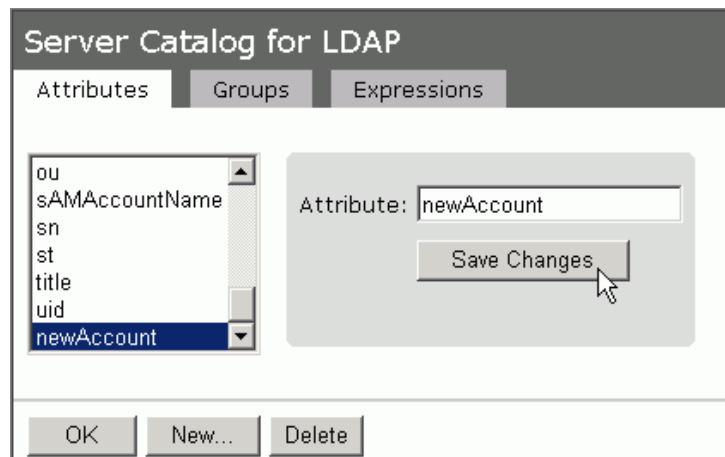
Figure99: Server Catalog > Groups Tab — Adding Active Directory Groups
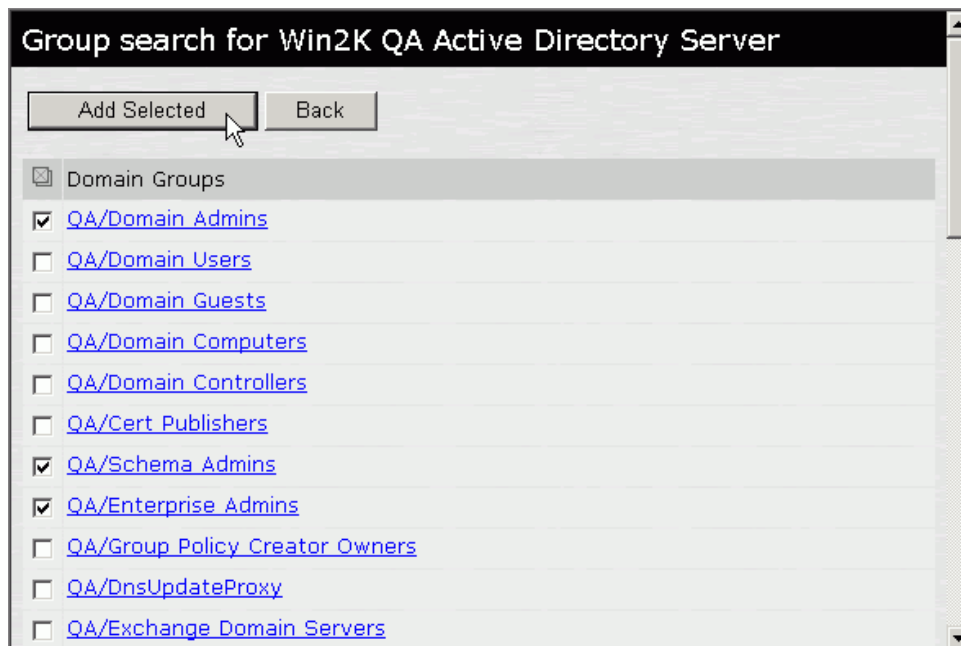
- **Expressions**—The Server Catalog Expressions tab provides a mechanism to write custom expressions for the role mapping rule.

To display the LDAP server catalog:

- After choosing the User attribute option on the Role Mapping Rule page, click **Update** to display the Attribute list and the Attributes button.

- Click the **Attributes** button to display the LDAP server catalog. (You can also click **Groups** after choosing the Group membership option, or click **Expressions** after choosing the Custom Expressions option.)

Related Documentation

- Specifying Role Mapping Rules for an Authentication Realm

- Using Custom Expressions in Rule Configuration

Customizing User Realm UI Views

You can use customization options on the User Authentication Realms page to quickly view the settings that are associated with a specific realm or set of realms. For instance, you can view the role-mapping rules that you have associated with all your user realms. Additionally, you can use these customized views to easily link to the authentication policies, servers, role-mapping rules, and roles associated with a user realms.

To view a sub-set of data on the User Authentication Realms page:

1. Select one of the following options from the View menu:

   - **Overview**—Displays the authentication servers and dynamic policy evaluation settings that you have set for the specified user realms. You may also use this setting to link to the specified server configuration pages.

   - **Authentication Policy**—Displays Host Checker and Cache Cleaner restrictions that you have enabled for the specified user realms. You may also use this setting to link to the specified Host Checker and Cache Cleaner configuration pages.

   - **Role Mapping**—Displays rule conditions and corresponding role assignments that you have enabled for the specified user realms. You may also use this setting to link to the specified rule conditions and role assignments configuration pages.

   - **Servers**—Displays authentication server names and corresponding types that you have enabled for the specified user realms. You may also use this setting to link to the specified server configuration pages.

   - **Roles**—Displays role assignments and corresponding permissive merge settings that you have enabled for the specified user realms.

2. Select one of the following options from the for list:

   - **All realms**—Displays the selected settings for all user realms.

   - **Selected realms**—Displays the selected settings for the user realms you choose. If you select this option, select one or more of the check boxes in the Authentication Realm list.

3. Click **Update**.

CHAPTER 12 Sign-In Policies

About Sign-In Policies

Sign-in policies define the URLs that users and administrators use to access the device and the sign-in pages that they see. The system has two types of sign-in policies—one for users and one for administrators. When configuring sign-in policies, you associate realms, sign-in pages, and URLs.

For example, in order to allow all users to sign in to the device, you must add all user authentication realms to the user sign-in policy. You may also choose to modify the standard URL that the end-users use to access the system and the sign-in page that they see. Or, if you have the proper license, you can create multiple user sign-in policies, enabling different users to sign into different URLs and pages.

Additionally, systems come with a meeting URL. You can use this URL to control the sign-in page that users see when they sign into a meeting on the device. You can also create additional meeting sign-in pages, enabling different Pulse Collaboration users to sign into different URLs and pages.

You can create multiple sign-in policies, associating different sign-in pages with different URLs. When configuring a sign-in policy, you must associate it with a realm or realms. Then, only members of the specified authentication realm(s) may sign in using the URL defined in the policy. Within the sign-in policy, you may also define different sign-in pages to associate with different URLs.

For example, you can create sign-in policies that specify:

- Members of the "Partners" realm can sign in to the device using the URLs: **partner1.yourcompany.com** and **partner2.yourcompany.com**. Users who sign into the first URL see the "partners1" sign-in page; users who sign into the second URL see the "partners2" sign-in page.

- Members of the "Local" and "Remote" realms can sign into the device using the URL: **employees.yourcompany.com**. When they do, they see the "Employees" sign-in page.

- Members of the "Admin Users" realm can sign into the device using the URL: **access.yourcompany.com/super**. When they do, they see the "Administrators" sign-in page.

When defining sign-in policies, you may use different hostnames (such as **partners.yourcompany.com** and **employees.yourcompany.com**) or different paths (such as **yourcompany.com/partners** and **yourcompany.com/employees**) to differentiate between URLs.

NOTE: If a user attempts to sign in while there is another active user session with the same sign-in credentials, the system displays a warning page showing the IP address of the existing session and two buttons: Continue and Cancel. By clicking the Cancel button, the user terminates the current sign-in process and redirects the user back to the Sign-in page. By clicking the Continue button, the system creates the new user session and terminates the existing session.

NOTE: When enabling multiple sign-in URLs, note that in some cases the system must use cookies on the user's machine to determine which sign-in URL and corresponding sign-in page to display to the user. The system creates these cookies when the user signs into the device. (When a user signs into the device, the system responds with a cookie that includes the sign-in domain of the URL. The system then attaches this cookie to every system request the user makes.) Generally, these cookies ensure that the system displays the correct sign-in URL and page to the user. For example, if a user signs into the device using the URL **http://yourcompany.net/employees** and then her session times out, the system uses the cookie to determine that it must display the **http://yourcompany.net/employees** sign-in URL and corresponding page to the user when she requests another system resource.

However, in isolated cases, the cookie on the user's machine may not match the resource she is trying to access. The user may sign into one URL and then try to access a resource that is protected by a different URL. In this case, the system displays the sign-in URL and corresponding sign-in page that the user signed into most recently. For example, a user may sign into the device using the sign-in URL **http://yourcompany.net/employees**. Then she may try to access the system resource using a link on an external server, such as **https://yourcompany.net/partners/dana/term/winlaunchterm.cgi?host=<termsrvIP>**. Or, she may try to open a bookmark that she created during a different session, such as **https://yourcompany.net/partners/,DanaInfo=.awxyBmszGr3xt1r5O3v.,SSO=U+**. In these cases, the system would display the **http://yourcompany.net/employees** sign-in URL and page to the user, rather than the sign-in URL or page that is associated with the external link or saved bookmark that she is trying to access.

Sign-in policies and pages are an integral part of the access management framework, and therefore are available on all Connect Secure products.

Related Documentation

- Task Summary: Configuring Sign In Pages
- Defining Meeting Sign-In Policies
- Configuring User Sign In Policies

Task Summary: Configuring Sign In Pages

To configure sign-in policies, you must:

1. Create an authentication realm through the Administrators > Admin Realms or the Users > User Realms page of the admin console.

2. (Optional) Modify an existing sign-in page or create a new one using options in the Authentication > Signing In > Sign-in Pages page of the admin console.

3. (Optional) Modify an existing sign-in page or create a new one using options in the Authentication > Signing In > Sign-in Pages page of the admin console.

4. Specify a sign-in policy that associates a realm, sign-in URL, and sign-in page using settings in the Authentication > Signing In > Sign-in Policies page of the admin console.

5. If you differentiate between URLs using hostnames, you must associate each hostname with its own certificate or upload a wildcard certificate into the system using options in the System > Configuration > Certificates > Device Certificates page.

Related Documentation

- About Configuring Sign In Policies
- Configuring User Sign In Policies
- Defining Meeting Sign-In Policies

About Configuring Sign In Policies

User sign-in policies also determine the realm(s) that users and administrators can access.

Depending on whether a sign-in policy is for endpoints (users) or administrators, the configuration options are different. For users, different authentication protocol sets can be configured, and realm selection is based on the authentication method that is associated with the realm.

Related Documentation

- Configuring User Sign In Policies
- Defining Meeting Sign-In Policies

Configuring User Sign In Policies

To create or configure user sign-in policies:

1. In the admin console, select **Authentication > Signing In > Sign-in Policies**.

2. To create a new sign-in policy, click **New URL**. Or, to edit an existing policy, click a URL in the **Administrator URLs** or **User URLs** column.

3. Select Users or Administrators to specify which type of user can sign in using the access policy.

4. In the **Sign-in URL** field, enter the URL that you want to associate with the policy. Use the format <host>/<path> where <host> is the hostname of the device, and <path> is any string you want users to enter. For example: partner1.yourcompany.com/outside. To specify multiple hosts, use the * wildcard character.

   To specify that all administrator URLs should use the sign-in page, enter */admin.

   > **NOTE:**
   > - You may only use wildcard characters (*) in the beginning of the hostname portion of the URL. The system does not recognize wildcards in the URL path.
   > - SAML authentication does not support sign-in URLs that contain multiple realms. Instead, map each sign-in URL to a single realm.

5. (optional) Enter a Description for the policy.

6. From the Sign-in Page list, select the sign-in page that you want to associate with the policy. You may select the default page that comes with the system, a variation of the standard sign-in page, or a custom page that you create using the customizable UI feature.

7. (User URLs only) In the Meeting URL field, select the meeting URL that you want to associate with this sign-in policy. The system applies the specified meeting URL to any meeting created by a user who signs into this user URL.

8. Under Authentication realm, specify which realm(s) map to the policy, and how users and administrators should pick from amongst realms. If you select:

   - User types the realm name—The system maps the sign-in policy to all authentication realms, but does not provide a list of realms from which the user or administrator can choose. Instead, the user or administrator must manually enter his realm name into the sign-in page.

   - User picks from a list of authentication realms—The system only maps the sign-in policy to the authentication realms that you choose. The system presents this list of realms to the user or administrator when he signs-in to a device and allows him to choose a realm from the list. (Note that the system does not display a drop-down list of authentication realms if the URL is only mapped to one realm. Instead, it automatically uses the realm you specify.)

     > **NOTE:** If you allow the user to pick from multiple realms and one of those realms uses an anonymous authentication server, the system does not display that realm in the drop-down realm list. To effectively map your sign-in policy to an anonymous realm, you must add only that realm to the Authentication realm list.

9. Click **Save Changes**.

Enabling and Disabling Sign-In Policies

To enable and disable sign-in policies:

1. In the admin console, choose **Authentication > Signing In > Sign-in Policies**.

2. To enable or disable:

- **An individual policy**—Select the check box next to the policy that you want to change, and then click Enable or Disable.

- **All user and meeting policies**—Select or deselect the **Restrict access to administrators only** check box at the top of the page.

  If you select this option, all user sessions are immediately terminated. If this device is part of a cluster, all user sessions across all nodes in the cluster are immediately terminated.

3. Click **Save Changes**.

Specifying the Order in Which Sign-In Policies are Evaluated

The system evaluates sign-in policies in the same order that you list them on the Sign-in Policies page. When it finds a URL that matches exactly, it stops evaluating and presents the appropriate sign-in page to the administrator or user. For example, you may define two administrator sign-in policies with two different URLs:

- The first policy uses the URL */admin and maps to the default administrator sign-in page.

- The second policy uses the URL yourcompany.com/admin and maps to a custom administrator sign-in page.

If you list the policies in this order on the Sign-in Policies page, the system never evaluates or uses the second policy because the first URL encompasses the second. Even if an administrator signs in using the yourcompany.com/admin URL, the system displays the default administrator sign-in page. If you list the policies in the opposite order, however, the system displays the custom administrator sign-in page to those administrators who access the system using the yourcompany.com/admin URL.

Note that the system only accepts wildcard characters in the hostname section of the URL and matches URLs based on the exact path. For example, you may define two administrator sign-in policies with two different URL paths:

- The first policy uses the URL */marketing and maps to a custom sign-in page for the entire Marketing Department.

- The second policy uses the URL */marketing/joe and maps to a custom sign-in page designed exclusively for Joe in the Marketing Department.

If you list the policies in this order on the Sign-in Policies page, the system displays Joe's custom sign-in page to him when he uses the yourcompany.com/marketing/joe URL to access the device. He does not see the Marketing sign-in page, even though it is listed and evaluated first, because the path portion of his URL does not exactly match the URL defined in the first policy.

To change the order in which administrator sign-in policies are evaluated:

1. In the admin console, choose **Authentication > Signing In > Sign-in Policies**.

2. Select a sign-in policy in the **Administrator URLs** , **User URLs** or **Meeting URLs** list.

3. Click the up and down arrows to change the selected policy's placement in the list.

4. Click **Save Changes**.

Related Documentation

- About Configuring Sign In Policies

About Sign-In Notifications

With sign-in notifications, you can create and configure detailed notification messages that appear for Pulse clients and for agentless access endpoints when the user attempts to sign in. For example, you could configure a notification message that explains terms of use, company-specific policies, a welcome page, an end user license agreement (EULA), or a message of the day (MOTD).

For a browser-based (agentless) login, the notification message appears in a separate page either before (pre-auth) or after (post-auth) user authentication during the sign-in process. For a Pulse client login, the notification messages appear in a Pulse message box. The user is expected to read the content of the sign-in notification message and acknowledge by clicking a Proceed button. The user may indicate disagreement by clicking a Decline button, which ends the login attempt.

You can configure a sign-in policy to use a sign-in notification either as pre-auth or post-auth (or both). In the case of post-auth configuration, you can either use a common message for all roles or use separate messages for each role.

You can create a multi-language sign-in notification package that relies on the language setting of the endpoint. You can customize the sign-in notification page appearance for browser-based logins by modifying the related fields in a sign-in page in the Admin UI or by using a custom sign-in page.

Notes:

- Sign-in notifications are supported on Windows, Mac, and for browser-based access on mobile devices. However, sign-in notifications might not work well with all mobile devices due to device limitations.

- Sign-in notifications (including uploaded packages) are included in XML exports.

- If a Pulse session is resumed or extended, the pre-auth notification message is not shown again. However, if the user switches roles when resuming a session, and that role change results in a new notification, Pulse displays the message. You can configure the post-auth message to be skipped if it has already been seen. If the post-auth message is not marked to be skipped, then it always appears.

Related Documentation

- Configuring and Implementing Sign-in Notifications

Configuring and Implementing Sign-in Notifications

Sign-in notifications appear for Pulse client and for browser-based logins when the user attempts to sign in.

To configure and implement sign-in notifications:

1. In the admin console, select **Authentication > Signing In > Sign-in Notifications**.

2. Click **New Notification**.

3. Specify a Name for the notification. This name appears in the sign-in policies page, and in the UI Options page for a selected role.

4. Select **Text** or **Package** in the Type box.

   • If you select **Text**, type the desired sign-in notification message, or copy and paste the relevant text into the Text field.

   • If you select **Package**, click the **Browse** button and navigate to a previously prepared .zip file. A package is typically used to provide different language versions of the notification message.

      • The zip file should include a default.txt file and one or more <language>.txt files (Example: en.txt).

      • Language-abbreviations should be strings that can appear in Accept-Language header of an HTTP request.

      • The character encoding supported is UTF-8.

         ⓘ   NOTE: When you create a zip file, do not add the folder containing the files, but add the files directly.

5. Click **Save Changes**.

To enable sign-in notifications:

1. In the admin console, click **Authentication > Signing In > Sign-in Policies**.

2. Select an existing URL or create a new URL.

3. Under Configure Sign-in Notifications, select the check box for **Pre-Auth Sign-in Notification**, **Post-Auth Sign-in Notification**, or both.

   • After Pre-Auth Sign-in Notification, select a previously configured sign-in notification from the drop-down menu.

   • After Post-Auth Sign-in Notification, select the option for **Use a common Sign-in Notification for all roles** or **Use the Sign-in Notification associated to the assigned role**.

   • If you select **Use a common Sign-in Notification for all roles**, select a previously configured sign-in notification from the drop-down menu.

   • If you select **Use the Sign-in Notification associated to the assigned role**, the sign-in notification configured for the assigned role will be used.

   • Prevent the Post-Auth sign-in notification from being displayed to users who have seen it before, by selecting the **Skip if already shown** check box. (This is only a hint to the system and might not be honored in all environments.)

4. Click **Save Changes**.

5. You can customize the appearance of the sign-in notification message by selecting **Authentication > Signing In > Sign-in Pages** and creating a sign-in page or using an existing page.

6. Under Sign-in Notification appearance, customize UI options for Pre-Auth Notifications and Post-Auth Notifications by changing the following items:

- For **Notification Title** enter the text that appears at the top of the sign-in notification page.

- In the **Proceed Button** box, enter the text for the button that the user clicks to proceed with the sign-in.

  This text applies to browser-based logins only. A Pulse client login always displays Proceed.

- Optionally, clear the check box for **Display "Decline" Button**. If this box is not checked, the user does not have the option to decline.

- In the **Decline Button** box, enter the text for the button that the user clicks to decline.

  This text applies to browser-based logins only. A Pulse client login always displays Decline.

- In the **Message on Decline** box, enter the text that you would like to appear when a user clicks the Decline button.

7. Click **Save Changes**.

> ℹ️ NOTE: If you enabled **Use the Sign-in Notification associated to the assigned role** you must complete the implementation by selecting the sign-in notification on the Users > User Roles > *Role Name* > General > UI Options page or Administrators > Admin Roles > *Role Name* > General > UI Options page, as applicable.
>
> If more than one role is available to a user, the sign-in notification associated with the first role assigned is displayed.

8. Add the sign-in page in which you have customized the sign-in notification appearance to the sign-in policy.

Related Documentation

- About Sign-In Notifications

Defining Authorization-Only Access Policies

Authorization-only access is similar to a reverse proxy. Typically, a reverse proxy is a proxy server that is installed in front of webservers. All connections coming from the Internet addressed to one of the webservers are routed through the proxy server, which may either deal with the request itself or pass the request wholly or partially to the main webserver.

With an authorization-only access, you select a user role. the system then acts as reverse proxy server and performs authorization against the Netegrity SiteMinder server for each request.

For example, the authorization-only access feature satisfies the following business needs:

- If you have a third-party AAA policy management server (like Netegrity), the system acts as an authorization-only agent.

- If your user sessions are managed by a third-part session management system, there is no need to duplicate the user session management in the system.

With authorization-only access, there is no SSO from the system. SSO is controlled by your third-party AAA infrastructure.

> NOTE: Before defining this policy, you must first configure your Netegrity server and define your hostnames in the Network Configuration page.
>
> You must also specify settings in the SiteMinder authorization settings section of the SiteMinder authentication server page. Users are redirected to the URL specified in the If Automatic Sign In fails, redirect to field when the SMSESSION cookie validation fails or if no SMSESSION cookie exists. Users are redirected to the URL specified in the If authorization fails, redirect to field when an access denied error occurs.

To create or configure authorization-only access policies:

1. In the admin console, choose **Authentication > Signing In > Sign-in Policies**.

2. To create a new authorization only access policy, click **New URL** and select **authorization only access**. Or, to edit an existing policy, click a URL in the Virtual Hostname column.

3. In the Virtual Hostname field, enter the name that maps to the system's IP address. The name must be unique among all virtual hostnames used in pass-through proxy's hostname mode. The hostname is used to access backend application entered in the Backend URL field. Do not include the protocol (for example, http:) in this field.

   For example, if the virtual hostname is myapp.ivehostname.com, and the backend URL is http://www.xyz.com:8080/, a request to https://myapp.ivehostname.com/test1 via the system is converted to a request to http://www.xyz.com:8080/test1. The response of the converted request is sent to the original requesting web browser.

4. In the Backend URL field, enter the URL for the remote server. You must specify the protocol, hostname and port of the server. For example, http://www.mydomain.com:8080/*.

   When requests match the hostname in the Virtual Hostname field, the request is transformed to the URL specified in the Backend URL field. The client is directed to the backend URL unaware of the redirect.

5. (optional) Enter a Description for this policy.

6. Select the server name or **No Authorization** from the Authorization Server drop-down menu. If you select a server, ensure that the front-end server provides the SMSESSION cookie otherwise you will receive an error.

7. Select a user role from the Role Option drop-down menu.

   Only the following user role options are applicable for authorization-only access.

   - Allow browsing un-trusted SSL web sites (Users > User Roles > *RoleName* > Web > Options > View advanced options)

   - HTTP Connection Timeout (Users > User Roles > *RoleName* > Web > Options > View advanced options)

- Source IP restrictions (Users > User Roles > *RoleName* > General > Restrictions)

- Browser restrictions (Users > User Roles > *RoleName* > General > Restrictions)

  Ensure the user role you select has an associated Web Access policy.

8. Select the **Allow ActiveSync Traffic only** option to perform a basic of validation of the HTTP header to ensure the request is consistent with ActiveSync protocol. If you select this option only ActiveSync protocol requests can be processed. If validation fails, a message is created in the user's event log. If you do not select this option, both ActiveSync and non-ActiveSync requests are processed.

9. Click **Save Changes** to save your edits.

The System Status Overview page displays the number of current active concurrent connections and a histogram of the active concurrent connections (Authorization Only Access Active Connections plot in the Concurrent SSL Connections graph).

Related Documentation

- Using a SiteMinder Server

- Specifying Web Browsing Options

- Specifying Browser Access Restrictions

Defining Meeting Sign-In Policies

To create or configure meeting sign-in policies:

1. In the admin console, choose **Authentication > Authentication > Signing In Policies**.

2. To create a new sign-in policy, click **New URL**. Or, to edit an existing policy, click a URL in the Meeting URLs column.

3. Select **Meeting**.

4. In the Sign-in URL field, enter the URL that you want to associate with the meeting policy. Use the format <host>/<path> where <host> is the hostname of the device, and <path> is any string you want users to enter. For example: Partner1.YourCompany.com/OnlineConference. When creating the meeting URL, note that:

   - You cannot modify the URL of the default meeting URL (*/meeting) that comes with the product.

   - If you want to enable users to sign into meetings using all of the hostnames defined in the associated user URL, use the * wildcard character in your meeting URL definition. For example, you might associate the following hosts with your user URL:

     - YourInternalServer.YourCompany.net

     - YourExternalServer.YourCompany.com

     Then, if you create an */OnlineConference meeting URL definition and associate it with the user URL, users can access the meeting sign-in page using either of the following URLs:

     - http://YourInternalServer.YourCompany.net/OnlineConference

- http://YourExternalServer.YourCompany.com/OnlineConference

- If you create a meeting URL that includes the * wildcard character and enable e-mail notifications, the system constructs the meeting URL in the notification e-mail using the hostname specified by the user when signing into the device. For instance, a user might sign into the device using the following URL from the previous example:

  http://YourInternalServer.YourCompany.net

  Then, if the user creates a meeting, the system specifies the following sign-in URL for that meeting in the e-mail notification:

  http://YourInternalServer.YourCompany.net/OnlineConference

  Note that since the e-mail link references an internal server, out-of-network users cannot access the meeting.

- If you only want to enable users to sign into meetings using a sub-set of the hostnames defined in the associated user URL, or if you want to require users to use a completely different URL to sign into meetings, do not include the * wildcard character in your meeting URL definition. Instead, create a unique and specific meeting URL definition.

  For instance, you can create the following meeting URL definition and associate it with the user URL from the previous example in order to specify that all meetings contain links to the external server only:

  YourExternalServer.YourCompany.com/OnlineConference

5. (optional) Enter a Description for the policy.

6. From the Sign-in Page list, select the sign-in page(s) that you want to appear to users who access meetings using this policy. You may select the default pages that come with the system, a variation of the standard sign-in pages, or customized pages that you create using the customizable UI feature.

7. Click **Save Changes**.

Related Documentation

- Configuring Sign-In Pages

Configuring Sign-In Pages

A *sign-in page* defines the customized properties in the end-user's welcome page such as the welcome text, help text, logo, header, and footer. The system allows you to create two types of sign-in pages to present to users and administrators:

- Standard sign-in pages—Standard sign-in pages are produced by Pulse Secure and are included with all versions of the Connect Secure software. You can modify standard sign-in pages through the Authentication > Signing In > Sign-in Pages tab of the admin console.

- Customized sign-in pages—Customized sign-in pages are THTML pages that you produce using the Template Toolkit and upload to the system in the form of an archived ZIP file. The customized sign-in pages feature enables you to use your own pages rather than having to modify the sign-in page included with the system.

Configuring Standard Sign-In Pages

Standard sign-in pages that come with the system include:

- Default Sign-In Page—the system displays this page to users when they sign into the device.

- Meeting Sign-In Page—the system displays this page to users when they sign into a meeting.

You can modify the default sign-in page that the system displays to users when they sign into the device. You can also create new standard sign-in pages that contain custom text, logo, colors, and error message text using settings in the Authentication > Signing In > Sign-in Pages tab of the admin console.

To create or modify a standard sign-in page:

1. In the admin console, select **Authentication > Signing In > Sign-in Pages**.

2. If you are:

   - **Creating a new page**—**Click New Page**.

   - **Modifying an existing page**—Select the link corresponding to the page you want to modify.

3. (New pages only) Under Page Type, specify whether this is an administrator/user access page or a meeting page.

4. Enter a name to identify the page.

5. In the Custom text section, revise the default text used for the various screen labels as desired. When adding text to the Instructions field, note that you may format text and add links using the following HTML tags: <i>, <b>, <br>, <font>, and <a href>. However, the system does not rewrite links on the sign-in page (since the user has not yet authenticated), so you should only point to external sites. Links to sites behind a firewall will fail.

   If you use unsupported HTML tags in your custom message, the system may display the end-user's home page incorrectly.

6. In the Header appearance section, specify a custom logo image file for the header and a different header color.

7. In the Custom error messages section, revise the default text that is displayed to users if they encounter certificate errors.

   You can include <<host>>, <<port>>, <<protocol>>, and <<request>> variables and user attribute variables, such as <<userAttr.cn>> in the custom error messages. Note that these variables must follow the format <*variable*> to distinguish them from HTML tags which have the format <*tag*>.

8. To provide custom help or additional instructions for your users, select **Show Help button**, enter a label to display on the button, and specify an HTML file to upload to the system. Note that the system does not display images and other content referenced in this HTML page. (Not available for the Pulse Collaboration sign-in page.)

9. Click **Save Changes**. The changes take effect immediately, but users with active sessions might need to refresh their Web browsers.

Click **Restore Factory Defaults** to reset the sign-in page, the user home page, and admin console appearance.

Related Documentation

- *Custom Text HTML Tags*

CHAPTER 13 Single Sign-On

About Single Sign-On

Single sign-on (SSO) is a process that allows pre-authenticated Connect Secure users to access other applications or resources that are protected by another access management system without having to re-enter their credentials.

The system provides several integration mechanisms that allow you to configure SSO connections from the system to other servers, applications, and resources. SSO mechanisms include:

- Remote SSO—The system provides loose integration with any application that uses a static POST action within an HTML form to sign in users. You can configure the system to post system credentials, LDAP attributes, and certificate attributes to a Web-enabled application, as well as set cookies and headers, allowing users to access the application without re-authenticating.

- SAML—The system provides loose integration with selected access management systems that use the Security Assertion Markup Language (SAML) to communicate with other systems. You can enable users to sign in to the system and then sign in to and access resources protected by the access management system without re-authenticating. You can also enable users to sign in to another access management system and then access resources protected by the system, without re-authenticating.

- Basic authentication and NTLM intermediation to Intranet sites—The system allows you to automatically submit user credentials to other web sites and proxies within the same Intranet zone. When you enable basic authentication intermediation through the Users > Resource Profiles > Web Applications/Pages page of the admin console, the system submits the cached credentials to Intranet web sites whose hostnames end in the DNS suffix configured in the System > Network > Overview page. To maximize security, you may also configure the system to use base-64 encoding to protect the cached credentials.

- Active Directory server—The system allows you to automatically submit Active Directory SSO credentials to other web sites and Windows file shares within the same Intranet zone that are protected by native NTLM authentication. When you enable this option, the system submits cached credentials to NTLM-protected web sites whose hostnames end in the DNS suffix configured in the System > Network > Overview page of the admin console.

- eTrust SiteMinder policy server—When you authenticate system users using a eTrust SiteMinder policy server, you can enable them access to SiteMinder protected resources without re-authenticating (provided they are authorized with the correct protection level). Additionally, you can re-authenticate users through the system if they request resources for which their current protection level is inadequate and you can enable users to sign into the policy server first and then access the system without re-authenticating.

- Terminal Sessions—When you enable the Terminal Services feature for a role, you allow users to connect to applications that are running on a Windows terminal server or Citrix MetaFrame server without re-authenticating. You may also pass a username to the Telnet/SSH server.

- Email clients—When you enable the Email Client feature for a role and then create a corresponding resource policy, you allow users to access standards-based e-mail such as Outlook Express, Netscape Communicator, or Qualcomm's Eudora without re-authenticating.

The system determines which credentials to submit to the SSO-enabled server, application, or resource based on the mechanism you use to connect. Most mechanisms allow you to collect user credentials for up to two authentication servers in the system sign-in page and then submit those credentials during SSO.

The remaining mechanisms (SAML, eTrust SiteMinder, and the Email Client) use unique methods for enabling SSO from Connect Secure to the supported application.

About Multiple Sign-In Credentials

When configuring an authentication realm, you can enable up to two authentication servers for the realm. Enabling two authentication servers allows you to require two different sets of credentials—one for Connect Secure and another for your SSO-enabled resource—without requiring the user to enter the second set of credentials when accessing the resource. It also allows you to require two-factor authentication in order to access the device.

Related Documentation

- Remote SSO Overview

- Understanding SAML 1.1

- Defining a Single Sign-On Autopolicy

- *Defining an Active Directory or Windows NT Domain Server Instance*

- Using a SiteMinder Server

- About Terminal Services

- About the E-Mail Client

- Multiple Sign-In Credentials Execution

Task Summary: Configuring Multiple Authentication Servers

To enable multiple authentication servers:

1. Create authentication server instances through the Authentication > Auth. Servers page of the admin console.

2. Associate the authentication servers with a realm using settings in the following pages of the admin console:

- Users > User Realms > *Select Realm*> General

- Administrators > Admin Realms > *Select Realm* > General

3. (Optional) Specify password length restrictions for the secondary authentication server using settings in the following pages of the admin console:

- Users > User Realms > *Select Realm* > Authentication Policy > Password

- Administrators > Admin Realms > *Select Realm* > Authentication Policy > Password

Related Documentation

- Understanding Authentication Realms

- Specifying Password Access Restrictions

Task Summary: Enabling SSO to Resources Protected by Basic Authentication

To enable single sign-on to Web servers and Web proxies that are protected by basic authentication, you must:

1. Specify a hostname that ends with the same prefix as your protected resource using settings in the System > Network > Overview page of the admin console. (The system checks the hostnames to ensure that it is only enabling SSO to sites within the same Intranet.)

2. Enable users to access Web resources, specify the sites to which you want the system to submit credentials, create autopolicies that enable basic authentication intermediation single sign-on, and create bookmarks to the selected resources using settings in the Users > Resource Profiles > Web Application/Pages > [Profile] page of the admin console.

3. If you want users to access Web servers through a proxy, configure the system to recognize the appropriate servers and proxies using settings in the following pages of the admin console:

- Use settings in Users > Resource Policies > Web > Web proxy > Servers page to specify which Web servers you want to protect with the proxy.

- Use settings in the Users > Resource Policies > Web > Web proxy > Policies page to specify which proxies you want to use and which servers (above) you want the proxies to protect. You may specify individual resources on the server or the entire server.

Related Documentation

- Task Summary: Enabling SSO to Resources Protected by NTLM

Task Summary: Enabling SSO to Resources Protected by NTLM

i NOTE: The system supports web proxies that perform NTLM authentication. However, the following case is not supported: a proxy exists between the system and the back-end server and the back-end server performs the NTLM authentication.

To enable single sign-on to Web servers, Windows file servers, and Web proxies that are protected by NTLM, you must:

1. Specify a hostname that ends with the same suffix as your protected resource using settings in the System > Network > Overview page of the admin console. (The system checks the hostnames to ensure that it is only enabling SSO to sites within the same Intranet.)

2. Enable users to access the appropriate type of resource (Web or file), specify the sites or servers to which you want the system to submit credentials, create autopolicies that enable NTLM single sign-on, and create bookmarks to the selected resources using settings in the following pages of the admin console:

   - Users > Resource Profiles > Web Application/Pages > [*Profile*]

   - Users > Resource Profiles > File Browsing Resource Profiles> [*Profile*]

3. If you want users to access Web servers through a proxy, configure the system to recognize the appropriate servers and proxies using settings in the following pages of the admin console:

   1. Use settings in Users > Resource Policies > Web > Web proxy > Servers page to specify which Web servers you want to protect with the proxy.

   2. Use settings in the Users > Resource Policies > Web > Web proxy > Policies page to specify which proxies you want to use and which servers (above) you want the proxies to protect. You may specify individual resources on the server or the entire server.

Related Documentation

   - Task Summary: Enabling SSO to Resources Protected by Basic Authentication

Multiple Sign-In Credentials Execution

The following diagram illustrates the process that the system uses to collect and authenticate multiple user credentials and submit them to SSO-enabled resources. Each of the steps in the diagram are described in further detail in the sections that follow.

Figure100: Collecting and Submitting Credentials from Multiple Servers
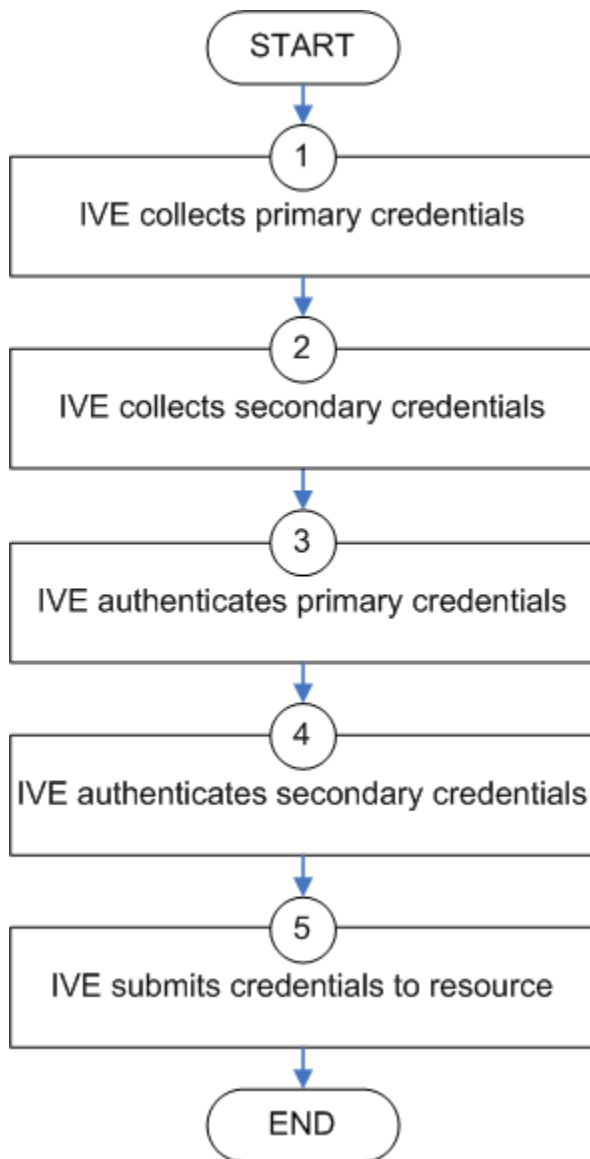
Step 1: Collect the User's Primary Credentials

When the user signs in to a device, the system prompts him to enter his primary server credentials. The system saves these credentials to submit to the SSO resource later, if necessary. Note that the system saves the credentials exactly as the user enters them—it does not pre-pend or append them with additional information such as the user's domain.

Step 2: Collect or Generate the User's Secondary Credentials

You may configure the system to either manually collect or automatically generate the user's secondary set of credentials. If you configure the system to:

- Manually collect the user's secondary credentials—The user must enter his secondary credentials directly after entering his primary credentials.

- Automatically generate the user's credentials—The system submits the values you specified in the administration console during setup. By default, the system uses the <username> and <password> variables, which hold the username and passwod entered by the user for the primary authentication server.

For example, you may configure an LDAP server as your primary authentication server and an Active Directory server as your secondary authentication server. Then, you may configure the system to infer the user's Active Directory username but require the user to manually enter his Active Directory password. When the system infers the Active Directory username, it simply takes the name entered for the LDAP server (for example, JDoe@LDAPServer) and resubmits it to the Active Directory (for example, JDoe@ActiveDirectoryServer).

Step 3: Authenticate the Primary Credentials

After the system collects all required credentials, it authenticates the user's first set of credentials against the primary authentication server. Then:

- If the credentials successfully authenticate, the system stores them in the <username> and <password> session variables and continues on to authenticate the secondary credentials.

  NOTE: If you authenticate against a RADIUS server that accepts dynamic, time-sensitive passwords, you may choose to not store user passwords using the system session variable.

- If the credentials do not successfully authenticate, the system denies the user access to the device.

Step 4: Authenticate the Secondary Credentials

After authenticating the primary credentials, the system authenticates the secondary credentials. Then:

- If the credentials successfully authenticate, the system stores them in the <username[2]> and <password[2]> session variables and allows the user access to the device. You may also access these variables using the syntax <username@SecondaryServer> and <password@SecondaryServer>.

  NOTE: If you authenticate against a RADIUS server that accepts dynamic, time-sensitive passwords, you may choose to not store user passwords using the system session variable.

- If the credentials do not successfully authenticate, the system does not save them. Depending on how you configure your authentication realm, the system may allow or deny the user access to a device if his secondary credentials do not successfully authenticate.

  You can detect that secondary authentication failed by creating a custom expression that checks for an empty user@secondaryAuth variable. You may want to do this so that you can assign users to roles based on successful authentication. For example, the following expression assigns users to the "MoreAccess" role if they successfully authenticate against the "ACE server" secondary authentication server:

user@{ACE Server} != "" then assign role MoreAccess

Note "Ace server" is shown in curly braces since the authentication server's name contains spaces.

Step 5: Submit Credentials to an SSO-Enabled Resource

After the user successfully signs in to a device, he may try to access an SSO-enabled resource using a pre-configured bookmark or other access mechanism. Then, depending on which type of resource the user is trying to access, the system submits different credentials. If the user is trying to access a:

- Web SSO, Terminal Services, or Telnet/SSH resource—The system submits the credentials that you specify through the admin console, such as <username> (which submits the user's primary credentials to the resource) or <username[2]> (which submits the user's secondary credentials to the resource). Or, if the user has entered a different username and password through the end user console, the system submits the user-specified credentials.

    NOTE: The system does not support submitting ACE server, certificate server, or anonymous server credentials to a Web SSO, terminal services, or Telnet/SSH resource. If you configure the system to submit credentials from one of these types of primary authentication servers, it submits credentials from the user's secondary authentication server instead. If these credentials fail, the system prompts the user to manually enter his username and password.

- Resource protected by a Web server, Windows server, or Web proxy that is using NTLM authentication—The system submits credentials to the backend server or proxy that is protecting the Web or file resource. Note that you cannot disable NTLM authentication through the system—If a user tries to access a resource that is protected by NTLM, the system automatically intermediates the authentication challenge and submits credentials in the following order:

    - (Windows file resources only) Administrator-specified credentials—If you create a resource profile that specifies credentials for a Windows file resource and the user then accesses the specified resource, the system submits the specified credentials.

    - Cached credentials—If the system does not submit administrator-specified credentials or the credentials fail, it determines whether it has stored credentials for the specified user and resource in its cache. (See below for information about when the system caches credentials.) If available, the system submits its stored credentials.

    - Primary credentials—If the system does not submit cached credentials or the credentials fail, it submits the user's primary system credentials provided that following conditions are true:

        - The resource is in the same Intranet zone as the device (that is, the resource's hostname ends in the DNS suffix configured in the System > Network > Overview page of the admin console).

        - (Web proxies only) You have configured the system to recognize the Web proxy through settings in the Users > Resource Policies > Web > Web Proxy pages of the admin console.

        - The credentials are not ACE credentials.

- (RADIUS credentials only) You specify in the RADIUS configuration page that the RADIUS server does not accept one-time passwords.

- Secondary credentials—If the primary credentials fail, the system determines whether it has secondary credentials for the user. If available, the system submits the user's secondary credentials provided that the conditions described for primary credentials are true.

- Last-entered credentials—If the system does not submit secondary credentials or if the credentials fail, it determines whether it has stored credentials for the specified user and a different resource in its cache. (See below for information about when the system caches credentials.) If available, the system submits its stored credentials provided the conditions described for primary credentials are true.

- User-specified credentials (prompt)—If the system does not submit last-entered credentials or if the credentials fail, it prompts the user to manually enter his credentials in the intermediate sign-in page. If the user selects the "Remember password?" check box, The system caches the user-specified credentials and, if necessary, resubmits them when the user tries to access the same resource again. Note that when the system caches these credentials, it remembers the specific user and resource, even after the user signs out of the device.

- Resource protected by a Web server or Web proxy using basic authentication—The system submits credentials in the following order to the backend server or proxy that is protecting the Web resource:

  - Cached credentials—If the system does not submit administrator-specified credentials or the credentials fail, it determines whether it has stored credentials for the specified user and resource in its cache. If available, the system submits its stored credentials.

  - Primary credentials—If the system does not submit cached credentials or the credentials fail, it submits the user's primary system credentials provided that following conditions are true:

    - The resource is in the same Intranet zone as the device (that is, the resource's hostname ends in the DNS suffix configured in the System > Network > Overview page of the admin console).

    - (Web proxies only) You have configured the system to recognize the Web proxy through settings in the Users > Resource Policies > Web > Web Proxy pages of the admin console.

    - The credentials are not ACE credentials.

    - (RADIUS credentials only) You specify in the RADIUS configuration page that the RADIUS server does not accept one-time passwords.

  - Secondary credentials—If the primary credentials fail, the system determines whether it has secondary credentials for the user. If available, it submits the user's secondary system credentials provided that the conditions described for primary credentials are true.

  - Last-entered credentials—If the system does not submit secondary credentials or if the credentials fail, it determines whether it has stored credentials for the specified user and a different resource in its cache. If available, the system submits its stored credentials provided the conditions described for primary credentials are true.

- User-specified credentials (prompt)—If the system does not submit last-entered credentials or if the credentials fail, it prompts the user to manually enter his credentials in the intermediate sign-in page. If the user selects the "Remember password?" check box, the system caches the user-specified credentials and, if necessary, resubmits them when the user tries to access the same resource again. Note that when the system caches these credentials, it remembers the specific user and resource, even after the user signs out of the device.

NOTE: The system does not support the multiple credential authentication mechanism described in this section with the e-mail client and SAML SSO mechanisms.

You cannot define an anonymous server, certificate server, SAML or eTrust SiteMinder server as a secondary authentication server.

If you define an eTrust SiteMinder server as your primary authentication server, you cannot define a secondary authentication server.

The system supports basic authentication and NTLM challenge/response scheme for HTTP when accessing web applications, but does not support HTTP-based cross-platform authentication via the negotiate protocol.

Related Documentation

- Using a RADIUS Server

CHAPTER 14 Synchronizing User Records

About User Record Synchronization

The user record synchronization feature promotes a more consistent user experience by allowing users to retain their bookmarks and individual preferences regardless of which device they log in to.

User record synchronization relies on client-server pairings. The client is the device that users log in to start their remote access. Each client is associated with one primary server and one backup server to store user record data. Clients can be individual appliances or a node within a cluster.

A server in this instance is the device that stores the user data records. Each server can be configured to replicate its user record data to one or more peer servers. Servers are identified by a user-defined logical name. The same logical name can be assigned to more than one authentication server to let you associate authentication servers of different types to the same user. For example, SA1 is an ACE authentication server with user1 who creates a bookmark to www.juniper.net. SA2 is an Active Directory authentication server with the same user1. For the www.juniper.net bookmark to be transferred from SA1/ACE/user1 to SA2/AD/user1 you would assign the logical name "Logical1" to both the ACE server on SA1 and the Active Directory server on SA2.

> **ⓘ** NOTE: Cluster VIPs cannot be used as the IP for synchronizing between clients and peers servers.

As long as the logical name is the same, the authentication servers can be different types and different server names and still be associated with a common user. The username must be the same for user record data to be synchronized across the servers. The logical authentication server (LAS) and username combination is what uniquely identifies a user record.

The following user records are synchronized between the client and server:

- Bookmarks
  - Web
  - File
  - Terminal Services
  - JSAM
- Preferences
- Persistent cookies
- Cached passwords

User session data is not synchronized. Persistent cookies, if changed, are synchronized when the user session terminates. All other modifications to the user records are synchronized immediately. User records are stored in cache on the client node prior to being pushed to the servers.

When a user logs in to a client, their data is pulled from the associated server. The pull is performed in the background and does not delay the login process. Users using browsers that do not support JavaScript must manually refresh the index page for updated bookmarks and preferences to appear. For browsers that support JavaScript, users may see a spinning progress indicator and their home page will refresh automatically with updated bookmarks and preferences.

Clients and servers need not be installed with the same Connect Secure software version as long as they are using version 7.2 or later.

> NOTE: User record synchronization uses port 17425. This port number is not configurable. If you are deploying across a firewall, configure your firewall to allow traffic on this port.

To set up user record synchronization, you perform the following tasks:

1. Enable user record synchronization for each participating client and server, identify which ones are the client and which ones are the server and assign a node name to each client and server.

2. Create a shared secret which is used to authenticate the client with the server and the server to its peer servers.

3. On each server, define which clients and peers are allowed to communicate with the server.

4. On each client, define the servers that handle records for each LAS server.

When enabling this feature, you have several options to initialize the user record database. You can:

- populate the database using user records located in the cache of the client systems.

- populate the database use user records located in the cache of the server systems.

- don't pre-populate the database but populate it as users log in and out of the client system.

If you choose the last option, users may not be able to view their saved bookmarks and preferences until the next time they log in, depending on which client they log in to.

> **NOTE:** User records may not synchronize if the time clocks on the devices are not in sync. We recommend that you use the same NTP server for each node participating in user record synchronization to keep system times accurately adjusted.
>
> The user record synchronization feature will not start automatically after importing a system configuration that has this feature enabled. The workaround is to disable user record synchronization and then enable user record synchronization from the user interface after the configuration import.

Related Documentation

- Enabling User Record Synchronization
- Configuring the User Record Synchronization Authentication Server
- Configuring the User Record Synchronization Server
- Configuring the User Record Synchronization Client
- Configuring the User Record Synchronization Database

Enabling User Record Synchronization

The first step in enabling user record synchronizing is to define the node name and the shared secret used to authenticate between the clients and the servers:

1. Select **System > Configuration > User Record Synchronization > General**.

2. Select the **Enable User Record Synchronization** check box.

3. Enter a unique node name. This name is used when associating a client with a server and is different from the logical name assigned to a server. This node name is also not the same as the cluster node name.

4. Enter the shared secret and confirm it.

   The shared secret is the password used to authenticate the client with its servers and the primary server with its peer servers. Use the same shared secret for all clients and servers participating in user record synchronization.

5. Select whether this node is client only or if this node acts as both a client and server.

6. Click **Save Changes**.

> **NOTE:** If you need to make any changes in this window at a later time, you must deselect the Enable User Record Synchronization check box and click Save Changes. Make your edits, select the Enable User Record Synchronization check box and save your changes.
>
> Once you enter a name and shared secret, you cannot clear these fields.

Related Documentation

- Configuring the User Record Synchronization Authentication Server

- Configuring the User Record Synchronization Server
- Configuring the User Record Synchronization Client
- Configuring the User Record Synchronization Database

Configuring the User Record Synchronization Authentication Server

To set up the authentication server you must define it's logical name:

1. Select **Authentication > Auth Servers**.

2. Click the name of the authentication server you want assign a LAS name.

   By assigning the authentication server a LAS name, all users that authenticate using the authentication server are associated with this LAS. In this instance, we are referring to the client nodes, not the user record synchronization server nodes.

3. Select the **User Record Synchronization** check box.

4. Enter a logical name to identify this server.

   This allows you to share user record data across authentication servers on different devices. By assigning a LAS name to an authentication server, you are implicitly assigning it to all users that authenticate with that auth server. The combination of the user's login name and their LAS name uniquely identifies the user's user record across all user record synchronization servers.

5. Click **Save Changes**.

Related Documentation

- Configuring the User Record Synchronization Client
- Configuring the User Record Synchronization Database

Configuring the User Record Synchronization Server

To set up the user record synchronization server you must define it's peer nodes (optional) and the clients that can access this server.

1. Select **System > Configuration > User Record Synchronization > This Server**.

2. Enter the peer server's node name and IP address, then click **Add**. To specify more than one peer server, enter each server's node name and IP address individually and click **Add**. There is no limit on the number of peer servers you can add.

   Data is replicated from the primary or backup server to its peer servers. If the primary is not available, user data is sent to the backup. User data is then replicated to the peer servers.

3. For each client you want synchronized with this server, enter the client's name and IP address and click **Add**.

Once added, peer servers will have a colored icon next to their name indicating their connection status. Node status is provided to client nodes and LAS mapping servers as well.

Color	Description
Green	Connected
Yellow	Connecting
Gray	Not connected

Related Documentation

- Configuring the User Record Synchronization Authentication Server
- Configuring the User Record Synchronization Client
- Configuring the User Record Synchronization Database

Configuring the User Record Synchronization Client

To set up the client, you select the primary and backup server you want this client to synchronize with:

1. Select **System > Configuration > User Record Synchronization > This Client**.

2. Select the LAS name you want to synchronize and enter the primary IP of the user record server that will server the user records. If you prefer to synchronize with any available server, select **Any LAS**.

3. Enter the primary and optionally a backup server's IP address and then click **Add**.

   Even if you select Any LAS, you must enter a primary server IP address.

Once added, the primary and backup servers have a colored icon next to their name indicating their connection status.

Related Documentation

- Configuring the User Record Synchronization Authentication Server
- Configuring the User Record Synchronization Server
- Configuring the User Record Synchronization Database

Configuring the User Record Synchronization Database

With the Database tab, you can delete inactive records from the client cache, retrieve statistics about the database, export and import the data and remove user data from the server's database.

To configure the database:

1. Select **System > Configuration > User Record Synchronization > Database**.

2. Select **Auto-delete inactive synchronized user records from the Cache** to remove inactive user records from the cache. This option does not remove user records from the user record database.

   When this option is selected, the system performs a check every 15 minutes and deletes user records that meet all of the following criteria:

- There are no active user sessions associated with the user record.

- The user record does not have any custom settings or the latest version of the user record has been synchronized with the user record database.

- The authentication server associated with the user record database does not have type "local". For example, the "System Local" auth server that is part of the default configuration of the system has a "local" type, so any user records associated with that auth server will not be auto-deleted. However, user records associated with external authentication servers like Radius or LDAP may be deleted, depending on the two prior criteria.

3. Select **Auto-delete user records from the local synchronization database that have been idle for X days** to permanently remove user records from the database located on the server. Enter the number of days user records must be inactive before being deleted.

   In this instance, "inactive" means that no client as pulled the user record or pushed any modifications to the user record in X days.

4. Click **Retrieve Statistics** to display the number of records in the database. You cannot edit or view records in the database.

5. Under Export, you export user records to a file. The user records can be exported from the user record database, or from the cache. The exported file can be used to pre-populate the user record database on another node.

   - Enter the LAS name of the user records you want to export. If you leave this field blank, all user records are exported. If you enter a LAS name, only user records with the entered LAS name are exported.

   - To encrypt the exported data, select the **Encrypt the exported data with password** check box and enter the password.

   - Click **Export** to export the user records from the specified source (cache or database). You will be prompted where to save the file.

6. Under Import, you import user records into the synchronization database. The user records can be imported from a file or from the cache. Use the Import operation to pre-populate the user record database with user records exported from another node, or with user records from the cache.

   - Click **Browse** to locate the exported file and enter the password if the exported file was encrypted with a password.

   - Select the **Override Logical Auth Servers in imported user records with** check box to replace the LAS name in each imported user record with the LAS name entered.

     For example, you change the LAS name, use this option to update the user records with the new name.

   - Click **Import**.

7. Under Delete, specify which user records to permanently remove from the user record database. The options you select apply only to the user record database associated with this server.

- Select **User record with login name and Logical Auth Server** to remove a specific record. The login name and LAS name together uniquely identify a user record. Select this option to remove that record (if it exists).

- Select **User records with Logical Auth Server** to delete all user records with the specified LAS name.

- Select **All user records** to permanently remove user records from the database on this node.

- Click **Delete**.

Related Documentation

- Configuring the User Record Synchronization Authentication Server

- Configuring the User Record Synchronization Server

- Configuring the User Record Synchronization Client

Scheduling User Record Synchronization Backup

You can configure periodic backups of the user record database. User record synchronization backup can be enabled only on a user record synchronization server.

To backup the user record database:

1. Ensure the system is set up as a user record synchronization server. See System > Configuration > User Record Synchronization.

2. Select **Maintenance > Archiving > Archiving Servers**.

3. Select the **Archive User Record Synchronization Database** check box.

4. Specify an archive schedule. Through the options, schedule archives on any combination of weekdays including weekends.

   > NOTE: If you schedule an archival operation to occur during the hour that your system switches to Daylight Savings Time (DST) the operation may not occur as scheduled. For example, if your system is set to change to DST at 1:00 a.m. and you have scheduled an archival operation to occur at anytime between 1:01 a.m. and 1:59 a.m., the operation is not accomplished, because at 1:00 a.m. the system clock is moved forward to 2:00 a.m. and the system never reaches your archival time for that date.

5. Define a specific time when you want the system to archive data or elect to archive data every hour, which produces twenty-four files with unique timestamps.

   > NOTE: We recommend you schedule an archival operation during hours when traffic is light in order to minimize its impact to your users. The automatic archiving process compresses files and, if the system is busy, can degrade performance for users. Also, a cluster node may appear unresponsive if the system is busy with traffic and performing archiving simultaneously.

6. Provide a password if you want to encrypt system configuration or user account archives with a password (optional).

7. Click **Save Changes**.

Related Documentation

- Configuring the User Record Synchronization Database

399

PART 3 Endpoint Defense

- Host Checker
- Cache Cleaner

CHAPTER 15 Host Checker

Host Checker and Trusted Network Computing

Host checker is a client-side agent that performs endpoint health and security checks for hosts that attempt to connect to a Connect Secure device.

The Trusted Computing Group (TCG) is a not-for-profit organization formed in 2003 to develop, define, and promote open standards for hardware-enabled trusted computing and security technologies across multiple platforms, peripherals, and devices. The TCG has over 100 members that include component vendors, software developers, systems vendors, and network companies.

Trusted Network Connect (TNC) is a subgroup of the TCG that created an architecture and set of standards for verifying endpoint integrity and policy compliance during or after a network access request. Many of the TCG members participated in the definition and specification of the TNC's architecture. The TNC defined several standard interfaces that enable components from different vendors to securely operate together. The TNC architecture is designed to build on established standards and technologies, such as 802 1X, RADIUS, IPsec, EAP, and TLS/SSL. For more information about TNC, see www.trustedcomputinggroup.org.

Using technology based on the TNC architecture and standards, the Host Checker component of the Policy Secure solution provides a comprehensive approach to assess the trust worthiness of endpoints.

You can use Host Checker to perform endpoint checks on hosts before allowing them to connect to the system and access protected resources. Host Checker can check for third party applications, files, process, ports, registry keys, and custom DLLs on hosts. Based on the results of the checks, it can then deny or allow access to protected resources. For example, you may choose to check for virus detection software before allowing a user access to any of the realms, launch the software on the user's system if necessary, map the user to roles based on individual policies defined in your own DLL, and then further restrict access to individual resources based on the existence of spyware detection software. When a user's computer does not meet the requirements you specify, you can display remediation instructions to users so they can bring their computers into compliance.

> **NOTE:** If you configure a large number of Host Checker policies, and the system is under a heavy load, the server process inside the device could get overloaded. When this happens, a message appears in the event log (Log/Monitoring > Events > Log).

Host Checker supports TNC-based integrity measurement collectors (IMCs) and integrity measurement verifiers (IMVs). IMCs are software modules that run on the host and collect information such as antivirus, antispyware, patch management, firewall, and other configuration and security information about the host. IMVs are software modules that run on the device and verify a particular aspect of an host's integrity. Each IMV on the device works with the corresponding IMC on the host to verify that the host meets the requirements of the integrity measurement custom rule(s) that you configure. IMCs frequently scan the client machine for changes in security status. Some IMCs can detect a change in status (for example, if the user turns off virus checking) and then trigger a new check to make sure the modified system complies with the requirements of the Host Checker policy. You can configure Host Checker to monitor third-party IMCs installed on client computers by using third-party IMVs that are installed on a remote IMV server.

You can invoke Host Checker at the role level, or the realm level to specify access requirements for endpoints attempting to authenticate.

All Host Checker rules are implemented through IMCs and IMVs based on the TNC open architecture. IMCs are software modules that Host Checker runs on the client machine.

IMCs are responsible for collecting information, such as antivirus, antispyware, patch management, firewall, and other configuration and security information for a client machine.

IMVs are software modules running on the device that are responsible for verifying a particular aspect of an endpoint's integrity.

The system and Host Checker manage the flow of information between the corresponding pairs of IMVs and IMCs. Each IMV on the device works with the corresponding IMC on the client machine to verify that the client meets the Host Checker rules.

You can also configure Host Checker to monitor third-party IMCs installed on client computers by using third-party IMVs that are installed on a remote IMV server.

Related Documentation

- Creating Global Host Checker Policies

- Task Summary: Configuring Host Checker

Task Summary: Configuring Host Checker

> **NOTE:** Ensure that user endpoints have signed ActiveX components or signed Java applets enabled within their browsers to permit Host Checker to download, install, and launch.

To configure a Host Checker policy, perform these tasks:

1. Create and enable Host Checker policies through the **Authentication > Endpoint Security > Host Checker** page of the admin console.

    1. In the admin console, select **Authentication > Endpoint Security > Host Checker**.

    2. Under Policies, click **New**.

    3. Enter a name in the Policy Name field and then click **Continue**. (Users see this name on the Host Checker remediation page if you enable custom instructions for this policy.)

    4. Create one or more rules to associate with the policy.

5. Configure additional system-level options on the **Authentication > Endpoint Security > Host Checker** page of the admin console as necessary:

    - If you want to display remediation information to users if they fail to meet the requirements of a Host Checker policy, configure remediation options through the **Authentication > Endpoint Security > Host Checker** page of the admin console.

    - For Windows clients, determine whether you need to use a pre-authentication access tunnel between the clients and policy server(s) or resources. If necessary, create a manifest.hcif file with the tunnel definition and upload it through the **Authentication > Endpoint Security > Host Checker** page of the admin console.

    - To change default Host Checker settings, configure settings through the **Authentication > Endpoint Security > Host Checker** page of the admin console.

6. Determine the level you that you want to enforce Host Checker policies:

    - To enforce Host Checker policies when the user initially accesses the device, implement the policy at the realm level by selecting the policy at the **Users > User Realms >** *Select Realm* **> Authentication Policy > Host Checker** page of the admin console.

    - To allow or deny users access to specific roles based on compliance with Host Checker policies, implement the policies at the role level by using the **Users > User Roles >** *Select Role* **> General > Restrictions > Host Checker** page of the admin console.

- To map users to roles based on their compliance with Host Checker policies, use custom expressions in the **Users > User Realms >** *Select Realm* **> Role Mapping** page of the admin console.

- To allow or deny users access to individual resources based on their compliance with Host Checker policies, use conditions in the Users > Resource Policies > *Select Resource* > *Select Policy* > Detailed Rules > *Select|Create Rule* page of the admin console.

7. Specify how users can access the Host Checker client-side agent that enforces the policies you define:

   - To enable automatic installation of the Host Checker client-side agent on all platforms, use the Administrators > Admin Realms > Select Realm > Authentication Policy > Host Checker page or the Users > User Realms > *Select Realm* > Authentication Policy > Host Checker page of the admin console.

   - To download the Host Checker installer and manually install it on your Windows users' systems, use the Maintenance > System > Installers page of the admin console.

8. Determine whether you want to create client-side logs. If you enable client-side logging through the System > Log/Monitoring > Client Logs page of the admin console, the system creates log files on your users' systems and writes to the file whenever Host Checker runs.

If more than one valid session exists from the same system, and Host Checker is used in those sessions, all of the valid sessions are terminated if a user signs out from any of the sessions. To prevent this, turn off Host Checker for those sessions that do not need Host Checker.

Related Documentation

- Introducing Pulse Secure Client

- Creating Global Host Checker Policies

Creating Global Host Checker Policies

To use Host Checker as a policy enforcement tool for managing endpoints, you create Host Checker policies through the **Authentication > Endpoint Security > Host Checker** page of the admin console, and then implement the policies at the realm, role, and resource policy levels.

The system provides many options that you can use to enable, create, and configure Host Checker policies:

- **Predefined policies (prevent in-network attacks or downloads malware detection software)**—The system comes equipped with a predefined client-side Host Checker policy that you simply need to enable. The Connection Control policy prevents attacks on Windows client computers from other infected computers on the same network.

- **Predefined rules (check for third party applications)**—Host Checker contains a wide array of predefined rules that check for antivirus software, firewalls, malware, spyware, and specific operating systems from a variety of industry leaders. You can enable one or more of these rules within a Host Checker client-side policy to ensure that the integrated third party applications that you specify are running on your users' computers.

- **Custom rules (check for additional requirements)**—In addition to Predefined rules, you can create custom rules within a Host Checker policy to define requirements that user endpoints must meet. Using custom rules, you can:

  - Configure Host Checker to check for custom third party DLLs that perform customized client-side checks.

  - Verify that certain ports are open or closed on the user's computer.

- Confirm that certain processes are or are not running on the user's computer.

- Check that certain files are or are not present on the client machine.

- Evaluate the age and content of required files through MD5 checksums.

- Confirm that registry keys are set on the client machine (Windows only).

- Check the NetBIOS name, MAC addresses, or certificate of the client machine (Windows only).

- Assess the client operating system and application service packs to ensure they are up to date (Windows only).

- Perform application and version checks to ensure that endpoints are running the correct software (Windows only).

- **Custom integrated applications (implement through server API)**—For Windows clients, you can upload a third-party J.E.D.I. DLL to the system.

- Within a single policy, you can create different Host Checker requirements for Windows, Macintosh and Linux, checking for different files, processes, and products on each operating system. You can also can combine any number of host check types within a single policy and check for alternative sets of rules.

Related Documentation

- Configuring a Remote IMV Server

- Implementing Host Checker Policies

Enabling Connection Control Host Checker Policies (Windows Only)

The predefined connection control Host Checker policy prevents attacks on Windows client computers from other infected computers on the same physical network.

> NOTE: The Host Checker connection control policy is not supported on Windows Vista or Windows 7.

> The Host Checker connection control policy blocks all incoming TCP, UDP and ICMP connections. This policy allows all outgoing TCP and VPN Tunneling traffic, as well as all connections to DNS servers, WINS servers, DHCP servers, proxy servers, and the system.

> NOTE: Users must have administrator privileges in order for Host Checker to enforce the connection control policy on the client computer.

To enable the predefined Host Checker connection control policy:

1. In the admin console, choose **Authentication > Endpoint Security > Host Checker**.

2. Under Options, select the **Create Host Checker Connection Control Policy** check box.

3. Click Save Changes. The system enables the Host Checker connection control policy.

   > NOTE: Note that you cannot modify this policy—only enable or disable it. Also note that since you cannot modify this policy, the system does not display it in the Policies section of the Authentication > Endpoint Security > Host Checker page with other configurable policies.

4. Implement the Host Checker connection control policy at the realm, role, or resource policy levels.

You must evaluate or enforce the connection control policy at the realm level to make the policy effective on client computers.

Related Documentation

- Configuring Host Checker Restrictions

Creating and Configuring New Client-side Host Checker Policies

You can create a variety of policies through the Host Checker client that check for antivirus software, firewalls, malware, spyware, and specific operating systems from a wide variety of industry leaders. You can also create checks for custom third-party DLLs, ports, processes, files, registry keys and the NetBIOS name, MAC addresses, or certificate of the client machine.

> *i*   NOTE: We recommend you check for multiple MAC addresses in a single policy instead of creating a policy for each MAC address. If you create policies for each MAC address, unexpected results may occur if there are more than 100 policies due to browser cookie size limitations.

When creating the policies, you must define the policy name, and either enable predefined rules, or create custom rules that run the specified checks. Optionally, you can specify how Host Checker should evaluate multiple rules within a single policy.

To create a standard client-side policy:

1. In the admin console, choose **Authentication > Endpoint Security > Host Checker**.

2. Under Policies, click **New**.

3. Enter a name in the Policy Name field and then click **Continue**. (Users see this name on the Host Checker remediation page if you enable custom instructions for this policy.)

4. Create one or more rules to associate with the policy.

5. Specify how Host Checker should evaluate multiple rules within the policy.

6. (Recommended) Specify remediation options for users whose computers do not meet the requirements specified in the policy. (If you do not create remediation instructions and the policy fails, your users will not know why they cannot access their resources.)

7. Implement the policy at the realm, role, or resource policy levels.

Related Documentation

- Checking for Third-Party Applications Using Predefined Rules (Windows Only)

- Specifying Customized Requirements Using Custom Rules

- Configuring General Host Checker Remediation

- Configuring Host Checker Restrictions

Checking for Third-Party Applications Using Predefined Rules (Windows Only)

Host Checker comes pre-equipped with a vast array of predefined rules that check for antivirus software, firewalls, malware, spyware, and specific operating systems from a wide variety of industry leaders. You can enable one or more of these rules within a Host Checker client-side policy to ensure that the integrated third party applications that you specify are running on your users' computers in accordance with your specifications. For firewall and antivirus rules, you can specify remediation actions to automatically bring the endpoint into compliance.

To view the currently supported applications, go to Authentication > Endpoint Security > Host Checker and create a new policy. You can choose predefined rule types from the Select Rule

Type drop down list box to see a list of the supported applications within that category. The lists of applications can be quite extensive and are updated at each support release, so it is useful to check the list periodically.

The following predefined rule types are available:

- **Predefined: AntiVirus**—Select this option to create a rule that checks for the antivirus software that you specify, and to specify remediation options.

- **Predefined: Firewall**—Select this option to create a rule that checks for the firewall software that you specify, and to specify remediation options.

- **Predefined: Malware**—Select this option to create a rule that checks for the malware protection software that you specify.

- **Predefined: AntiSpyware**—Select this option to create a rule that checks for the anti-spyware protection software that you specify.

- **Predefined: OS Checks**—Select this option to create a rule that checks for the Windows operating systems and minimum service pack versions that you specify. (Any service pack whose version is greater than or equal to the version you specify satisfies the policy.)

> NOTE: If the underlying TNCC service is killed or stopped, the endpoint can remain on the network, potentially out of compliance, until the next Host Checker policy refresh.

This section details Predefined Malware and Predefined OS check. Predefined Antivirus, Firewall and Malware checks are defined in sections that follow.

To create a Host Checker rule using Predefined Malware or Predefined OS Check rules:

1. In the admin console, select **Authentication > Endpoint Security > Host Checker**.

2. Create a new policy or click on an existing policy in the Policies section of the page.

3. Under Rule Settings, choose one of the following options and click **Add**:

   - Predefined Malware

   - Predefined OS Checks

   The predefined rule page opens.

4. In the **Rule Name field**, enter an identifier for the rule.

5. Under Criteria, select the specific malware or operating systems that you want to check for and click **Add**. (When checking for an operating system, you may also specify a service pack version.)

   > NOTE: When you select more than one type of software within a predefined rule, Host Checker considers the rule satisfied if any of the selected software applications are present on the user's machine.

6. Under Optional, select **Monitor this rule for change in result** to continuously monitor the policy compliance of endpoints. If this check box is selected, and a change in compliance status on an endpoint that has successfully logged in occurs, the system initiates a new handshake to re-evaluate realm or role assignments.

   > NOTE: Use this option only for dynamic rules, such as checking whether Real Time Protection is enabled on the antivirus software. Use the host checker update frequency to monitor other rules periodically.

7. Click **Save Changes**.

8. Optionally add additional rules to the policy, specify how Host Checker should evaluate multiple rules within the policy, and define remediation options.

Related Documentation

- Creating and Configuring New Client-side Host Checker Policies

- Configuring a Predefined Firewall Rule with Remediation Options (Windows Only)

- Configuring a Predefined Antivirus Rule with Remediation Options

- Implementing Host Checker Policies

Configuring a Predefined Antivirus Rule with Remediation Options

You can configure antivirus remediation actions with Host Checker. You can specify a requirement for the age (in days) of the last successful virus scan, and you can specify that virus signatures installed on client machines should not be older than a specified number of updates.

You can also monitor policies to ensure that logged-in endpoints maintain compliance status, and remediate the endpoint to another role or realm depending on the current status.

If a client attempts to log in, and the client machine does not meet the requirements you specify, Host Checker can attempt to correct the deficiencies to allow the client to successfully log in. With Host Checker antivirus remediation, you can prompt the endpoint to download the latest virus signature files, turn on antivirus protection, and initiate an antivirus scan.

All of the remediation options are not supported for all antivirus software vendors' products. All available vendors and products that are supported are displayed when you select the Require any supported product option button.

Alternately, you can select the Require specific products/vendors option button and select either the Require any supported product from a specific vendor or Require specific products check boxes, then add an available type to Selected Types. The remediation options appear, and you can determine which remediation options are available for specific products or vendors

To configure a Predefined Antivirus rule:

1. In the admin console, select **Authentication > Endpoint Security > Host Checker**.

2. Create a new policy or click on an existing policy in the **Policies** section of the page.

3. Under Rule Settings, choose **Predefined: Antivirus** and click **Add**.

4. Enter the **Rule Name** for this antivirus rule.

5. To determine if your software vendor's product is supported for the System Scan check, click **these Antivirus products**. A new window will open with a list of all of the products that support the feature.

6. Select or clear the check box next to **Successful System Scan must have been performed in the last _ days**, and enter the number of days in the field.

   If you select this check box, a new option appears. If the remediation action to start an antivirus scan has been successfully begun, you can override the previous check.

7. Select or clear the check box next to **Consider this rule as passed if 'Full System Scan' was started successfully as remediation**.

8. Select or clear the check box next to **Virus definition files should not be older than _** **updates**. Enter a number between 1 and 10. If you enter 1, the client must have the latest update. You must import the virus signature list for the supported vendor.

9. Select your antivirus vendor(s) and product(s) by using either the **Require any supported product** or **Require specific products/vendors** option buttons.

   Require any supported product allows you to check for any product (rather than requiring you to select every product separately). This option button reveals a list of products in the remediation section to allow you to enable remediation options which are product specific.

   Require specific products/vendors allows you to define compliance by allowing any product by a specific vendor (for example, any Symantec product).

   Require specific products provides functionality that allows you to select individual products to define compliance.

   After you select your vendor(s) and product(s), remediation options will appear on the page.

   For each of the following remediation actions:

   - **Download latest virus definition files**—obtains the latest available file for the specified vendor from the vendor's web site

   - **Turn on Real Time Protection**—launches the virus scanning mechanism for the specified vendor

   - **Start Antivirus Scan**—performs a real-time virus scan for the specified vendor

   The check box is active (clickable) if the action is supported for your product.

   If your antivirus product is not supported, you can click the remediation column headers to determine what vendors and products are supported.

10. If your product is supported, select the check box for any or all of the remediation actions that you want to apply.

11. Under Optional, select **Monitor this rule for change in result** to continuously monitor the policy compliance of endpoints. If this check box is selected, and a change in compliance status on an endpoint that has successfully logged in occurs, the system initiates a new handshake to re-evaluate realm or role assignments.

    NOTE: Use this option only for dynamic rules, such as checking whether Real Time Protection is enabled on the antivirus software. Use the host checker update frequency to monitor other rules periodically.

12. Click **Save Changes** to save the antivirus rule and enforce antivirus remediation.

13. Optionally add additional rules to the policy, specify how Host Checker should evaluate multiple rules within the policy, and define remediation options.

Related Documentation

- Configuring Virus Signature Version Monitoring and Patch Assessment Data Monitoring

- Creating and Configuring New Client-side Host Checker Policies

- Implementing Host Checker Policies

Configuring a Predefined Firewall Rule with Remediation Options (Windows Only)

You can configure firewall remediation actions with Host Checker after you create a Host Checker firewall rule that requires the endpoint to have a specific firewall installed and running prior to connecting to the network.

After you enforce the Host Checker rule with firewall remediation actions, if an endpoint attempts to log in without the required firewall running, Host Checker can attempt to enable the firewall on the client machine.

The remediation option is not supported for all firewall products. All available products are displayed by using the Require any supported product or Require specific products/vendors option buttons.

To configure a Host Checker Predefined Firewall rule:

1. In the admin console, choose **Authentication > Endpoint Security > Host Checker**.

2. Create a new policy or click an existing policy in the Policies section of the page.

3. Under Rule Settings, choose **Predefined: Firewall** and click **Add**.

4. Enter a **Rule Name** for the firewall rule.

5. Select your firewall vendor(s) and product(s) by using either the **Require any supported product** or **Require specific products/vendors** option buttons.

   Require any supported product allows you to check for any product (rather than requiring you to select every product separately). This option button reveals a list of products in the remediation section to allow you to enable remediation options which are product specific.

   When you add an available product to Selected Products, the remediation option appears, and you can determine if the remediation option is available for your selected firewall.

   Require specific products/vendors allows you to define compliance by allowing any product by a specific vendor (for example, any Symantec product).

   Require specific products provides functionality that allows you to select individual products to define compliance.

   After you select your vendor(s) and product(s), the remediation options will appear on the page. The Turn on Firewall check box is active (clickable) if the action is supported for your product.

6. If your firewall is supported, select the check box to **Turn on Firewall**.

7. Under Optional, select **Monitor this rule for change in result** to continuously monitor the policy compliance of endpoints. If this check box is selected, and a change in compliance status on an endpoint that has successfully logged in occurs, the system initiates a new handshake to re-evaluate realm or role assignments.

8. Click **Save Changes** to save the firewall rule and enforce firewall remediation.

9. Optionally add additional rules to the policy, specify how Host Checker should evaluate multiple rules within the policy, and define remediation options.

Related Documentation

- Creating and Configuring New Client-side Host Checker Policies

- Implementing Host Checker Policies

Configuring a Predefined AntiSpyware Rule (Windows Only)

You can configure Host Checker to check for installed antispyware on endpoints.

After you enforce the Host Checker rule, if an endpoint attempts to log in without the required spyware, the Host Checker rule will fail.

The option is not supported for all spyware products. All available products are displayed by using the Require any supported product or Require specific products/vendors option buttons.

To configure a Host Checker Predefined Spyware rule:

1. In the admin console, select **Authentication > Endpoint Security > Host Checker**.

2. Create a new policy or click an existing policy in the Policies section of the page.

3. Under Rule Settings, choose **Predefined: AntiSpyware** and click **Add**.

4. Enter a **Rule Name** for the firewall rule.

5. Select one of the following options:

   - Select the **Require any supported product** option button to check for any product (rather than requiring you to select every product separately).

   - Select the **Require specific products/vendors** option button to specify the spyware that you want to check for.

     - Choose either the **Require any supported product from a specific vendor** or **Require specific products** to specify spyware.

     - Add antispyware from Available Products to Selected Products.

6. Under Optional, select **Monitor this rule for change in result** to continuously monitor the policy compliance of endpoints. If this check box is selected, and a change in compliance status on an endpoint that has successfully logged in occurs, the system initiates a new handshake to re-evaluate realm or role assignments.

7. Click **Save Changes**.

8. Optionally add additional rules to the policy, specify how Host Checker should evaluate multiple rules within the policy, and define remediation options.

Related Documentation

- Creating and Configuring New Client-side Host Checker Policies

- Implementing Host Checker Policies

Configuring Virus Signature Version Monitoring and Patch Assessment Data Monitoring

You can configure Host Checker to monitor and verify that the virus signatures, operating systems, software versions, and patches installed on client computers are up to date, and remediate those endpoints that do not meet the specified criteria. Host Checker uses the current virus signatures and patch assessment versions from the vendor(s) you specify for predefined rules in a Host Checker policy.

You can automatically import the current Virus signature version monitoring or Patch Management Info Monitoring lists from the Pulse Secure staging site at a specified interval, or you can download the files from Pulse Secure and use your own staging server.

You can also configure a proxy server as a staging site between the system and the Pulse Secure site. To use a proxy server, you enter the servers network address, port and authentication credentials, if applicable.

To access the Pulse Secure staging site for updates, you must enter the credentials for your Pulse Secure Support account.

To configure the system automatically import the current virus signature version monitoring and patch management version monitoring list(s) from the Pulse Secure staging site:

1. Choose **Authentication > Endpoint Security > Host Checker**.

2. Click **Virus signature version monitoring**, or **Patch Management Info Monitoring**.

3. Select **Auto-update virus signatures list** or **Auto-update Patch Management data**.

4. For Download path, leave the existing URL(s) of the staging site(s) where the current list(s) are stored. The default URLs are the paths to the Pulse Secure staging site:

   https://download.juniper.net/software/av/uac/epupdate_hist.xml

   (for auto-update virus signatures list)

   https://download.juniper.net/software/hc/patchdata/patchupdate.dat

   (for auto-update patch management)

5. For **Download interval**, specify how often you want the system to automatically import the current list(s).

6. For **Username** and **Password**, enter your Pulse Secure Support credentials.

7. Click **Save Changes**.

To manually import the current virus signature version monitoring and patch management version monitoring list(s):

1. Choose **Authentication > Endpoint Security > Host Checker**.

2. Click **Virus signature version monitoring**, or **Patch Management Info Monitoring**.

3. Download the list(s) from the Pulse Secure staging site to a network server or local drive on your computer by entering the Pulse Secure URLs in a browser window.

   http://www.pulsesecure.net/support

4. Under Manually import virus signatures list, click **Browse**, select the list, and then click **OK**.

5. Click **Save Changes**.

   NOTE: If you use your own staging site for storing the current list(s), you must upload the trusted root certificate of the CA that signed the staging's server certificate to the system.

To use a proxy server as the auto-update server:

1. Choose **Authentication > Endpoint Security > Host Checker**.

2. Click **Virus signature version monitoring**, or **Patch Management Info Monitoring**.

3. Select **Auto-update virus signatures list** or **Auto-update Patch Management data**.

4. For Download path, leave the existing URL(s) of the staging site(s) where the current list(s) are stored. The default URLs are the paths to the Pulse Secure staging site:

   http://www.pulsesecure.net/support

   (for auto-update virus signatures list)

http://www.pulsesecure.net/support

(for auto-update patch management)

5. For **Download interval**, specify how often you want the system to automatically import the current list(s).

6. For **Username** and **Password**, enter your Pulse Secure Support credentials.

7. Select the check box for **Use Proxy Server**.

8. Enter the **IP Address** of your proxy server.

9. Enter the **Port** that the Pulse Secure Support site will use to communicate with your proxy server.

10. If your proxy server is password protected, type the **Username** and **Password** of the proxy server.

11. Click **Save Changes**.

Related Documentation

- Using Trusted Server CAs

- Implementing Host Checker Policies

Host Checker Statement of Health for Pulse Connect Secure Overview

You can use the open standard Statement of Health (SoH) rule in a Host Checker policy for the Pulse for Windows client and for the Windows in-box Pulse client. SoH components evaluate an endpoint's state of health and make policy decisions for network access based on the result of the health check. To use SoH with the Windows in-box Pulse client, you must also enable the SoH functionality on the endpoint.

NOTE: The Statement of Health rule is the only Host Checker rule supported by the Windows in-box Pulse client, and the Statement of Health rule is available only with Connect Secure release 8.0 and later.

You can use the SoH health state validation to determine which roles or realms can be accessed by endpoints. If an endpoint fails the SoH check, or if the SoH cannot be negotiated successfully, the Host Checker policy fails.

You can check the following system health indicators:

- Antivirus is enabled.

- Antivirus is up to date.

- Antispyware is enabled.

- Antispyware is up to date.

- Firewall is enabled.

- Automatic updating is enabled

Related Documentation

- *Microsoft Windows In-Box Pulse Secure Client Overview*

- *Microsoft Windows In-Box Pulse Secure Client User Interface*

- *Windows PowerShell Script Examples for Microsoft Windows In-Box Pulse Secure Client*

- *Host Checker for Windows In-Box Pulse Secure Client Overview*

- *Enabling Statement of Health on the Windows Endpoint for the Windows In-Box Pulse Secure Client*

- *Viewing Windows In-Box Pulse Secure Client Log Messages*

Configuring a Statement of Health Host Checker Policy for Pulse Connect Secure

You can use the open standard Statement of Health rule in a Host Checker policy for both the Pulse for Windows client and the Windows in-box Pulse client.
To configure a Statement of Health rule in a Host Checker policy:

1. In the admin console, select **Authentication > Endpoint Security > Host Checker**.

2. In the Policies section, click **New** to create a new policy, or click an existing policy.

3. For a new policy, specify a name for the policy and then click **Continue**.

4. Click the **Windows** tab. Statement of Health is available for Windows endpoints only.

5. Under Rule Settings, select **Custom: Statement of Health**, and then click **Add**.

6. Type a Rule Name for this rule.

To configure the SoH rule, you must select one or more of the Statement of Health parameters.

1. Under Criteria, enter a **Label** for the selected SoH parameter, or accept the default.

2. Select an SoH policy option from the Parameter menu, and then click **Add** for the following types:

   - Antivirus Enabled

   - Antivirus up to date

   - Antispyware enabled

   - Antispyware up to date

   - Firewall Enabled

   - Automatic Updating Enabled

3. Select additional options from the Parameter list to add additional SoH parameters.

4. (Optional) For each rule, select the **Enable automatic remediation** check box. If you select this option for a rule, the user receives a remediation message from the SoH agent, and appropriate remediation is performed, if possible. If the box is not selected, the user receives a remediation message, but no remediation action is performed.

   NOTE: Automatic remediation works for the Pulse for Windows client only. The Windows in-box Pulse client does not support automatic remediation.

5. Click **Save Changes**.

Related Documentation

- *Microsoft Windows In-Box Pulse Secure Client Overview*

- *Microsoft Windows In-Box Pulse Secure Client User Interface*

- *Windows PowerShell Script Examples for Microsoft Windows In-Box Pulse Secure Client*

- Host Checker Statement of Health for Pulse Connect Secure Overview

- *Enabling Statement of Health on the Windows Endpoint for the Windows In-Box Pulse Secure Client*

- *Viewing Windows In-Box Pulse Secure Client Log Messages*

Patch Management Info Monitoring and Patch Deployment

You can configure Host Checker policies that check for a Windows endpoint's operating system service pack, software version, or desktop application patch version compliance. Host Checker uses a list of the most current patch versions from the vendor for predefined rules in the Host Checker policy. Host Checker does not scan for non-security patches.

You obtain the most current patch version information from a Pulse Secure staging site. You can manually download and import the list into the system, or you can automatically import the list from the Pulse Secure staging site or your own staging site at a specified interval.

Monitoring is based on either one or more specified products or on specific patches, though not in the same policy. For example, you could check for Internet Explorer Version 7 with one policy, and Patch MSOO-039: SSL Certificate Validation Vulnerabilities with a second policy. Then, apply both policies to endpoints at the role or realm level to ensure that the user has the latest browser version with a specific patch. In addition, for Microsoft products, you can specify the severity level of patches that you wish to ignore. For example, you could ignore low or moderate threats.

The system can send remediation instructions (such as. a message describing what patches or software are non-compliant, and a link to where the endpoint can obtain the patch). The system does not autoremediate in the event of a non-compliant endpoint. However, you can send the items to the client for manual remediation of managed machines.

When an endpoint first connects to the device, the latest versions of the data files and libraries of the IMC are downloaded to the host computer. The initial check takes 10- 20 seconds to run, depending on the link speed. If the files are outdated, they are automatically updated at subsequent checks. If this is the first time the endpoint has connected to the device with the patch assessment policy, and the connection is a Layer 2 connection, the IMC required to run the Patch Assessment check cannot download. In this case, you should configure a remediation role that displays instructions to direct the user to retry with a Layer 3 connection or contact the administrator.

Note that in non-English installations, the English version of local patches is displayed.

Endpoints with Pulse that are not in compliance with specified patch policies can be updated with the required patches and brought into compliance automatically. This is achieved through a new patch deployment engine. The patch deployment engine executes on endpoints, downloads specified patches, and installs patches that are required through the Host Checker policy. The patch deployment engine provides a new means of remediating endpoints that do not meet the patch assessment policies defined on the system. This functionality is available for Windows 32-bit and 64-bit versions.

The Host Checker IMC on the endpoint interfaces with the patch deployment engine to download and install missing patches reported by the IMV. When the patch installation is complete, the IMC signals the TNC client to start a new handshake with the device, and enables the system to make access control decisions based on the result of the handshake.

Endpoints without Pulse can still use the legacy basic patch remediation mechanism, in which a pre-installed SMS/SCCM client is triggered to get patches from a pre-configured SMS/SCCM server. This mechanism installs only those patches that are published on the SMS/SCCM server. If the SMS/SCCM client is not installed, or the server doesn't host the patches required by the policies on the system, the endpoint cannot be fully remediated.

The patch deployment engine is an executable file that is hosted on the device. The executable can be downloaded to any endpoint that you would like to remediate. Unlike the SMS/SCCM client, one can specify what patches need to be applied. The patch deployment engine directly

downloads missing patches from vendor web sites, without going through the system. Therefore, Internet connectivity is needed for Shavlik remediation to work. The patch deployment engine does not work with Layer 2 without Layer 3 connectivity. You can configure a remediation VLAN for post authentication. Once Layer 3 connectivity is obtained, endpoints can remediate successfully. With Layer 3 connectivity, the patch deployment engine downloads missing patches.

A separate license is required for patch info monitoring and deployment functionality. It is not available as part of the endpoint solution.

All of the files required for patch deployment are a part of a ESAP packages beginning with software version 7.1. The default ESAP package shipped with software version 7.1 contains the required patch deployment files. Any older ESAP packages fail to update on these devices.

The IMC and IMV for patch monitoring is backward compatible. Since this feature is available from Pulse 2.0 onward, a new Pulse communicating with an older IMV (with Pulse support), or a new IMV communicating to an older IMC exhibit the same behavior as today. There should be no change in the patch assessment, and Shavlik's deployment engine is not invoked for remediation.

User Experience

Patch remediation can take a good deal of time, and policies continue to fail until the process is completed. When an update is required, the user is given an option to proceed with patch deployment. If the user decides not to deploy the patch(es) and proceed, the user may not have connectivity or may have limited connectivity, depending on the administration configuration. If any patches require a reboot subsequent to installation, the application informs Host Checker, and Pulse notifies the user. In this case, until the machine is rebooted, patches continue to be reported as missing in subsequent patch assessments. If a reboot is required, any further patch deployment is not carried out until the machine is rebooted. The user is notified if a reboot is required.

Pulse notifies the user that patches need to be installed, and provides status as the download is occurring. When the installation is complete, the client presents the login dialogue.

Using a System Management Server

You can use Microsoft System Management Server or Microsoft System Center Configuration Manager (SMS/SCCM) to provide a method for automatic updates to non-compliant software.

Pulse can support either the SMS/SCCM download method or the patch deployment engine for patch deployment, depending on the configuration on the system. If the system is configured for the SMS/SCCM method for patch deployment, the client machine should have the SMS/SCCM client already installed in the machine for deployment to begin , otherwise remediation fails.

Endpoints configured with SMS/SCCM for software management typically poll the server for updates every fifteen minutes or longer. In a worst-case scenario, clients that are not in compliance with existing Host Checker software requirements might have to wait until the next update interval to login.

Using the SMS/SCCM download method, you can force the client to initiate the software update immediately after the patch assessment check.

If a user attempts to log in, and the endpoint does not have a required software version for compliance with a Host Checker patch assessment policy, Host Checker immediately notifies the client to poll the server for an immediate update. The client receives notification that an SMS/SCCM update has started.

To configure SMS/SCCM to update the client when notified, set the advertisement time on the SMS/SCCM to As soon as possible. The following process then occurs:

- The patch assessment policy specifies the required software.

- When an endpoint attempts to authenticate, Host Checker evaluates the client and sends the results back to the system.

- The system evaluates the results and sends reason strings and remediation information to the client, including a message that directs the client to poll the server for software advertisements immediately.

- The SMS/SCCM client queries the SMS/SCCM server for software advertisements.

- The server identifies what patches should be advertised to the client (this is configured on the server, Host Checker does not interact with the server).

- The SMS/SCCM client receives the advertisement and applies the required patch(es).

You assign clients to a particular group or collection on the server. Then the SMS/SCCM server can advertise patches for that collection. You can configure roles that correspond to collections, and SMS/SCCM can send the appropriate patches for a particular role.

You must have the SMS/SCCM client installed and configured correctly on endpoints, and the SMS/SCCM server must be reachable. In a Layer 2 network, Host Checker is performed before the endpoint is connected to the network. Host Checker can obtain the IP address of the SMS/SCCM server configured for the client. If the endpoint is out of compliance and remediation is necessary, Host Checker pings the server IP address every 15 seconds until the server can be notified to update the client.

It is important as an administrator to inform users of the expected behavior if this feature is enabled, as there is no notification to the user until the SMS/SCCM sends back the advertisement.

Specifying Customized Requirements Using Custom Rules

In addition to the predefined policies and rules that come with the system, you can create custom rules within a Host Checker policy to define requirements that your users' computers must meet. Using custom rules, you can:

- Configure remote integrity measurement verifiers (IMVs) to perform customized client-side checks.

- Configure Host Checker to check for custom DLLs that perform customized client-side checks.

- Verify that certain ports are open or closed on the user's computer.

- Confirm that certain processes are or are not running on the user's computer.

- Check that certain files are or are not present on the client machine.

- Evaluate the age and content of required files through MD5 checksums.

- Confirm that registry keys are set on the client machine.

- Confirm the NETBIOS name of the client machine.

- Confirm the MAC addresses of the client machine.

- Check the validity of the machine certificate that is installed on the user's computer.

> NOTE: You can only check for registry keys, third-party DLLs, NETBIOS names, MAC addresses, and machine certificates on Windows computers.

To create a client-side Host Checker policy:

1. In the admin console, select **Authentication > Endpoint Security > Host Checker**.

2. Create a new policy or click an existing policy in the Policies section of the page.

3. Click the tab that corresponds to the operating system for which you want to specify Host Checker options—Windows, Mac, Linux or Solaris. In the same policy, you can specify different Host Checker requirements on each operating system. For example, you can create one policy that checks for different files or processes on each operating system.

   > NOTE: You must explicitly create policies for each operating system you want to allow. For example, if you create a Windows Host Checker policy, but don't create one for Mac or Linux, users who sign into the device from a Mac or Linux machine will not comply with the Host Checker policy and therefore will not be able to access the realm, role, or resource on which you enforce Host Checker.

4. Under Rule Settings, choose the options in the following sections and click **Add**. The Add Custom Rule page for the rule type appears.

   - **Custom: Remote IMV**—Use this rule type to configure integrity measurement software that a client must run to verify a particular aspect of the client's integrity, such as the client's operating system, patch level, or virus protection.

   - **3rd Party NHC Check**—Use this rule type to specify the location of a custom DLL (Windows only). Host Checker calls the DLL to perform customized client-side checks. If the DLL returns a success value to Host Checker, then the system considers the rule met. In the 3rd Party NHC Check configuration page:

     1. Enter a name and vendor for the 3rd Party NHC Check rule

     2. Enter the location of the DLL on client machines (path and file name).

     3. Click **Save Changes**.

        The 3rd Party NHC Check feature is primarily provided for backwards compatibility. We recommend that you use IMCs and IMVs instead

   - **Ports**—Use this rule type to control the network connections that a client can generate during a session. This rule type ensures that certain ports are open or closed on the client machine before the user can access the device. In the Ports configuration page:

     1. Enter a name for the port rule.

     2. Enter a comma delimited list (without spaces) of ports or port ranges, such as: 1234,11000-11999,1235.

     3. Select **Required** to require that these ports are open on the client machine or Deny to require that they are closed.

     4. Under Optional, select **Monitor this rule for change in result** to continuously monitor the policy compliance of endpoints. If this check box is selected, and a change in compliance status on an endpoint that has successfully logged in occurs, the system initiates a new handshake to re-evaluate realm or role assignments.

     5. Click **Save Changes**.

   - **Process**—Use this rule type to control the software that a client may run during a session. This rule type ensures that certain processes are running or not running on the client machine before the user can access resources protected by the system. In the Processes configuration page:

     1. Enter a name for the process rule.

2. Enter the name of a process (executable file), such as: good-app.exe.

> ℹ️ NOTE: For Linux, Macintosh and Solaris systems, the process that is being detected must be started using an absolute path.

You can use a wildcard character to specify the process name.

For example: good*.exe

3. Select **Required** to require that this process is running or **Deny** to require that this process is not running.

4. Specify the MD5 checksum value of each executable file to which you want the policy to apply (optional). For example, an executable may have different MD5 checksum values on a desktop, laptop, or different operating systems. On a system with OpenSSL installed—many Macintosh, Linux and Solaris systems have OpenSSL installed by default—you can determine the MD5 checksum by using this command: openssl md5 <processFilePath>

5. Click **Save Changes**.

- **File**—Use this rule type to ensure that certain files are present or not present on the client machine before the user can access the device. You may also use file checks to evaluate the age and content (through MD5 checksums) of required files and allow or deny access accordingly. In the Files configuration page:

1. Enter a name for the file rule.

2. Enter the name of a file (any file type), such as: c:\temp\bad-file.txt or /temp/bad-file.txt.

You can use a wildcard character to specify the file name. For example:

*.txt

You can also use an environment variable to specify the directory path to the file. (You cannot use a wildcard character in the directory path.) Enclose the variable between the <% and %> characters. For example:

<%windir%>\bad-file.txt

3. Select **Required** to require that this file is present on the client machine or **Deny** to require that this file is not present.

4. Specify the minimum version of the file (optional). For example, if you require notepad.exe to be present on the client, you can enter 5.0 in the field. Host Checker accepts version 5.0 and later, of notepad.exe.

5. Specify the maximum age (File modified less than n days) (in days) for a file (optional). If the file is older than the specified number of days, then the client does not meet the attribute check requirement.

> ℹ️ NOTE: You can use the maximum age option to check the age of virus signatures. Make sure you specify the path to a file in the File Name field whose timestamp indicates when virus signatures were last updated, such as a virus signature database or log file that updates each time the database updates. For example, if you use TrendMicro, you may specify:
>
> C:\Program Files\Trend Micro\OfficeScan Client\TmUpdate.ini.

6. Specify the MD5 checksum value of each file to which you want the policy to apply (optional). On Macintosh, Linux and Solaris, you can determine the MD5 checksum by using this command: openssl md5<filePath>

7. Select **Monitor this rule for change in result** to continuously monitor the policy compliance of endpoints. If this check box is selected, and a change in compliance status on an endpoint that has successfully logged in occurs, the system initiates a new handshake to re-evaluate realm or role assignments.

8. Click **Save Changes**.

- **Registry Setting**—Use this rule type to control the corporate PC images, system configurations, and software settings that a client must have to access the device (Windows only). This rule type ensures that certain registry keys are set on the client machine before the user can access the device. You may also use registry checks to evaluate the age of required files and allow or deny access accordingly. In the Registry Settings configuration page:

   1. Enter a name for the registry setting rule.

   2. Select a root key from the drop-down list.

   3. Enter the path to the application folder for the registry subkey.

   4. Enter the name of the key's value that you want to require (optional). This name appears in the Name column of the Registry Editor.

   5. Select the key value's type (String, Binary, or DWORD) from the dropdown list (optional). This type appears in the Type column of the Registry Editor.

   6. Specify the required registry key value (optional). This information appears in the Data column of the Registry Editor.

      If the key value represents an application version, select Minimum version to allow the specified version or newer versions of the application. For example, you can use this option to specify version information for an antivirus application to make sure that the client antivirus software is current. The system uses lexical sorting to determine if the client contains the specified version and later. For example:

      3.3.3 is newer than 3.3

      4.0 is newer than 3.3

      4.0a is newer than 4.0b

      4.1 is newer than 3.3.1

      > **(i)** NOTE: If you specify only the key and subkey, Host Checker simply verifies the existence of the subkey folder in the registry.

   7. Under Optional, select **Monitor this rule for change in result** to continuously monitor the policy compliance of endpoints. If this check box is selected, and a change in compliance status on an endpoint that has successfully logged in occurs, the system initiates a new handshake to re-evaluate realm or role assignments.

      You can configure registry setting remediation actions with Host Checker. If a client attempts to login, and the client machine does not meet the requirements you specify, Host Checker can attempt to correct the discrepancies to allow the client to login.

   8. Select the check box for **Set Registry value specified in criteria**.

   9. Click **Save Changes**.

- **NetBIOS (Windows only, does not include Windows Phone)**—Use this rule type to check the NetBIOS name of the client machine before the user can access the device. In the NetBIOS configuration page:

   1. Enter a name for the NetBIOS rule.

2. Enter a comma-delimited list (without spaces) of NetBIOS names. The name can be up to 15 characters in length. You can use wildcard characters in the name and it is not case-sensitive. For example, md*, m*xp and *xp all match MDXP.

3. Select **Required** to require that NETBIOS name of the client machine match one of the names you specify, or **Deny** to require that the name does not match any name.

4. Click **Save Changes**.

- **MAC Address (Windows only)**—Use this rule type to check the MAC addresses of the client machine before the user can access the device. In the MAC Address configuration page:

1. Enter a name for the MAC address rule.

2. Enter a comma-delimited list (without spaces) of MAC addresses in the form XX:XX:XX:XX:XX:XX where the X's are hexadecimal numbers. For example:

   00:0e:1b:04:40:29

   You can use a * wildcard character to represent a two-character section of the address. For example, you can use a * to represent the "04", "40", and "29" sections of the previous example address:

   00:0e:1b:*:*:*

   But you cannot use a * to represent a single character. For example, the * in the following address is not allowed:

   00:0e:1b:04:40:*9

3. Select **Required** to require that a MAC address of the client machine matches any of the addresses you specify, or **Deny** to require that the all addresses do not match. A client machine will have at least one MAC address for each network connection, such as Ethernet, wireless, and VPN. This rule's requirement is met if there is a match between any of the addresses you specify and any MAC address on the client machine.

4. Click **Save Changes**.

   NOTE: Since the MAC address is changeable on some network cards, this check may not guarantee that a client machine meets the requirements of your Host Checker policy.

- **Machine Certificate (Windows only)**—Use this rule type to check that the client machine is permitted access by validating the machine certificate stored on the client machine. In the Machine Certificate configuration page:

1. Enter a name for the machine certificate rule.

2. From the Select Issuer Certificate list, select the certificate that you want to retrieve from the user's machine and validate. Or, select Any Certificate to skip the issuer check and only validate the machine certificate based on the optional criteria that you specify below.

3. From the Optional fields (Certificate field and Expected value), specify any additional criteria that Host Checker should use when verifying the machine certificate.

4. Click **Save Changes**.

   NOTE: If more than one certificate is installed on the client machine that matches the specified criteria, The Host Checker client passes the first certificate it finds to the system for validation.

5. Optionally add additional rules to the policy, specify how Host Checker should evaluate multiple rules within the policy, and define remediation options.

Related Documentation

- Implementing Host Checker Policies

- Task Summary: Configuring Host Checker

- Using a Wildcard or Environment Variable in a Host Checker Rule

Using a Wildcard or Environment Variable in a Host Checker Rule

You can use the following wildcards to specify a file name in a **Custom File** rule or a process name in a **Custom Process** rule:

Table74: Wildcard Characters for Specifying a File Name or Process Name

Wildcard Character	Description	Example
*	Matches any character	*.txt
?	Matches exactly one character	app-?.exe

In a **Custom File** rule for Windows, you can use the following environment variables to specify the directory path to a file:

Table75: Environment Variables for Specifying a Directory Path on Windows

Environment variable	Example Windows Value
<%APPDATA%>	C:\Documents and Settings\jdoe\Application Data
<%windir%>	C:\WINDOWS
<%ProgramFiles%>	C:\Program Files
<%CommonProgramFiles%>	C:\Program Files\Common Files
<%USERPROFILE%>	C:\Documents and Settings\jdoe
<%HOMEDRIVE%>	C:
<%Temp%>	C:\Documents and Settings \<username>\Local Settings\Temp

In a **Custom File** rule for Linux and Solaris, you can use the following environment variables to specify the directory path to a file:

Table76: Environment Variables for Specifying a Directory Path on Linux and Solaris

Environment variable	Example Linux and Solaris Values
<%java.home%>	/local/local/java/j2sdk1.4.1_02/jre
<%java.io.tmpdir%>	/tmp
<%user.dir%>	/home-shared/cknouse
<%user.home%>	/home/cknouse

In a **Custom File** rule for Macintosh, you can use the following environment variables to specify the directory path to a file.

Table77: Environment Variables for Specifying a Directory Path on Macintosh

Environment variable	Example Macintosh Value
<%HOME%>	/Users/*admin* where *admin* is the logged in username
<%USER%>	Maps to the login name of the MAC machine

> NOTE: Although environment variables are formatted in the same way as Toolkit Template directives, they are not interchangeable and you should not confuse them.

Related Documentation

- Specifying Customized Requirements Using Custom Rules

- Implementing Host Checker Policies

Configuring Patch Assessment Policies

You can configure Host Checker policies that check for Windows endpoint's operating system service pack, software version, or desktop application patch version compliance.

Host Checker uses a list of the most current patch versions from the vendor for predefined rules in the Host Checker policy.

You obtain the most current patch version information from a staging site at Pulse Secure. You can manually download and import the current list into the system, or you can automatically import the current list from the Pulse Secure staging site or your own staging site at the specified interval.

Checks can be based on one or more specified products or on specific patches, though not in the same policy. For example, you could check for Internet Explorer version 7 with one policy, and Patch MSOO-039: SSL Certificate Validation Vulnerabilities with a second policy. Then, apply both policies to endpoints at the role or realm level to ensure that the user has the latest browser version with a specific patch. Additionally, you can specify the severity level of patches that you wish to ignore for Microsoft products; for example, you could choose to ignore low or moderate threats.

The system can send remediation instructions (for example, a message describing what patches or software are non-compliant, and a link to where the endpoint can obtain the patch). The system does not auto-remediate in the event of a non-compliant endpoint. However, you can choose to send the items to the client for manual remediation of managed machines.

When an endpoint first connects to the device, the latest versions of the data files and libraries of the IMC are downloaded to the host computer. The initial check takes 10- 20 seconds to run, depending on the link speed. If outdated, these files are automatically updated at subsequent checks. If this is the first time the endpoint has connected to the device with the patch assessment policy, and the connection is a Layer 2 connection, the IMC required to run the Patch Assessment check cannot download. In this case, you should configure a remediation role which displays instructions to direct the user to retry with a Layer 3 connection or contact the administrator.

Note that in non-English installations, the English version of local patches is displayed.

Using Microsoft System Management Server or Microsoft System Center Configuration Manager (SMS/SCCM)

For Windows clients, you can use Microsoft System Management Server or Microsoft System Center Configuration Manager (SMS/SCCM) to provide a method for automatic updates to non-compliant software.

Endpoints configured with SMS/SCCM for software management typically poll the server for updates every fifteen minutes or longer. In a worst-case scenario, clients that are not in compliance with existing Host Checker software requirements may have to wait until the next update interval to login.

Using the SMS/SCCM remediation feature, you can force the client to initiate the software update immediately after the patch assessment check.

If a user attempts to log in, and the endpoint does not have a required software version for compliance with a Host Checker patch assessment policy, Host Checker immediately notifies the client to poll the server for an immediate update. The client receives notification that an SMS/SCCM update has started.

To have SMS/SCCM update the client when notified, set the advertisement time on the SMS/SCCM to As soon as possible.

- The patch assessment policy specifies the required software.

- When an endpoint attempts to authenticate, Host Checker evaluates the client and sends the results back to the system.

- The system evaluates the results and sends reason strings and remediation information to the client, including a message that directs the client to poll the server for software advertisements immediately.

- The SMS/SCCM client queries the SMS/SCCM server for software advertisements.

- The server identifies what patches should be advertised to the client. This information is configured on the server, Host Checker does not interact with the server.

- The SMS/SCCM client receives the advertisement and applies the required patch(es).

You assign clients to a particular group or collection on the server, then SMS/SCCM can advertise patches for that collection. You can configure roles that correspond to collections, and SMS/SCCM can send the appropriate patches for a particular role.

You must have the SMS/SCCM client installed and configured correctly on endpoints, and the SMS/SCCM server must be reachable. In a Layer 2 network, Host Checker is performed before the endpoint is connected to the network. Host Checker can obtain the IP address of the SMS/SCCM server configured for the client. If the endpoint is out of compliance and remediation is necessary, Host Checker pings the server IP address every 15 seconds until the server can be notified to update the client.

It is important as an administrator to inform users of the expected behavior if this feature is enabled, as there is no notification to the user until SMS sends back the advertisement.

Related Documentation

- Configuring Patch Assessment Rules

Configuring Patch Assessment Rules

To configure a predefined patch management rule:

1. In the admin console, select **Authentication > Endpoint Security > Host Checker**.

2. Create a new policy, or click an existing policy in the Policies section of the page.

3. Click the Windows tab.

4. Under Rule Settings, select **Predefined: Patch Management**.

Figure 101 shows the rule settings page for host checker.

Figure101: Rule Settings for Host Checker Policy
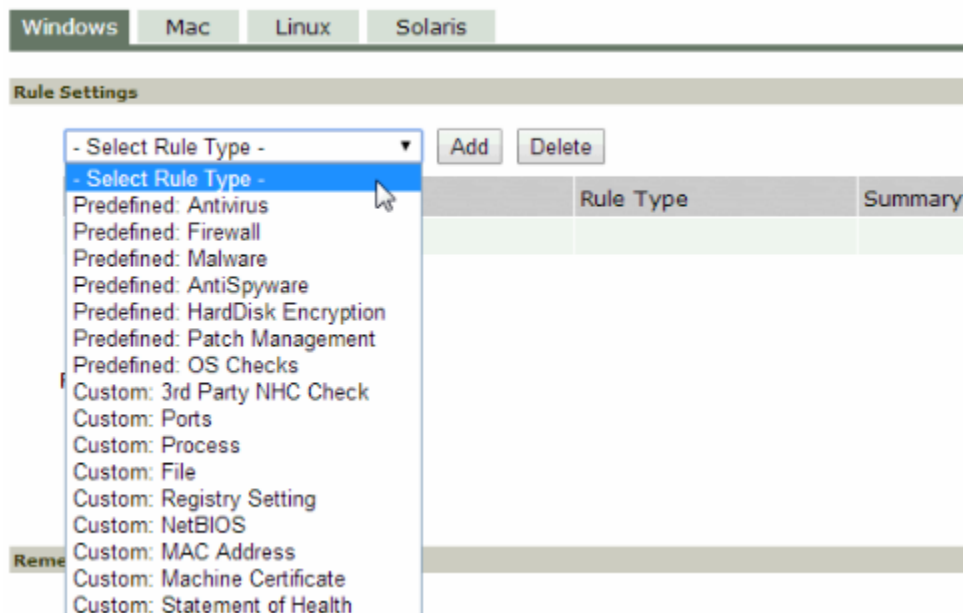
5. Under Rule Settings, click **Add**. The Add Predefined Rule:Patch Management page is displayed.

Figure 102 shows the configuration page to you use to add a patch management rule to the Host Checker policy.

Figure102: Patch Management

6. In the Rule Name box, enter a name for the integrity measurement rule.

> ℹ️ NOTE: If a policy includes a selection that does not apply (for example, if the target software application is not installed on the endpoint), the check for that selection is not performed.

7. Under Criteria, select the product name.

Figure 103 shows the different product names that you can select.

Figure103: Patch Management

8. To automatically enable patch deployment, select **Enable Automatic Patch Deployment**.

> ℹ️ NOTE: Only the SMS/SCCM patch deployment method is used.

9. Click **Save Changes**.

Related Documentation

- Configuring Patch Assessment Policies

- Implementing Host Checker Policies

Using Third-party Integrity Measurement Verifiers

The Trusted Network Connect (TNC) standard enables the enforcement of security requirements for endpoints connecting to networks. The client-side components of the TNC are the IMCs and the TNC-client (TNCC). The TNCC compiles the IMC measurements and sends them to the server. At the server, there is a corresponding set of components: the TNC-server (TNCS) and the IMVs. The TNCS manages the messages between the IMVs and the IMCs and sends the recommendations, based on the IMVs, to the policy engine. This type of rule is available for Host Checker policies on all platforms.

Connect Secure and Host Checker comply with the standards produced by the TNC. For more information about the TNC, IMVs and IMCs, see www.trustedcomputinggroup.org.

You can configure Host Checker to monitor third-party TNC-compliant IMCs installed on client computers. To do so, you must:

1. Run the Third-party Integrity Measurement Verifier (IMV) Server installer on the system designated as the remote IMV server. Install the third-party IMVs and create the server certificates.

2. Specify the remote IMV server so that the system can communicate with it.

3. Implement the Host Checker policy.

Related Documentation

- Configuring a Remote IMV Server

- Implementing the Third-Party IMV Policy

Configuring a Remote IMV Server

ⓘ NOTE:

- In an Active/Passive cluster, the Active/Passive nodes' individual IP addresses must be added to the RIMV as the system IP addresses.

- The successful addition of remote IMV server is not logged in the event log.

- When Host Checker fails, custom instructions are not displayed. There is no user access log on the system about Host Checker failure.

During this step, you install third-party IMVs. Third-party IMVs are installed on the remote IMV server, not on the device.

During this step, you also obtain a server certificate for the remote IMV server. You import the trusted root CA certificate of the CA that generated the server certificate onto the device. The system then authenticates with the remote IMV server through the certificate. If you do not have a certificate authority, install and use OpenSSL to generate a CA certificate.

To install, configure, and implement the server software:

1. In the system admin console, choose **Maintenance > System > Installers** and download the Third-party Integrity Measurement Verifier (IMV) Server installer.

2. Run the installer on the system designated as the remote IMV server.

3. Install the third-party IMVs on the remote IMV server and the corresponding IMCs on the client systems.

4. Generate a server certificate from a certificate authority for the remote IMV server. The server's certificate Subject CN value must contain the actual hostname or IP address of the remote IMV server.

   The server certificate and the private key must be combined into a single PKCS#12 file and encrypted with a password.

   If you do not have a certificate authority, you can use the following steps to create a CA and then create a server certificate for the remote IMV server.

   ⓘ NOTE: Install the full version of OpenSSL. The "light" version of OpenSSL will not work.

   Follow the steps below to set up OpenSSL:

1. Download and install OpenSSL from this site:

   http://www.slproweb.com/products/Win32OpenSSL.html

2. At the Windows command prompt, type the following commands:

   cd \openssl
   md certs
   cd certs
   md demoCA
   md demoCA\newcerts
   edit demoCA\index.txt

3. Press the **ALT-F** keys and then the **S** key to save the file.

4. Press the **ALT-F** keys and then the **X** key to exit the editor.

5. At the Windows command prompt, type the following commands:

   edit demoCA\serial

6. Type the following in the document window: **01**

7. Press the **ALT-F** keys and then the **S** key to save the file.

8. Press the **ALT-F** keys and then the **X** key to exit the editor.

9. At the Windows command prompt, type the following commands:

   set path=c:\openssl\bin;%path%

Follow the steps below to create a CA key:

1. To create a CA key, type the following command at the Windows command prompt in the c:\openssl\certs directory:

   openssl genrsa -out ca.key 1024

   The following output should appear:

   ```
   Loading 'screen' into random state - done

   Generating RSA private key, 1024 bit long modulus

   ........+++++

   .+++++

   e is 65537 (0x10001
   ```

Follow the steps below to create a CA Certificate:

1. Type the following command at the Windows command prompt in the c:\openssl\certs directory:

   openssl req -new -x509 -days 365 -key ca.key -out
   demoCA/cacert.pem

2. Enter the appropriate Distinguished Name (DN) information for the CA certificate. You can leave some fields blank by entering a period.

   For example:

   Country Name: US
   State or Province Name: CA
   Locality Name: Sunnyvale
   Organization Name: XYZ
   Org. Unit Name: IT
   Common Name: ic.xyz.com
   Email Address: user@xyz.com

3. To set up the CA, type the following command at the Windows command prompt in the directory c:\openssl\certs:

   copy ca.key demoCA
   notepad demoCA.cnf

4. When prompted to create a new file, press the **yes** button.

5. Type the following lines in the document, pressing the Enter key at the end of each line.

   [ca]
   default_ca = demoCA
   [demoCA]
   dir = ./demoCA
   database = $dir/index.txt
   new_certs_dir = $dir/newcerts
   certificate = $dir/cacert.pem
   serial = $dir/serial
   private_key = $dir/ca.key
   default_days = 365
   default_md = md5
   policy = policy_any
   email_in_dn = no
   name_opt = ca_default
   name_opt = ca_default
   copy_extensions = none
   [policy_any]
   countryName = supplied
   stateOrProvinceName = optional
   organizationName = optional
   organizationalUnitName = optional
   commonName = supplied
   emailAddress = optional

6. Save the file and close notepad.

7. Type the following command to generate an RSA private key for the remote IMV server:

   openssl genrsa –out rimvs_key.pem 1024

8. Type the following command to generate a CSR for the remote IMV server:

   openssl req –new –key rimvs_key.pem –out rimvs_csr.pem

9. Type the following lines:

   Country Name:
   State or Province Name:
   Locality Name:
   Organization Name:
   Organizational Unit Name:
   Common Name: [IPAddress]
   Email Address:
   A challenge password:
   An optional company name:

   You may enter any value you like for most fields, but the Common Name field must contain the IP address of the machine running the remote IMV server. This machine should have a static IP address.

10. Type the following command to generate a certificate for the remote IMV server:

    openssl ca –config demoCA.cnf –in rimvs_csr.pem –out rimvs_cert.pem

11. Type 'y' twice when prompted to generate the certificate. This certificate is valid for 365 days by default. If you want a different certificate lifetime, change the default_days parameter in the demoCA.cnf file, or use the – days parameter to the openssl ca command to specify a different lifetime.

12. Type the following command to place the remote IMV server key and certificate in a PKCS#12 file (substitute your password):

    openssl pkcs12 –export –in rimvs_cert.pem –inkey rimvs_key.pem –passout
    pass:<password>-out rimvs_p12.pem

13. On the remote IMV server, choose **Programs > Pulse Secure > Remote IMV Server > Remote IMV Server Configurator** from the Start menu.

14. Under Client Info, click **Add**.

15. Configure the port to service SOAP requests.

16. Enter the client's IP address, the number of addresses to use, and the shared secret used by both the system and the remote IMV server.

17. Change logging settings if you choose (log is generated in the install directory).

18. Browse and find the PKCS#12 file you generated in the filesystem.

19. Specify the password associated with the certificate.

20. In the Connect Secure admin console, use the **System > Configuration > Certificates > Trusted Server CAs** tab to import the trusted root CA certificate of the CA that issued the certificate for the remote IMV server.

    If you used OpenSSL to generate the Remote IMV Server's server certificate is: demoCA\cacert.pem.

    If you did not use OpenSSL to generate this certificate, ensure that the file you import has the CA certificate (not the root certificate).

21. Click **Import Trusted Server CA** and browse for the server certificate used on the remote IMV server.

22. Add the new remote IMV server:

    To specify the remote IMV server so that Connect Secure can communicate with it:

    1. In the admin console, choose **Authentication > Endpoint Security > Host Checker**.

    2. Under Remote IMV, click **New Server**.

    3. In the New Server page:

       1. Create a label for the server using the **Name** and (optional) **Description** fields.

       2. In the **Hostname** field, enter either the IP address or hostname as defined in the server certificate.

       3. In the **Port** field, enter the unique port number the system uses to communicate with the remote IMV server. Ensure that no other service is using this port number.

          The default port number is the same as the default https port number. If you are running a web server on the same system as the Remote IMV Server, enter a new port number in the Port field.

       4. In the **Shared Secret** field, enter the same shared secret used in the client information entry on the remote IMV server.

       5. Click **Save Changes**.

    6. Under Remote IMV, click **New IMV** to specify the third-party IMV.

    7. In the New IMV page:

       1. Create a label for the IMV using the **Name** and (optional) **Description** fields.

       2. In the **IMV Name** field, enter the name of the IMV. This name must match the "human readable name" in the IMV's well-known registry key on the remote IMV server. For more information about human readable names and the well-known registry key, see www.trustedcomputinggroup.org.

       3. From the Primary Server pop-up menu, select the remote IMV server where this IMV is installed.

4. (Optional) From the Secondary Server pop-up menu, select the secondary remote IMV server where this IMV is installed. The secondary server acts as a failover in case the primary server becomes unavailable.

   The system continues to try to re-establish connection to the primary remote IMV Server, and uses the primary Remote IMV Server on subsequent handshakes once it becomes available.

5. Click **Save Changes**.

6. Click **Save Changes**.

Related Documentation

- Implementing the Third-Party IMV Policy

- Using Third-party Integrity Measurement Verifiers

Implementing the Third-Party IMV Policy

To use Host Checker as a policy enforcement tool for managing endpoints, you must create global Host Checker policies at the system level through the Authentication > Endpoint Security > Host Checker page of the admin console, and then implement the policies at the realm and role levels.

> ℹ️ NOTE: The **Custom: Remote IMV** option does not appear until you add the Remote IMV New Server and New IMV on the main Host Checker page.

To implement the third-party IMV policy:

1. In the admin console, select **Authentication > Endpoint Security > Host Checker**.

2. Under Policies, click **New**.

3. Enter a name in the **Policy Name** field and then click **Continue**. (Users see this name on the Host Checker remediation page if you enable custom instructions for this policy.)

4. Under Rule Settings, choose **Custom: Remote IMV** and click **Add**.

5. In the Add Custom Rule: Remote IMV page:

   1. In the **Rule Name** field, enter an identifier for the rule.

   2. Under **Criteria**, select the third-party IMV to be associated with this rule.

   3. Click **Save Changes**.

4. Specify how Host Checker should evaluate multiple rules within the policy.

5. (Recommended) Specify remediation options for users whose computers do not meet the requirements specified in the policy

6. Click **Save Changes**.

7. Implement the policy at the realm or role level.

Related Documentation

- Using Third-party Integrity Measurement Verifiers

- Configuring a Remote IMV Server

- Implementing Host Checker Policies

Implementing Host Checker Policies

After you create global policies through the Authentication > Endpoint Security > Host Checker page of the admin console, you can restrict the system and resource access by requiring Host Checker in a:

- **Realm authentication policy**—When administrators or users try to sign in to the device or launch a Virtual Workspace session, the system evaluates the specified realm's authentication policy to determine if the pre-authentication requirements include Host Checker. You can configure a realm authentication policy to download Host Checker, launch Host Checker and enforce Host Checker policies specified for the realm, or not require Host Checker. The user must sign in using a computer that adheres to the Host Checker requirements specified for the realm. If the user's computer does not meet the requirements, the system denies access to the user unless you configure remediation actions to help the user bring his computer into compliance. You can configure realm-level restrictions through the Administrators > Admin Realms > *SelectRealm* > Authentication Policy > Host Checker page or the Users > User Realms > *SelectRealm* > Authentication Policy > Host Checker page of the admin console.

- **Role**—When the system determines the list of eligible roles to which it can map an administrator or user, it evaluates each role's restrictions to determine if the role requires that the user's computer adheres to certain Host Checker policies. If it does and the user's computer does not follow the specified Host Checker policies, then the system does not map the user to that role unless you configure remediation actions to help the user bring his computer into compliance. You can configure role-mapping using settings in the Users > User Realms > SelectRealm > Role Mapping page. You can configure role-level restrictions through the Administrators > Admin Roles > SelectRole > General > Restrictions > Host Checker page of the admin console or the Users > User Roles> SelectRole > General > Restrictions > Host Checker page. If you have enabled Advanced Endpoint Defense Malware Protection, you can select to implement this feature for any role.

- **Resource policy**—When a user requests a resource, the system evaluates the resource policy's detailed rules to determine if the resource requires that the user's computer adheres to certain Host Checker policies. The system denies access to the resource if the user's computer does not follow the specified Host Checker policies unless you configure remediation actions to help the user bring his computer into compliance. To implement Host Checker restrictions at the resource policy level, use settings in the Users > Resource Policies > *SelectResource* > *SelectPolicy* > Detailed Rules page.

You may specify that the system evaluate your Host Checker policies only when the user first tries to access the realm, role, or resource that references the Host Checker policy. Or, you may specify that the system periodically re-evaluate the policies throughout the user's session. If you choose to periodically evaluate Host Checker policies, the system dynamically maps users to roles and allows users access to new resources based on the most recent evaluation.

Executing Host Checker Policies

When the user tries to access the system, Host Checker evaluates its policies in the following order:

1. **Initial evaluation**—When a user first tries to access the system sign-in page, Host Checker performs an initial evaluation. Using the rules you specify in your policies, Host Checker verifies that the client meets your endpoint requirements and returns its results to the system. Host Checker performs an initial evaluation regardless of whether you have implemented Host Checker policies at the realm, role, or resource policy level.

If the user navigates away from the system sign-in page after Host Checker starts running but before signing in, Host Checker continues to run on the user's machine until the Host Checker process times out.

If the system does not receive a result from Host Checker for any reason (including because the user manually terminated Host Checker), it displays an error and directs the user back to the sign-in page.

Otherwise, if the Host Checker process returns a result, the system goes on to evaluate the realm level policies.

2. **Realm-level policies**—The system uses the results from Host Checker's initial evaluation to determine which realms the user may access. Then, the system displays or hides realms from the, only allowing the user to sign into those realms that you enable for the sign-in page, and if the Host Checker requirements for each realm are met. If the user cannot meet the Host Checker conditions required by any of the available realms, the system does not display the sign-in page. Instead, it displays an error stating the user has no access unless you have configured remediation actions to help the user bring the endpoint into compliance.

   Note that Host Checker only performs realm-level checks when the user first signs in. If the state of the user's system changes during his session, the system does not remove him from the current realm or allow him access to a new realm based on his new system state.

3. **Role-level policies**—After the user signs into a realm, the system evaluates role-level policies and maps the user to the role or roles if he meets the Host Checker requirements for those role(s). Then, the system displays the homepage to the user and enables those options that the mapped role(s) allow.

   If Host Checker returns a different status during a periodic evaluation, the system dynamically remaps the user to roles based on the new results. If the user loses rights to all available roles during one of the periodic evaluations, the system disconnects the user's session unless you have configured remediation actions to help the user bring the endpoint into compliance.

4. **Resource-level policies**—After allowing the user to access the homepage, the user may try to access a resource that is controlled by a resource policy. When he does, the system determines whether or not to perform the action specified in the resource policy based on the last status returned by Host Checker.

   If Host Checker returns a different status during a periodic evaluation, the new status only impacts new resources that the user tries to access. For example, if the user successfully initiates a VPN Tunneling session and then fails his next resource-level host check, he may continue to access the open VPN Tunneling session. The system only denies him access if he tries to open a new VPN Tunneling session. The system checks the last status returned by Host Checker whenever the user tries to access a new Web resource or open a new Secure Application Manager, VPN Tunneling, or Secure Terminal Access session.

With either a success or fail result, Host Checker remains on the client. Windows users may manually uninstall the agent by running uninstall.exe in the directory where Host Checker is installed. If you enable client-side logging through the System > Log/Monitoring > Client Logs page, this directory also contains a log file, which the system rewrites each time Host Checker runs.

If you enable dynamic policy evaluation for Host Checker, the system evaluates resource policies implemented at the realm level whenever a user's Host Checker status changes. If you do not enable dynamic policy evaluation for Host Checker, it does not evaluate resource policies but it does evaluate the authentication policy, role mapping rules, and role restrictions whenever a user's Host Checker status changes.

Related Documentation

- Configuring Host Checker Restrictions

- Specifying General Host Checker Options

- Role Restrictions

- Defining Authentication Access Policies

Configuring Host Checker Restrictions

To specify Host Checker restrictions:

1. Navigate to: **Authentication > Endpoint Security > Host Checker** and specify global options for Host Checker to apply to any user for whom Host Checker is required in an authentication policy, a role mapping rule, or a resource policy.

2. If you want to implement Host Checker at the realm level:

    1. Navigate to:

        - **Administrators > Admin Realms >** *Select Realm* **> General > Restrictions > Host Checker**.

        - **Users > User Realms >** *Select Realm* **> General > Restrictions > Host Checker**.

    2. Choose one of the following options for either all available policies or for individual policies listed in the **Available Policies** column:

        - **Evaluate Policies**—Evaluates without enforcing the policy on the client and allows user-access. This option does not require Host Checker to be installed during the evaluation process; however, Host Checker is installed once the user signs in to the system.

        - **Require and Enforce**—Requires and enforces the policy on the client in order for the user to log in to the specified realm. Requires that Host Checker is running the specified Host Checker policies in order for the user to meet the access requirement. This option requires the system to download Host Checker to the client machine. If you choose this option for a realm's authentication policy, then the system downloads Host Checker to the client machine after the user is authenticated and before the user is mapped to any roles in the system. Selecting this option automatically enables the Evaluate Policies option.

    3. Select the **Allow access to realm if any ONE of the selected "Require and Enforce" policies is passed** check box if you do not want to require users to meet all of the requirements in all of the selected policies. Instead, the user can access the realm if he meets the requirements of any one of the selected Host Checker policies. Note that Cache Cleaner policies are not part of the "requirement" decision process. Users can access the realm as long as they meet the other requirements regardless of whether they meed the Cache Cleaner policy.

4. If you want to implement Host Checker at the role level:

    1. Navigate to:

        - **Administrators > Admin Roles >** *Select Role* **> General > Restrictions > Host Checker**.

        - **Users > User Roles >** *Select Role* **> General > Restrictions > Host Checker**.

    2. Choose one of the following options:

        - **Allow all users** — Does not require Host Checker to be installed in order for the user to meet the access requirement.

- **Allow only users whose workstations meet the requirements specified by these Host Checker policies** — Requires that Host Checker is running the specified Host Checker policies in order for the user to meet the access requirement.

- Select the **Allow access to role if any ONE of the selected "Require and Enforce" policies is passed** check box if you do not want to require users to meet all of the requirements in all of the selected policies. Instead, the user can access the role if he meets the requirements of any one of the selected Host Checker policies.

3. If you want to create role-mapping rules based on a user's Host Checker status:

   1. Navigate to: **Users > User Realms >** *Select Realm* **>Role Mapping**.

   2. Click **New Rule**, select **Custom Expressions** from the Rule based on list, and click **Update**. Or, to update an existing rule, select it from the **When users meet these conditions** list.

   3. Click **Expressions**.

   4. Write a custom expression for the role mapping rule to evaluate Host Checker's status using the hostCheckerPolicy variable. For help writing the custom expressions, use tips in the Expressions Dictionary.

   5. In the **...then assign these roles** section, select the roles to map users to when they meet the requirements specified in the custom expression and click **Add**.

   6. Select the **Stop processing rules when this rule matches** if you want to stop evaluating role mapping rules if the user successfully meets the requirements defined in this rule.

7. If you want to implement Host Checker at the resource policy level:

   1. Navigate to: **Users > Resource Policies > Select Resource > Select Policy > Detailed Rules**.

   2. Click **New Rule** or select an existing rule from the Detailed Rules list.

   3. Write a custom expression for the detailed rule to evaluate Host Checker's status using the hostCheckerPolicy variable.

These options allow you to control which version of an application or service runs on client machines.

Remediating Host Checker Policies

You can specify general remediation actions that you want Host Checker to take if an endpoint does not meet the requirements of a policy. For example, you can display a remediation page to the user that contains specific instructions and links to resources to help the user bring their endpoint into compliance with Host Checker policy requirements.

You can also choose to include a message to users (called a reason string) that is returned by Host Checker or an integrity measurement verifier (IMV) and explains why the client machine does not meet the Host Checker policy requirements.

For example, the user may see a remediation page that contains the following custom instructions, a link to resources, and reason strings:

Your computer's security is unsatisfactory.

Your computer does not meet the following security requirements. Please follow the instructions below to fix these problems. When you are done click Try Again. If you choose to **Continue** without fixing these problems, you may not have access to all of your intranet servers.

1. Symantec

Instructions: You do not have the latest signature files. **Click here to download the latest signature files**. Reasons: The AntiVirus Product Version is too low.

The age of the Virus Definitions is not acceptable.

For each Host Checker policy, you can configure two types of remediation actions:

- **User-driven**—Using custom instructions, you can inform the user about the failed policy and how to make his computer conform. The user must take action to successfully re-evaluate the failed policy. For instance, you can create a custom page that is linked to a policy server or Web page and enables the user to bring his computer into compliance.

- **Automatic (system-driven)**—You can configure Host Checker to automatically remediate the user's computer. For example, when the initial policy fails, you can kill processes, delete files, or allow automatic remediation by an IMV. On Windows, you can also call the HCIF_Module.Remediate () API function as part of a third-party J.E.D.I. DLL. Host Checker does not inform users when performing automatic actions. (You could, however, include information in your custom instructions about the automatic actions.)

General Host Checker Remediation User Experience

Users may see the remediation page in the following situations:

- Before the user signs in:

  - If you enable custom instructions for a policy that fails, the system displays the remediation page to the user. The user has two choices:

    - Take the appropriate actions to make the endpoint conform to the policy and then click the Try Again button on the remediation page. Host Checker checks the user's computer again for compliance with the policy.

    - Leave the endpoint in its current state and click the Continue button to sign in. The user cannot access the realm, role, or resource that requires compliance with the failed policy.

      If you do not configure the system with at least one realm that allows access without enforcing a Host Checker policy, the user must bring the endpoint into compliance before signing in.

  - If you do not enable custom instructions for a policy that fails, Host Checker does not display the remediation page to the user. Instead, the system displays the sign-in page but does not allow the user to access any realms, roles, or resources that have a failed Host Checker policy.

- After the user signs in:

  - (Windows only) During a session, if a user's Windows computer becomes non-compliant with the requirements of a Host Checker policy, an icon appears in the system tray along with a pop-up message that informs the user of the non-compliance. The user can then click the pop-up message to display the remediation page.

  - (Macintosh or Linux) During a session, if a user's Macintosh or Linux computer becomes non-compliant with the requirements of a Host Checker policy, the system displays the remediation page to inform the user of the non-compliance.

    > *i* NOTE: If the user hides the remediation page by setting a user preference, he may only continue using the secure gateway if you configure other realms and roles that do not enforce a Host Checker policy.

Related Documentation

Configuring General Host Checker Remediation

To specify remediation actions for a Host Checker policy:

1. In the admin console, select **Authentication > Endpoint Security > Host Checker**.

2. Create or enable Host Checker policies.

3. Specify the remediation actions that you want Host Checker to perform if a user's computer does not meet the requirements of the current policy:

   - **Enable Custom Instructions**—Enter the instructions you want to display to the user on the Host Checker remediation page. You can use the following HTML tags to format text and add links to resources such as policy servers or web sites: <i>, <b>, <br>, <font>, and <a href>. For example:

   You do not have the latest signature files.

   Click here to download the latest signature files.

   > NOTE: For Windows clients, if you include in the instructions a link to a system-protected policy server, define a pre-authentication access tunnel.

   - **Enable Custom Actions**—You can select one or more alternate policies that you want Host Checker to evaluate if the user's computer does not meet the current policy requirements. The alternate policy must be a third-party policy that uses a J.E.D.I. package. For example, you can use a J.E.D.I. package to launch an application if the user's computer does not meet the current policy requirements. Select the alternate policy in the HC Policies list and then click Add.

   - Remediate—(Third party DLLs only) You can select this option to perform remediation actions specified by means of the Remediate () API function in a third-party J.E.D.I. DLL.

   > NOTE: The Remediate feature is primarily provided for backwards compatibility. We recommend that you use IMCs and IMVs instead.

   - **Kill Processes**—On each line, enter the name of one or more processes you want to kill if the user's computer does not meet the policy requirements. You can include an optional MD5 checksum for the process. (You cannot use wildcards in the process name.) For example:

   keylogger.exe

   MD5: 6A7DFAF12C3183B56C44E89B12DBEF56

   - **Delete Files**—Enter the names of files you want to delete if the user's computer does not meet the policy requirements. (You cannot use wildcards in the file name.) Enter one file name per line. For example:

   c:\temp\bad-file.txt

   /temp/bad-file.txt

   - **Send reason strings**—Select this option to display a message to users (called a reason string) that is returned by Host Checker or integrity measurement verifier (IMV) and

explains why the client machine does not meet the Host Checker policy requirements. This option applies to predefined rules, custom rules, and to third-party IMVs that use extensions in the Pulse Secure TNC SDK. For example, an antivirus IMV might display the following reason string:

The AntiVirus Product Version is too low. The age of the Virus Definitions is not acceptable.

> NOTE: By sending reason strings, you are disclosing to users what the IMV is checking on the client machine.

4. Click **Save Changes**.

Related Documentation

- Remediating Host Checker Policies

Upgrading the Endpoint Security Assessment Plug-In

The Endpoint Security Assessment Plug-in (ESAP) on Connect Secure checks third-party applications on endpoints for compliance with the predefined rules you configure in a Host Checker policy. This plug-in is included in the system software package.

Pulse Secure frequently adds enhancements, bug fixes, and support for new third-party applications to the plug-in. New plug-in releases are available independently and more frequently than new releases of the system software package. If necessary, you can upgrade the plug-in independently of upgrading the system software package.

You can upload up to four versions of the plug-in to your system, but it uses only one version at a time (called the active version). If necessary, you can rollback to a previously active version of the plug-in.

To upgrade the Endpoint Security Assessment Plug-in:

1. Download the Endpoint Security Assessment Plug-in from the Pulse Secure Global Support Center (PSGSC) Center to your computer:

   1. Open the following page:

      http://www.pulsesecure.net/support

   2. Click the Software tab.

   3. Navigate to the ESAP release you want and click the link to download the package file to your computer.

4. Select **Authentication > Endpoint Security > Host Checker**.

5. At the bottom of the Host Checker page under Manage Endpoint Security Assessment Plug-In Versions:

   1. If you have previously uploaded four versions of the component software, you must delete one of the versions before you can upload another one. Select the version you want to delete and click **Delete**.

   2. If you want the system to actively begin using the new component software immediately after you upload it, select the Set as active after upload option.

   3. Click **Browse**, select the plug-in file you want to upload to the system, and click **OK**.

   4. Click **Upload**. While the system uploads and decrypts the plugin .zip file, the message "Loading..." appears in the plug-in list under Manage Endpoint Security Assessment Plug-In Versions. If the device is a member of a cluster, it displays the message

"Loading..." while the plug-in is transferred to the other cluster nodes. After the plug-in is installed, the date and time of the plug-in installation appears in the plug-in list.

5. If you did not select the Set as active after upload option, activate the plug-in you want to use by selecting the version in the plug-in list and clicking Activate.

> ℹ️ NOTE:
>
> - If you attempt to activate a version of the plug-in that does not support all of the predefined rules already configured in all Host Checker policies, the system does not allow activation of that plug-in version. For example, if a Host Checker policy is configured to use a predefined rule to check for a version of antivirus software, and you attempt to activate a plug-in version that does not support that particular version of the antivirus software, the system does not allow you to activate that plug-in version. To view the list of supported products for a particular plug-in version, click the plug-in's version number under Manage Endpoint Security Assessment Plug-In Versions.
>
> - You can rollback to an older plug-in version after upgrading to a later version by selecting the older version as the active version. But, if you modified any Host Checker policies after upgrading to the later version, the rollback may not succeed. Rollback is guaranteed to succeed only if the policies did not change.
>
> - If you upgrade the system software to a newer version, or you import a user configuration file, the currently active plug-in version does not change. If you want to use a different plug-in version after upgrading or importing a user configuration file, you must manually activate that plug-in version.
>
> - If the system already has four versions of the plug-in installed when you upgrade the system software to a newer version, it automatically deletes the oldest plug-in version and installs, but does not activate, the plug-in included with the new system software.

Related Documentation

- Implementing Host Checker Policies

Defining Host Checker Pre-Authentication Access Tunnels

If your policies require Host Checker rules or third-party J.E.D.I. DLLs to access a policy server (or other resource) to check compliance before users are authenticated, you can use one of the following methods to make the resource available to the Host Checker Windows clients:

- **Deploy the policy server in a DMZ where Host Checker rules or third-party J.E.D.I. DLLs can access the server directly instead of going through Connect Secure**—This deployment is the simplest solution because you do not have to define a Host Checker pre-authentication access tunnel through Connect Secure between clients and the policy server.

- **Deploy the policy server in a protected zone behind Connect Secure (Windows only)**— This deployment requires you to define a pre-authentication access tunnel. A pre-authentication access tunnel enables Host Checker rules or third-party J.E.D.I. DLLs to access the protected policy server or resource before the system authenticates users. To define a pre-authentication access tunnel, you associate a loopback address (or hostname) and port on the client with an IP address and port on the policy server. You add one or more tunnel definitions to a MANIFEST.HCIF file, which you then upload to Connect Secure. You can upload multiple MANIFEST.HCIF files to Connect Secure. For all third-party policies enabled on a realm, Host Checker creates tunnels for all of the tunnel definitions in all of the MANIFEST.HCIF files, assuming the definitions are unique.

While running on a Windows client, Host Checker listens for a connection on each loopback address and port you specify in the tunnel definitions. The connections can originate from the integrated Host Checker rules and from client-side or server-side J.E.D.I. DLLs. Host Checker uses the pre-authentication access tunnel(s) to forward the connections through Connect Secure to the policy server(s) or other resource.

Figure104: Host Checker Creates a Tunnel from a Client to a Policy Server Behind Connect Secure

Host Checker relays connections from a Host Checker policy or third party J.E.D.I. DLL to a protected policy server through a pre-authentication access tunnel.

Host Checker listens on the loopback address and port for connections from the Host Checker policy or third-party J.E.D.I. DLL.
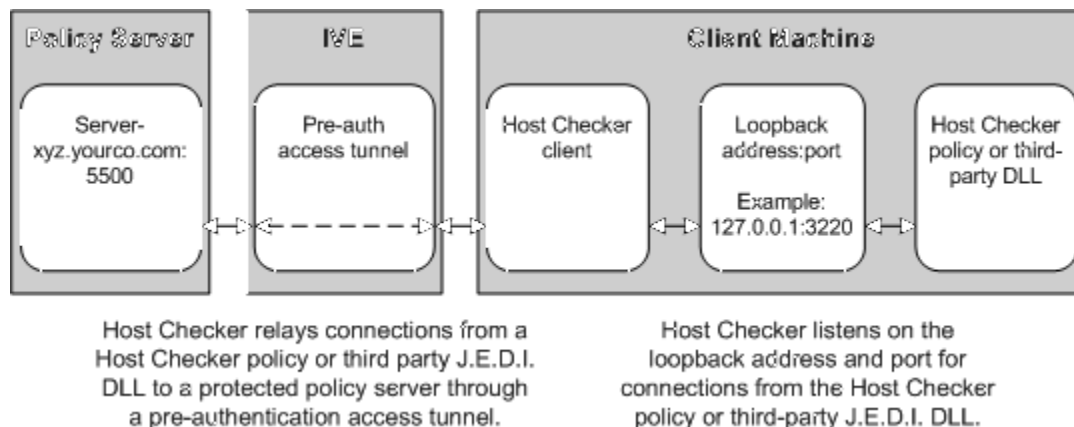
NOTE: Host Checker pre-authentication access tunnels are supported on Windows only.

Related Documentation

- Specifying Host Checker Pre-Authentication Access Tunnel Definitions

Specifying Host Checker Pre-Authentication Access Tunnel Definitions

For Windows clients, you can define a pre-authentication access tunnel that enables Host Checker methods or third-party J.E.D.I. DLLs to access a protected policy server (or other resource) before users are authenticated.

A definition for a Host Checker pre-authentication access tunnel configures access to one policy server or other resource. Each tunnel definition consists of a pair of IP addresses and ports: one loopback IP address and port on the client, and one IP address and port on the policy server.

You specify one or more tunnel definition(s) in a Host Checker policy package definition file. The package definition file, which must be named MANIFEST.HCIF, defines the name of an interface DLL, the Host Checker policies defined in the DLL, and the pre-authentication access tunnel definitions. Note that if you do not include policies in your package, Host Checker simply enforces that the package has run on the client. If you do declare policies through this file, they become available through the admin console where you can implement them at the realm, role, and resource policy levels.

Within the MANIFEST.HCIF file, you must include one definition per line, with a blank line between each definition, using the following format:

HCIF-Main: <DLLName>
HCIF-Policy: <PolicyName>
HCIF-IVE-Tunnel: <client-loopback>:port <policy-server>:port

where:

<DLLName> is the name of the interface DLL, such as myPestPatrol.dll. Even if you are not using an interface DLL, you must include a dummy DLL as a placeholder file that has this exact name.

<PolicyName> is the name of a policy defined in the DLL, such as myFileCheck. You can define multiple policies by using the HCIF-Policy statement for each policy. If you are not using an interface DLL, you can use any policy name as a placeholder.

The syntax of a Host Checker tunnel definition is:

HCIF-IVE-Tunnel: <client-loopback>:port <policy-server>:port

where:

<client-loopback> is a loopback address that begins with 127. and takes any of the following forms:

- An IP address and port that takes the form of 127.*.*.*:port. To avoid conflicts with JSAM, do not use 127.0.0.1 with port 80, but you can use 127.0.0.1 with other ports. For example: 127.0.0.1:3220

- A hostname that resolves to a loopback address that begins with 127. You can use a local hosts file on each client computer or a DNS server to resolve the loopback address.

- A hostname that does not resolve to a loopback address, or resolves to a non-loopback address. In these cases, Host Checker allocates a loopback address and updates the local hosts file on the client with the mapping. Note that the user must have administrator privileges in order for Host Checker to modify the local hosts file. If the user does not have administrator privileges, Host Checker cannot update the hosts file and cannot open the pre-authentication access tunnel. In that case, Host Checker logs an error.

<policy-server> is the IP address or hostname of the back-end policy server. Connect Secure resolves the hostname you specify.

For example, in the following tunnel definition, 127.0.0.1:3220 is the client loopback address and port, and mysygate.company.com:5500 is the policy server hostname and port:

HCIF-IVE-Tunnel: 127.0.0.1:3220 mysygate.company.com:5500

Or you can use a hostname for the client, as in this example:

HCIF-IVE-Tunnel: mysygate.company.com:3220 mysygate.company.com:5500

Keep the following in mind when specifying tunnel definitions:

- You must add a blank line between each line in the MANIFEST.HCIF file, and you can use a semi-colon at the beginning of a line to indicate a comment. For example:

  HCIF-Main: myPestPatrol.dll
  HCIF-Policy: myFileCheck
  HCIF-Policy: myPortCheck
  ; Tunnel definitions
  HCIF-IVE-Tunnel: 127.0.0.1:3220 mysygate.company.com:5500
  HCIF-IVE-Tunnel: 127.1.1.1:3220 mysygate2.company.com:5500
  HCIF-IVE-Tunnel: mysygate.company.com:3220 mysygate3.company.com:5500

- Host Checker pre-authentication access tunnels are supported on Windows only.

- If <client-loopback> is a non-loopback address, then Host Checker cannot open the pre-authentication access tunnel and logs an error instead.

Related Documentation

- Defining Host Checker Pre-Authentication Access Tunnels

Specifying General Host Checker Options

You can specify global options for Host Checker that apply to any user for whom Host Checker is required in an authentication policy, a role mapping rule, or a resource policy.

To specify general Host Checker options:

1. In the admin console, choose **Authentication > Endpoint Security > Host Checker**.

2. Under Options:

   - In the **Perform check every X minutes** field, specify the interval at which you want Host Checker to perform policy evaluation on a client machine. If the client machine fails to meet the requirements of the Host Checker policies required by a role or resource policy, then the system denies the associated user requests.

     For example, you may require that a user runs a specific third-party antivirus application in order to map to Role A, which enables network connections from an external location. If the user's client machine is running the required antivirus application when the user signs in, then the user maps to Role A and is granted all access features enabled for Role A. If the antivirus application stops running during the user session, however, the next time Host Checker runs, the user fails to meet the security requirements for Role A and therefore loses all access privileges for Role A.

     When an end-user logs into a Realm, Host Checker performs an initial policy check, regardless of whether or not the policy is enforced at the Realm, Role, and/or Resource level. The initial policy check establishes a start time. Host Checker evaluates policies at the frequency set by the Perform check every X minutes option starting the clock at the initial policy check. Although the frequency setting is set globally for all Host Checker policy checking, it is not synchronized for all end-user clients connected to the system. Each client performs its own initial policy check and starts its own X minute countdown. If you configure the authentication policy within a realm where Host Checker enforces policies (versus installing), the enforcement occurs only during the pre-authentication phase. After an end-user signs in and for the duration of the user's session, any subsequent Host Checker policy checks have no impact on realm access, meaning that there is no concept of removing an end-user session from a realm once an end-user successfully authenticates into that realm.

     If you configure a role restriction where Host Checker enforces policies, the enforcement occurs just after authentication during role mapping. Role restrictions are enforced periodically during the end-user session at an interval specified using the Host Checker frequency setting. If the end-user successfully passes the Host Checker evaluation during role mapping but later fails X minutes after login, that specific user loses rights to that role. If the end-user loses rights to all available roles due to Host Checker policy evaluation, the end-user session is disconnected.

     If you configure a resource-based policy rule where Host Checker enforces policies, the enforcement occurs when the end-user attempts to access the resource/backend server. For web resources, the Host Checker evaluation occurs at each request. For SAM and STA resources, the Host Checker evaluation occurs when the system activates the connection to the backend application/server. For VPN Tunneling access, the Host Checker evaluation occurs when the system initiates VPN Tunneling. Existing connections of applications running by way of SAM, Telnet/SSH connection, and VPN Tunneling connections are not affected by further Host Checker evaluations. Only new Web requests, new applications across SAM, new instances of STA, and launching VPN Tunneling are affected. The Host Checker evaluation is based on the most recent policy

check that occurred X minutes ago. Example, if you configure the frequency setting to Perform check every five minutes and the end-user attempts to access a protected resource or attempts to launch VPN Tunneling four minutes after the last check, then the policy evaluation is based on the state of the client machine four minutes ago, not at the moment the end-user attempted to access the resource.

> NOTE: If you enter a value of zero, Host Checker only runs on the client machine when the user first signs in.

- For the Client-side process, login inactivity timeout option, specify an interval to control timing out in the following situations:

  - If the user navigates away from the sign-in page after Host Checker starts running but before signing in to the device, Host Checker continues to run on the user's machine for the interval you specify.

  - If the user is downloading Host Checker over a slow connection, increase the interval to allow enough time for the download to complete.

- Select Perform dynamic policy reevaluation to automatically refresh the roles of individual users by enabling dynamic policy evaluation for Host Checker. Host Checker can trigger the system to evaluate resource policies whenever a user's Host Checker status changes. (If you do not select this option, the system does not evaluate resource policies but it does evaluate the authentication policy, role mapping rules, and role restrictions whenever a user's Host Checker status changes.)

3. Click **Save Changes**.

Related Documentation

- Configuring Host Checker Restrictions

Specifying Host Checker Installation Options

If you implement any policy at the realm, role, or resource policy level that requires Host Checker, you must provide a mechanism by which the system or the user can install Host Checker on the client machine. Otherwise, when the system evaluates the Host Checker policy, the user's machine fails because the Host Checker client is not available to return a success status.

You can use two methods to install Host Checker on a user's system:

- Connect Secure automatically installs Host Checker—Enable automatic installation through the Users/Administrators > User Realms/Administrator Realms > [Realm] > Authentication Policy > Host Checker page of the admin console. When you do, the system evaluates the realm-level option when the user accesses the sign-in page and then determines if the current version of Host Checker is installed on the user's machine. If Host Checker is not installed, the system attempts to install it using either an ActiveX or a Java delivery method.

  When a Windows user signs in to a device, the system attempts to install an ActiveX control on the user's system. If the system successfully installs the ActiveX control, the control manages the installation of the Host Checker program.

  If the system cannot install the ActiveX control because ActiveX is turned off on the user's system, it attempts to install Host Checker using Java. For Linux hosts, the system always uses the Java delivery method. The Java delivery method requires only user privileges, but Java must be enabled on the user's system. For the Firefox browser on Linux, the Java runtime and plug-in must be installed.

> **NOTE:** Due to some anomalies with Microsoft JVM, Host Checker may not install and an error box appears. If this occurs, click **Try Again**. The subsequent installation should succeed.

If the system cannot use the Java delivery method because Java is disabled on the user's system, it displays a no-access error message.

> **NOTE:** If Microsoft Vista is running on the user's system, the user must click the setup link that appears during the installation process to continue installing the setup client and Host Checker. On all other Microsoft operating systems, the setup client and Host Checker install automatically.

- The user or administrator manually installs Host Checker (Windows only)—Download the Host Checker installer from the Maintenance > System > Installers page of the admin console and use it to manually install Host Checker on the user's system.

> **NOTE:** To install Host Checker, users must have appropriate privileges, as described in the *Client-side Changes Guide* on the Pulse Secure Global Support Center (PSGSC) Center. If the user does not have these privileges, use the Juniper Installer Service available from the Maintenance > System > Installers page of the admin console to bypass this requirement.

Related Documentation

- Installing Host Checker Automatically or Manually

Client ActiveX Installation Delay

During end-user sign-in, the setup client is delivered through either ActiveX or Java, depending on the client system's capability. By default, Internet Explorer blocks ActiveX content and displays an information bar that lets the user decide whether to install the new ActiveX control.

> **NOTE:** For restricted users, the information bar displays help information only, it does not allow installation of new ActiveX controls.

The system displays to end-users an intermediate page with a 15-second delay to interact with the information bar content. End-users can choose to skip the installation (and the 15-second delay) by clicking the "click here" link. If end-users choose to skip the installation, they are not prompted again unless they clear their browser cookies.

Administrators can customize the message and locale displayed in this intermediate page by clicking the Custom Messages tab in the Default Options for User Roles page and filling out information under the User Login Messages section.

Related Documentation

- Customizing Messages

Using Host Checker with the GINA Automatic Sign-In Function

Using Host Checker in conjunction with the Windows Graphical Identification and Authorization (GINA) sign-in function for VPN Tunneling requires that you pay particular attention to the type, level, and number of items to verify on the client before granting or rejecting access to the system. Since the GINA sign-in function takes place before Windows has completely launched on the client, and therefore, before the user profile on Windows is created, we recommend you adopt the following practices when creating Host Checker policies you plan to use for Windows clients featuring the GINA sign-in function:

- You can check system-level processes at both realm enforce and realm evaluate. You can check user-level processes only at realm evaluate.

- If you have user-level processes at realm evaluate, create a separate VPN Tunneling role featuring only system-level policy checks that can be performed before Windows has completely launched on the client. Ensure that this role allows connectivity to the Windows Domain infrastructure in your secure network to support drive mapping, software updates, and group policies, for example. Mapping your users to this role allows the GINA authentication to complete. This role is in addition to the final role that you want the user to be mapped.

Related Documentation

- About VPN Tunneling

Installing Host Checker Automatically or Manually

To automatically install Host Checker on client computers:

1. In the admin console, choose **Authentication > Endpoint Security > Host Checker**.

2. Under Options, select **Auto-upgrade Host Checker** if you want the system to automatically download the Host Checker application to a client computer when the version of Host Checker on the system is newer than the version installed on the client. Here is a summary of what happens when the Auto-upgrade Host Checker option is selected or not selected:

   - If Host Checker is not installed on the client computer, Host Checker is installed automatically regardless of whether the Auto-upgrade Host Checker option is selected or not selected.

   - If the Auto-upgrade Host Checker option is selected and a previous version of Host Checker is installed, Host Checker is upgraded on the client automatically.

   - If the Auto-upgrade Host Checker option is not selected and a previous version of Host Checker is installed, Host Checker is not upgraded the client automatically.

   If you select the Auto-upgrade Host Checker option, note the following:

   - On Windows, the user must have administrator privileges in order for the system to automatically install the Host Checker application on the client. For more information, see the *Client-side Changes Guide* on the Pulse Secure Global Support Center (PSGSC) Center.

   - If a user uninstalls Host Checker and then signs in to a device for which the Auto-upgrade Host Checker option is not enabled, the user no longer has access to Host Checker.

3. Click **Save Changes**.

An administrator may choose to download and install Host Checker manually on their client systems. The Maintenance > System > Installers page of the admin console provides several applications and a service for download. You can download an application or service as a Windows executable file, which enables you to:

- Distribute the file to client machines using software distribution tools. This option enables you to install an application or service on client machines whose users do not have Administrator privileges, which are required to install the application or service.

- Post the executable in a secure repository so that users with the proper administrator right may download and install the appropriate version.

- Download and execute a script that automatically retrieves the proper version of the installer from an FTP server.

Related Documentation

- Specifying Host Checker Installation Options

Using Host Checker Reports and Logs

You can use the admin console to learn details about Host Checker policy violations.

Figure 105 shows the Compliance report. Note that the type of rule that has been violated is shown.

Figure105: Compliance Report

Figure 106 shows a log entry for a session that violates a Host Checker rule. Note the names of the configured rules that have been violated are shown in the log.

Figure106: User Access Log – Host Checker

To display the Compliance report:

1. Select **System > Reports > Compliance**.

2. Select a filter:

   - **Compliant**

   - **Non Compliant**

- **Remediated**

- **Not Assessed**

To display User Access logs:

1. Select **System Log/Monitoring**.

2. Click the **User Access** tab.

3. Click the **Log** tab.

Related Documentation

- *Host Checker Overview*

- Implementing Host Checker Policies

- Displaying System Logs

- Using Log Filters

Host Checker for Apple iOS

- Host Checker for Pulse iOS Clients

- Configuring Host Checker for Pulse iOS Clients

- Implementing Host Checker Policies for Pulse for iOS Devices

Host Checker for Pulse iOS Clients

Host Checker is a component of Pulse Secure client that reports the integrity of iOS endpoints that are attempting to connect to the system. Host Checker runs as a Trusted Network Connect (TNC) client on the endpoint. The client evaluates the endpoint according to predefined criteria and reports to the Trusted Network Connect server, which is a part of Connect Secure. If the endpoint is not in compliance with the Host Checker policies, then the user might not get access to the network or might get limited access to the network depending upon the enforcement policies configured by the administrator.

For iOS clients, Host Checker can evaluate client compliance based on the following predefined criteria:

- **OS Checks**—You can specify the iOS version or minimal version that must be installed on the device.

- **Jail Breaking Detection**—Jail breaking is a process that allows Apple iPhone, iPad and iPod Touch users to gain root access to the iOS operating system and bypass usage and access limitations imposed by Apple. With a jail broken device, an iOS user can install applications that are not available through the Apple App Store. Jail broken devices expose the device to a greater risk of running malicious applications.

- **Mobile Security Suite**—Pulse Mobile Security Suite (MSS) for iOS supports features that help protect the iOS device if it is lost or stolen: remote wipe, remote lock, and locate. Pulse MSS also allows the MSS administrator to provide VPN configurations to the device and enforce other usage practices such as requiring a lock password and specifying the password strength. Host Checker enables you to require that an iOS device have the Pulse Mobile Security Suite installed and enabled.

    i NOTE: The Pulse MSS check on the client is based on the presence of a valid client certificate issued by the Mobile Security Gateway. Therefore, to configure this Host Checker item, you need to add the CA used by MSS in the Certificates section of the system admin console.

Host Checker evaluation policies can be part of a larger Host Checker configuration that applies to many different types of clients or to iOS devices only.

Related Documentation

- Configuring Host Checker for Pulse iOS Clients

- Implementing Host Checker Policies for Pulse for iOS Devices

Configuring Host Checker for Pulse iOS Clients

Host Checker policies can be part of a larger Host Checker configuration that applies to many different types of clients or to iOS devices only. However, you might find it easiest to create a separate Host Checker policy specifically for iOS devices.

> **NOTE:** One type of iOS Host Checker rule, the Pulse Mobile Security Suite (MSS) check, requires that an iOS device be registered with the MSS. To create this type of rule, you must specify the same trusted client certificate authority (CA) that is configured on the Pulse Mobile Security Gateway (MSG). This means that the CA must already be defined in the Certificates section of the system admin console before you can create this type of Host Checker rule.

To create a Host Checker policy for iOS devices:

1. From the admin console, select **Authentication > Endpoint Security > Host Checker**.

2. In the Policies section, click **New** to open a New Host Checker Policy page.

3. Specify a name for the new policy and then click **Continue** to open the Host Checker Policy page.

   The name appears in lists when you implement the policy so be sure to use a descriptive name, such as **iOS HC Policy**.

4. Click the Mobile tab, and then click the iOS tab.

5. In the Rule Settings section, click **Select Rule Type** and select one of the following options and then click **Add**:

   - **OS Checks**—To specify the iOS version that must be installed on the device:

     1. Specify a descriptive name for this rule. For example **Must-Be-iOS-4.1-or-higher**. Rule names cannot include spaces.

     2. Specify the criteria. For example, to enforce iOS 4.1 and later, create two conditions: Equal to 4.1 and Above 4.1.

        Host Checker supports iOS versions 4.1 through 4.3.X.

     3. Click **Save Changes**.

   - **Jail Breaking Detection**—Jail breaking is a process that allows Apple iPhone, iPad and iPod Touch users to gain root access to the iOS operating system. and bypass usage and access limitations imposed by Apple. With a jail broken device, an iOS user can install applications that are not available through the Apple App Store. Jail broken devices possess a greater risk of running malicious applications.

     1. Specify a descriptive name for this rule. For example **No-iOS-Jailbreak**.

     2. The **Don't allow Jail Broken devices** check box is enabled by default.

     3. Click **Save Changes**.

   - **Mobile Security Suite**—Pulse Mobile Security Suite (MSS) for iOS supports features that help protect the iOS device if it is lost or stolen: remote wipe, remote lock, and locate.

Pulse MSS also allows the MSS administrator to provide VPN configurations to the device. An MSS rule allows you to require that an iOS device have MSS installed and enabled.

1. Specify a descriptive name for this rule. For example **iOS-MSS**.

2. In the Criteria section, the **Enable Mobile Security Suite on the device** check box is enabled by default.

3. Click the **MDM Trusted client CA** box to choose a certificate authority.
   MDM (Mobile Device Management) enables remote configuration and management of mobile devices. The system verifies registration with Pulse MSS through the presence of a valid client certificate issued by the Mobile Security Gateway. Therefore, to configure this Host Checker item, you need to know the client certificate authority that is used by the Pulse Mobile Security Gateway. The certificate authorities listed in the box are defined in the System > Configuration section of the system admin console.

4. Optionally, specify additional criteria (the certificate field and expected value), and then click **Add**.

5. Click **Save Changes**.

6. After you have configured all of your rules, specify how you want to enforce them by choosing one of the following options:

   - All of the rules

   - Any of the rules

   - Custom

   For Custom requirements, you can specify a custom expression using Boolean operators AND and OR and also group and nest conditions using parenthesis.

7. Specify remediation options:

   - **Enable custom instructions**—If you enable this check box, a text box appears and allows you to type information that appears on the user's device if Host Checker discovers an issue. For example, if you enabled the MSS rule that terminates the VPN session of Host Checker discovers a virus, you can instruct the user to run a virus scan to clear the issue before trying to connect.

   - **Send reason strings**—Select this option to display a message to users (called a reason string) that explains why the client machine does not meet the Host Checker policy requirements. For example, if the jailbreak detection policy fails, Pulse appears, **A jailbroken device is not allowed to access the network. Please contact your network administrator.**

8. When you are finished, click **Save Changes**.

A host checker policy configured for a VPN tunnel is not triggered if the VPN tunnel is launched automatically by VPN on Demand on an Apple iOS device. If the VPN session is started through the Pulse client, the host checker policy is applied correctly. A VPN on Demand configuration enables an iOS device to automatically initiate a VPN connection when any application running on the phone initiates a connection to a host in a predefined set of hosts. A VPN on Demand connection uses client certificate-based authentication so the user does not have to provide credentials every time a VPN connection is initiated.

Related Documentation

- Host Checker for Pulse iOS Clients

- Implementing Host Checker Policies for Pulse for iOS Devices

---

Implementing Host Checker Policies for Pulse for iOS Devices

After you create one or more Host Checker policies for iOS devices, you must implement them. The system can use Host Checker policies at the realm or the role level.

Realm Authentication—You can configure a realm authentication policy to download and run Host Checker with a particular Host Checker policy. If the iOS device does not meet the Host Checker requirements, then the system can deny access. You can provide remediation information in the Host Checker policy to describe the requirement and help users take steps to solve the issue.

To enable a Host Checker policy for a realm:

1. From the admin console, select **Users > User Realms > SelectRealm > Authentication Policy > Host Checker**. The Host Checker page displays all of the available Host Checker policies.

2. Select the check box next to each policy you want to include. Select one or both of the following check boxes next to the policy:

   - **Evaluate Policies**—Evaluates without enforcing the policy on the iOS device and allows access.

   - **Require and Enforce**—Requires that the iOS device be in compliance with the Host Checker policy. The system downloads Host Checker to the iOS device after the user is authenticated and before the user is mapped to any roles in the system. Selecting this option automatically enables the Evaluate Policies option.

3. Optionally select **Allow access to realm if any ONE of the selected "Require and Enforce" policies is passed**. This check box is available if you selected more than one Host Checker policy. If you enable this check box, an iOS device is allowed access if it passes any of the Require and Enforce policies. The Cache Cleaner policy does not apply to iOS devices.

4. Click **Save Changes**.

Role—You can configure a role to download and run Host Checker with a particular Host Checker policy. If the iOS device does not meet the Host Checker requirements, then the system can deny access or assign the user to a remediation role that has limited access. You can provide remediation information in the Host Checker policy to help users take steps to solve the issue.

To enable a Host Checker policy for a role:

1. From the admin console, select **Users > User Roles > SelectRole > General > Restrictions > Host Checker**. The Host Checker page displays all of the available Host Checker policies.

2. Select **Allow users whose workstations meet the requirements specified by these Host Checker policies**.

3. In the Available Policies list, select the policies that you want to apply to select them, and then click Add to move them to the Selected Policies list. To select a policy click it. So select more than one policy, use Ctrl+click.

4. Optionally select **Allow access to the role if any ONE of the selected policies (except cache-cleaner) is passed**. This check box is available if you selected more than one Host Checker policy. If you enable this check box, an iOS device is allowed access if it passes any of the Require and Enforce policies. The Cache Cleaner policy does not apply to iOS devices.

5. Click **Save Changes**.

Related Documentation

- Host Checker for Pulse iOS Clients

Using Proxy Exceptions

Connect Secure clients parse Internet Explorer's static proxy exception list. The system supports most exceptions that Internet Explorer supports with the following limitations:

- For IP address exception, we support n.*.*.*, n.n.*.*, n.n.n.*. For example, 10.*.*.*, 10.10.*.*, 10.10.10.*, or 10.10.10.10. We do not support 10* or 10.*.10.* even though Internet Explorer may support them.

- For string expression, we support specific strings such as my.company.net,or a wild card at front of the string, for example, *.my.company.net or *.company.net. We do not support *.company.*, *.company*, *.company. *.com, *.net *.com and so forth.

CHAPTER 16 Cache Cleaner

About Cache Cleaner

Cache Cleaner is a Host Checker policy that removes residual data, such as temporary files or application caches, left on a user's machine after a session. For example, when a user signs in to a device from an Internet kiosk and opens a Microsoft Word document using a browser plug-in, Cache Cleaner can remove the temporary copy of the Word file stored in the browser cache (Windows folder) when the session terminates. By removing the copy, Cache Cleaner prevents other kiosk users from finding and opening the Word document after the user concludes the session.

Cache Cleaner can also prevent Web browsers from permanently storing the usernames, passwords, and Web addresses that users enter in Web forms. By preventing browsers from improperly caching this information, Cache Cleaner keeps confidential user information from being stored on untrusted systems.

Related Documentation

- Setting Global Cache Cleaner Options

- Implementing Cache Cleaner Options

Setting Global Cache Cleaner Options

When you enable Cache Cleaner, it clears all content downloaded through the system's Content Intermediation Engine from a user's system. In addition, you can use settings in the Authentication > Endpoint Security > Cache Cleaner page of the admin console to clear content from the following places:

- Specified hosts and domains—If you enable WSAM or JSAM, you may want to configure Cache Cleaner to clear additional hosts and domains. When users browse the Internet outside the system using WSAM or JSAM, Internet files appear in their temporary Internet file folder. To delete these files using Cache Cleaner, you must specify the appropriate hostname (for example, www.yahoo.com).

- Specified files and folders—If you enable your users to access client-server applications on their local systems, you may want to configure Cache Cleaner to clear the temporary files and folders that the applications create on the users' systems.

NOTE: If you configure Cache Cleaner to remove files from a directory, Cache Cleaner clears all files, including those that the user has explicitly saved to the directory and files that were in the directory prior to the session.

Only one Cache Cleaner policy is allowed. You can neither delete the default Cache Cleaner policy (named "Cache Cleaner Policy") nor create a new one.

To specify global Cache Cleaner options:

1. Select **Authentication > Endpoint Security > Cache Cleaner** in the admin console.

2. Under Options:

    1. Specify how often Cache Cleaner runs in the Cleaner Frequency field. Valid values range from 1 to 60 minutes. Each time Cache Cleaner runs, it clears all content downloaded through the Content Intermediation Engine plus the browser cache, files, and folders you specify under the Browser Cache and Files and Folders sections.

    2. Select the **Disable AutoComplete of web addresses** check box to prevent the browser from using cached values to automatically fill in Web addresses during the user's session. When you select this option, the system sets the following Windows registry value to 0 during the user's session: HKEY_CURRENT_USER\Software\\Microsoft\\Windows\\CurrentVersion\\Explorer\ AutoComplete.

       Then, at the end of the session, the system restores the registry value to its original setting.

    3. Select the **Disable AutoComplete of usernames and passwords** check box to prevent Internet Explorer from automatically filling in user credentials in Web forms using cached values. Selecting this option also disables the "Save Password?" prompt on Windows systems. When you select this option, the system sets the following Windows registry values to 0:

       • HKEY_CURRENT_USER\Software\Microsoft\Internet Explorer\Main\FormSuggest Passwords

       • HKEY_CURRENT_USER\Software\Microsoft\Internet Explorer\Main\FormSuggest Passwords\FormSuggest PW Ask

       • HKEY_CURRENT_USER\SOFTWARE\Microsoft\ Windows\CurrentVersion\Internet Settings\ DisablePasswordCaching

    4. Select the **Flush all existing AutoComplete Passwords** check box to clear any cached passwords that Internet Explorer has cached on the user's system. When you select this option, the system sets the following Windows registry value to 0:

       HKEY_CURRENT_USER \Software\\Microsoft\\Internet Explorer\\ IntelliForms\\SPW

       Then, select one of the following options:

       • Select the **For Secure Gateway session only** option button to specify that the system should restore the user's cached passwords at the end of his session.

       • Select the **Permanently** option button to permanently delete the user's cached passwords.

5. Select the **Empty Recycle Bin and Recent Documents list** check box to empty the recycle bin and clear the recent documents list. The entire contents are removed, not just the files related to the user's sessions.

6. Under Browser Cache, enter one or more hostnames or domains (wildcards are permitted). When a user session ends, Cache Cleaner removes any content in the browser cache that originates from these servers. Cache Cleaner also removes this content when it runs at the specified cleaner frequency interval. Note that the system does not resolve hostnames, so enter all possible representations of a server, such as its hostname, FQDN, and IP address.

7. Under Files and Folders:

    1. Specify either:

        - The name of a file that you want Cache Cleaner to remove.

        - The complete directory path to a folder whose contents you want Cache Cleaner to remove. If you specify a directory, select Clear Subfolders to also clear the contents of any subdirectories within this directory.

    2. Select the **Clear folders only at the end of session** check box if you want Cache Cleaner to clear directory contents only at the end of the user session. Otherwise, Cache Cleaner also clears files and folders at the specified cleaner frequency interval

        > **NOTE:** When specifying files and folders to clear, note the following:
        >
        > Cache Cleaner uses a cookie called DSPREAUTH to send the client's status to the system. If you delete this cookie from the user's client, Cache Cleaner does not work properly. To avoid problems, do not specify Internet Explorer directories such as <userhome>\Local Settings\Temporary Internet Files* under File or folder path. Note that Cache Cleaner still clears all of the Internet Explorer cache downloaded from the system and the hosts specified in the Hostnames box, regardless of what directories you specify under Files and Folders.
        >
        > For the Firefox browser, Cache Cleaner clears only those directories you specify under Files and Folders.

    3. Click **Save Changes** to save these settings globally.

        If more than one valid session exists from the same system and Cache Cleaner is used in those sessions, all sessions are terminated when a user signs out from one of the sessions. To prevent this, turn off Cache Cleaner for those sessions that do not need Cache Cleaner.

        > **NOTE:** If multiple administrators or end users to a single system are signed in from the same client and at least one of them deploys Cache Cleaner, unexpected results may occur. For example, Cache Cleaner might shut down, role privileges might be lost, and forced disconnections might occur.

Related Documentation

- Implementing Cache Cleaner Options

Implementing Cache Cleaner Options

After you specify which hosts, domains, files, and folders to clear using settings in the Authentication > Endpoint Security > Cache Cleaner page of the admin console, you can restrict system and resource access by requiring Cache Cleaner in the following options:

- Realm authentication policy—When users try to sign in to a device, the system evaluates the specified realm's authentication policy to determine if the pre-authentication requirements include Cache Cleaner. You can configure a realm authentication policy to evaluate whether to require and enforce the Cache Cleaner policy in order for the user to log in to the specified realm. If the user's computer does not meet the requirements, then the user is denied access to the device. As a post-authentication requirement, you can evaluate without enforcing the Cache Cleaner policy on the client and allow user access. You configure realm-level restrictions through the Users > User Realms > *Realm*> Authentication Policy > Host Checker page of the admin console.

- Role—When the system determines the list of eligible roles to which it can map an administrator or user, it evaluates each role's restrictions to determine if the role requires Cache Cleaner to run on the user's workstation. If it does and the user's machine is not already running Cache Cleaner, then it does not map the user to that role. You can control which roles the system maps a user to by using settings in Users > User Realms > *Realm* > Role Mapping. Select or create a rule and then select Custom Expressions. You can configure role-level restrictions through the Users > User Roles > *Role* > General > Restrictions > Host Checker page of the admin console.

- Resource policy—When a user requests a resource, the system evaluates the resource policy's detailed rules to determine whether or not Cache Cleaner needs to be installed or running on the user's workstation. The system denies access to the resource if the user's machine does not meet the Cache Cleaner requirement. You can implement Cache Cleaner restrictions at the resource policy level through the Condition Field box of the Rules window. Select Users > Resource Policies > *Resource* > *Policy* > Detailed Rules and set hostCheckeryPolicy = 'Cache Cleaner policy'.

You may specify that the system evaluate your Cache Cleaner policies only when the user first tries to access the realm, role, or resource that references the Cache Cleaner policy. Or, you can use settings in the Authentication > Endpoint Security > Cache Cleaner tab to specify that the system periodically re-evaluate the policies throughout the user's session. If you choose to periodically evaluate Cache Cleaner policies, the system dynamically maps users to roles and allows users access to new resources based on the most recent evaluation.

When the user tries to access a device, Host Checker evaluates its policies (Cache Cleaner is a Host Checker policy) in the following order:

- Initial evaluation

- Realm-level policies

- Role-level policies

- Resource-level policies

Related Documentation

- Setting Global Cache Cleaner Options

Specifying Cache Cleaner Restrictions

To specify Cache Cleaner restrictions:

1. Select **Authentication > Endpoint Security > Cache Cleaner** and specify global options for Cache Cleaner to apply to any user for whom Cache Cleaner is required in an authentication policy, a role mapping rule, or a resource policy.

2. Implement Cache Cleaner at the realm level and role level as you would with Host Checker.

3. Create role-mapping rules based on a user's Cache Cleaner status as you would with Host Checker.

4. To implement Cache Cleaner at the resource policy level:

   1. Select **Users > Resource Policies > Select Resource > Select Policy > Detailed Rules**.

   2. Click **New Rule** or select an existing rule from the Detailed Rules list.

   3. Create a custom expression in a detailed rule that sets hostCheckeryPolicy = 'Cache Cleaner policy'.

Related Documentation

- Configuring Host Checker Restrictions
- Using Custom Expressions in Rule Configuration

About Cache Cleaner Logs

Since Cache Cleaner is a Host Checker policy, it is included in the Host Checker logs. Use the System > Log/Monitoring > Client Logs > Settings tab to enable client-side logging for Host Checker. When you enable this option, the system writes a client-side log to any client that uses Host Checker. The system appends to the log file each time the feature is invoked during subsequent user sessions. This feature is useful when working with the support team to debug problems with the respective feature.

PART 4 Remote Access

- Hosted Java Applets Templates
- Citrix Templates
- Lotus iNotes Templates
- Microsoft OWA Templates
- Microsoft Sharepoint Templates
- Web Rewriting
- File Rewriting
- Secure Application Manager
- Telnet/SSH
- Terminal Services
- Pulse Collaboration
- E-Mail Client
- VPN Tunneling
- Enterprise Onboarding and Secure Mail

CHAPTER 17 Hosted Java Applets Templates

About Hosted Java Applet Templates

The Java applet upload feature enables you to store the Java applets of your choice directly on the device without employing a separate Web server to host them. When you use this feature, you simply upload the applets to the device (along with additional files that the applets reference) and create a simple Web page through the system that references the files. Then, the system intermediates the Web page and Java applet content using its Content Intermediation Engine.

For example, you might want to use the system to intermediate traffic between an IBM AS/400 system on your network and individual 5250 terminal emulators on your users' computers. To configure the system to intermediate this traffic, obtain the 5250 terminal emulator's Java applet. Then you can upload this applet to the system and create a simple Web page that references the applet. After you create the Web page through the system, it creates a corresponding bookmark that users can access through their home pages.

The system enables you to host Java applets using Web resource profile templates (described in these topics) as well as through Terminal Services resource profiles.

The hosted Java applets feature is a standard feature on all Connect Secure devices.

Related Documentation

- Task Summary: Hosting Java Applets

Task Summary: Hosting Java Applets

The Java applet upload feature enables you to store the Java applets of your choice directly on the device without employing a separate Web server to host them.

To host Java applets on the device:

1. Specify which applets you want to upload, create bookmarks that reference the uploaded applets, and specify which roles can access the bookmarks using settings in the Users > Resource Profiles > Web page of the admin console.

2. (Optional.) To sign your Java applets, Select System > Configuration > Certificates > Code-Signing Certificates in the admin console to upload the Java certificate to the device. If you choose to skip this step, the user sees an untrusted certificate warning each time he accesses the corresponding bookmark.

3. (Optional.) To improve the performance of your Java applications:

   1. Select Enable Java instrumentation caching on the Maintenance > System > Options page of the admin console. This option can improve the performance of downloading Java applications.

   2. After you finish configuring the system, cache your Java applet and access it as an end user. This action eliminates the performance hit that occurs through the intermediation engine when the first end user accesses the applet.

Related Documentation

- Using Code-Signing CAs

- Uploading Java Applets to Connect Secure

- Signing Uploaded Java Applets

- Creating HTML Pages That Reference Uploaded Java Applets

Uploading Java Applets to Connect Secure

You can use Java applets to intermediate traffic to various types of applications through the system. For example, you can upload the 3270 applet, 5250 applet, or Citrix Java applet. These applets enable users to establish sessions to IBM mainframes, AS/400s, and Citrix MetaFrame servers through terminal emulators. (Note that to enable the Citrix Java ICA client through a session, you must upload multiple Citrix .jar and .cab files to the device.

The system enables you to upload individual .jar and .cab files or .zip, .cab, or .tar archive files. Archive files can contain Java applets and files referenced by the applets. Within the .zip, .cab, or .tar file, the Java applet must reside at the top level of the archive. You can upload any number of files to the system as long as their combined size does not exceed 100 MB.

To ensure compatibility with both Sun and Microsoft Java Virtual Machines (JVMs), you must upload both .jar and .cab files to the device. (The Sun JVM uses .jar files, whereas the Microsoft JVM uses .cab files.)

> NOTE: When you upload Java applets, the system asks you to read a legal agreement before it finishes installing the applets. Read this agreement carefully—it obligates you to take full responsibility for the legality, operation, and support of the Java applets that you upload.
>
> You can only upload 100 MB of Java applets to the system. The system displays the size of each applet that you upload on the Java Applets page so you can determine, if necessary, which applets you want to delete.
>
> Uploading Java applets requires signed ActiveX or signed Java applets to be enabled within the browser to download, install, and launch the client applications.

Related Documentation

- Task Summary: Hosting Java Applets

Signing Uploaded Java Applets

Unlike other Java applets that users can access through the system, you do not have to create a separate code-signing policy for the Java applets that you upload. The system automatically signs (or re-signs) them using the appropriate code-signing certificate. A code-signing certificate (also called an applet certificate) is a type of server-side certificate that re-signs Java applets intermediated by the system.

The system automatically signs (or resigns) your hosted Java applets with the code-signing certificate that you install through the System > Configuration > Certificates > Code-signing Certificates page of the admin console. If you do not install a code-signing certificate on the system, it uses its self-signed applet certificate to sign or re-sign the applets. In this case, users see an "untrusted certificate issuer" warning whenever they access the Java applets through the system.

> ℹ️ NOTE: The system re-instruments and re-signs your uploaded Java applets whenever you change (that is, import, renew, or delete) the corresponding code-signing certificate.

Related Documentation

- Task Summary: Hosting Java Applets

Creating HTML Pages That Reference Uploaded Java Applets

When uploading a Java applet to the system, you must create a simple Web page that references the applet. Users can access this Web page through a bookmark on their home pages or from external Web servers.

The Web page must contain a simple HTML page definition that references the uploaded Java applet. The Web page can also contain any additional HTML and JavaScript that you choose. The system can generate some of the Web page for you, including the HTML page definition and the references to your Java applet. (Note, however, that the system is not aware of all the applet-specific parameters that are required by your applet—you must find and fill these parameters in yourself.) When the system generates this HTML, it creates placeholders for any undefined values and prompts you to fill in the necessary values.

You can create these Web pages through Java applet upload resource profiles.

Related Documentation

- Task Summary: Hosting Java Applets

- Accessing Java Applet Bookmarks

Accessing Java Applet Bookmarks

Users can access the applets you upload to the system using two methods:

- Bookmarks on the end-user console—When you create a Web page that references your uploaded Java applets, the system creates a corresponding link to the Web page and displays that link in the Bookmarks section of the end-user console. Users who map to the appropriate role can simply click the link to access the Java applet.

- Links on external Web servers—Users can link to the Java applet bookmarks from an external Web server by simply using the correct URLs. When the user enters a bookmark's URL (or clicks an external link that contains the URL), the system prompts the user to enter his username and password. If he properly authenticates, it allows him to access the bookmark. You can construct the URL to the Java applet bookmark using the syntax described in either of the following lines:

https://SecureAccessGateway_hostname/dana/home/launchwebapplet.cgi?
bmname=bookmark Name

https://SecureAccessGateway_hostname/dana/home/launchwebapplet.cgi?
id=<resourceID>&bmname=bookmarkName

You can determine the ID for a Java applet bookmark by accessing it through the home page and then extracting the ID from the Web browser's address bar.

> NOTE: Although the system enables you to create multiple bookmarks with the same name, we strongly recommend that you use a unique name for each. If multiple bookmarks have the same name and a user accesses one of these bookmarks using a URL that includes the bmname parameter, the system randomly picks which of the identically named bookmarks to display to the user. Also note that the bmname parameter is case-sensitive.
>
> If you create links on external servers to Java applet bookmarks on the system and you are using multiple customized sign-in URLs, some restrictions occur.

Related Documentation

- Configuring Hosted Java Applet Resource Profile Bookmarks
- Creating Hosted Java Applets Bookmarks Through the User Roles Page

Creating a Hosted Java Applet Resource Profile

To create a hosted Java applet resource profile:

1. Select **Users > Resource Profiles > Web** in the admin console.

2. Click **New Profile**.

3. Select **Hosted Java Applet** from the Type list.

4. Enter a unique name and optionally a description for the resource profile.

5. Select the Java applet that you want to associate with the resource profile from the Applet to use list. Or, if the applet that you want to use is not currently available in the list, click Edit Applet. Then:

   1. Click New Applet to add an applet to this list. Or, select an existing applet and click Replace (to replace an existing applet with a new applet) or Delete (to remove an applet from the system.)

      > NOTE: If you replace an existing archive, make sure that the new applet archive contains all of the necessary files for the applet to successfully launch and run. If the associated HTML for the applet refers to files that do not exist in the new archive, then the applet will not function correctly.
      >
      > The system only allows you to delete applets that are not currently in use by a Web or Terminal Services resource profile.

2. Enter a name to identify the applet in the Name box (for new and replaced applets only).

3. Browse to the applet that you want to upload. You can upload applets (.jar or .cab files) or archives (.zip, .jar, and .tar files) that contain applets and all of the resources that the applets need (for new and replaced applets only).

4. Select the **Uncompress jar/cab file** check box if the file that you selected is an archive that contains the applet (New and replaced applets only).

5. Click **OK** and then click **Close Window**.

> ⓘ NOTE: When you select an applet in the Java Applets dialog box, you are loading third-party software onto the Pulse Secure product. By clicking OK, you are agreeing to the following terms on behalf of yourself (as purchaser of the equipment) or the organization that purchased the Pulse Secure product, as applicable.
>
> By loading third party software onto the Pulse Secure product, you are responsible for obtaining all rights necessary for using, copying, and/or distributing such software in or with the Pulse Secure product. Pulse Secure is not responsible for any liability arising from use of such third party software and will not provide support for such software. The use of third party software may interfere with the proper operation of the Pulse Secure product and/or Pulse Secure software, and may void any warranty for the Pulse Secure product and/or Pulse Secure software.

6. Use settings in the Autopolicy: Java Access Control section to enable access if your Java applets need to make socket connections.

7. Click **Save and Continue**.

8. Select the roles to which the resource profile applies In the Roles tab and click Add.

   The selected roles inherit the autopolicies and bookmarks created by the resource profile. If it is not already enabled, the system also automatically enables the Web option in the Users > User Roles > Select_Role > General > Overview page of the admin console and the Allow Java Applets option Users > User Roles > Select_Role > Web > Options page of the admin console for all of the roles you select.

9. Click **Save Changes**.

10. Create bookmarks in the Bookmarks tab.

Related Documentation

- Configuring Hosted Java Applet Resource Profile Bookmarks

Configuring Hosted Java Applet Resource Profile Bookmarks

You must create bookmarks to your hosted Java applets to enable end users to access the applets.

To configure hosted Java applet resource profile bookmarks:

1. Select **Users > Resource Profiles > Web >** *Select Resource Profile* **> Bookmarks** in the admin console.

2. Click the appropriate link in the Bookmark column if you want to modify an existing bookmark. Or, click New Bookmark to create an additional bookmark.

> ⓘ NOTE: Although it is generally easiest to create a resource profile session bookmark through the resource profile configuration page, you can choose to create one through the user roles page as well if you have already created a resource profile.

3. Enter a name and optionally a description for the bookmark. This information displays on the home page. (By default, the system names the bookmark the same name as the corresponding resource profile.)

> ⓘ NOTE: We strongly recommend that you use a unique name for each bookmark to make it clear to users which link they are accessing.

4. Click Generate HTML to create an HTML page definition that includes references to your Java applets. Then, fill in any required attributes and parameters.

    If you are using HTML generated by the system, make sure to search the HTML code for "__PLEASE_SPECIFY__" and update the code as necessary.

    You can also add more HTML or JavaScript to this Web page definition. the system rewrites all of the code that you enter in this field

> ⓘ NOTE: Make sure to enter unique HTML in this field. If you create two bookmarks with the same HTML code, the system deletes one of the bookmarks in the end-user view. You will still be able to see both bookmarks, however, in the administrator console.

5. List those attributes in the Multi-Valued User Attributes box if your HTML code contains attributes that may expand to multiple values (such as userAttr.hostname or userAttr.ports), . When the user signs into a device, the system evaluates these attributes and creates separate bookmarks as necessary based on each of the individual values. If you use an attribute that expands to multiple values, but do not enter that attribute in this box, the system creates a single bookmark based on the attribute's first value.

6. Under Display options, click Bookmark opens new window to automatically open the Web resource in a new browser window. Note that this functionality applies only to role bookmarks and not bookmarks created by users. Next, select the following options if you want to hide UI elements from the user:

    - Do not display the browser address bar—Select this option to remove the address bar from the browser window. This feature forces all Web traffic through the system by precluding users in the specified role from typing a new URL in the address bar, which circumvents the system.

    - Do not display the browser toolbar—Select this option to remove the menu and toolbar from the browser. This feature removes all menus, browsing buttons, and bookmarks from the browser window so that the user browses only through the system.

7. Under Roles, specify the roles to which you want to display the bookmark if you are configuring the bookmark through the resource profile pages:

    - ALL selected roles—Select this option to display the bookmark to all of the roles associated with the resource profile.

    - Subset of selected roles—Select this option to display the bookmark to a subset of the roles associated with the resource profile. Then select roles from the ALL Selected Roles list and click Add to move them to the Subset of selected roles list.

8. Click **Save Changes**.

Related Documentation

- Defining Resource Profile Bookmarks

- Creating Hosted Java Applets Bookmarks Through the User Roles Page

- Creating HTML Pages That Reference Uploaded Java Applets

- Required Attributes for Uploaded Java Applets

- Required Parameters for Uploaded Java Applets

Creating Hosted Java Applets Bookmarks Through the User Roles Page

It is generally easiest to create a hosted Java applets bookmark through the resource profile configuration pages, as explained in previous topic. However, you can choose to create a resource profile session bookmark through the user roles page using the following instructions:

1. Select **Users > Roles > Select_Role > Web > Bookmarks** in the admin console.

2. Click **New Bookmark**.

3. Select **Pick a Web Resource Profile** from the Type list. (The system does not display this option if you have not already created a hosted Java applet resource profile.)

4. Select an existing resource profile.

5. Click **OK**. (If you have not already associated the selected role with the resource profile, the system automatically makes the association for you. The system also enables any access control policies for the role that are required by the resource profile.)

6. If this role is not already associated with the selected resource profile, the system displays an informational message. If you see this message, click **Save Changes** to add this role to the resource profile's list of roles and to update the profile's autopolicies as required. Then, repeat the previous steps to create the bookmark.

7. Configure the bookmark settings.

   > NOTE: When you create a resource profile bookmark through the user roles page (instead of the standard resource profiles page), the system only associates the generated bookmark with the selected role. The system does not assign the bookmark to all of the roles associated with the selected resource profile.

Related Documentation

- Accessing Java Applet Bookmarks

- Configuring Hosted Java Applet Resource Profile Bookmarks

Required Attributes for Uploaded Java Applets

When you create a Java applets bookmark through the system, you must define the following attributes and their corresponding values. If you use the Generate HTML feature, it populates some of this information for you and adds PLEASE_SPECIFY to those attributes whose values you must specify. When specifying attributes and their corresponding values, use the attribute="value" format.

> NOTE: The system generates parameters that it knows are required. Note, however, that it is not aware of all the applet-specific parameters that are required by your applet—you must find and fill in these parameters yourself.

Attributes that are required by the system include:

- code—Indicates which class file to invoke in your Java applet. Use this value to point to your Java applet's main function. Example:

    applet code="com.citrix.JICA"

- codebase—Indicates where the Web browser can fetch the applet. Use the <<CODEBASE>> variable, which points to the location on the system where it stores the Java applet. When entering a path to a file, note that <<CODEBASE>> includes a trailing slash, which means the following example works:

    <img src="<<CODEBASE>>path/to/file">

    This example does not work:

    <img src="<<CODEBASE>>/path/to/file">

- archive—Indicates which archive file (that is, .jar, .cab, or .zip file) the Web browser should fetch. Example:

    archive="JICAEngN.jar"

In addition to the required attributes listed earlier, you may also use the following optional attributes when creating a Java applet bookmark:

- name—Specifies a label for the Java applet. Example:

    name="CitrixJICA"

- host—Specifies, for terminal sessions, the server to which the system should connect.

- port—Specifies, for terminal sessions, the port to which the system should connect.

- width and height—Indicates the size of the Java applet window. Example:

    width="640" height="480"

- align—Indicates the Java applet window's alignment within the browser window. Example:

    align="top"

> NOTE: When defining attributes and their corresponding values, note the following:
>
> - We strongly recommend that you not include useslibrarycabbase parameter in the HTML, because it causes the cab file to be permanently installed on the user's machine. If you later change a cab file on the system, all users will have to manually delete the cab files on their machines to get the new version from the system.
>
> - We do not support applet tags that are constructed through the document.write function because the dynamic HTML interferes with the system parser.
>
> - We do not support relative links to URLs, documents, or images in your HTML. If you do, the links will break when the user tries to access them from the end-user console. Instead, you should include absolute links. If you are linking to a document or image included in your zip file, use the <<CODEBASE>> variable to indicate that the system can find the file in the uploaded zip archive. For example:
>
>   <img src="<<CODEBASE>>*yourcompany_logo.gif*" alt="*YourCompany*">

Related Documentation

- Required Parameters for Uploaded Java Applets

Required Parameters for Uploaded Java Applets

When you create a Java applets bookmark through the system, you must specify parameters and values that should be passed to the Java applet. These parameters are completely applet-specific. When specifying parameters and their corresponding values, use the following format:

<param name="*parameterName*" value="*valueName*">

Where all of the text is literal except *parameterName* and *valueName*.

You can use variables to pass values to the Java applet by enclosing the variable names in double-brackets. For example, you might choose to pass the <<username>> and <<password>> values to the Java applet.

> **i** NOTE: When using the Java applet upload feature, if you include the <password> token within the generated HTML, it appears as cleartext if you view the source in the browser window that launches the applet. This behavior cannot be changed because the system does not control how the Java applet processes the password. We strongly discourage the use of the <password> token in the HTML code.

If you find a Web page that contains an applet that you want to use, go to the demonstration site and view the source on the page that runs the Java applet. Within the source, look at the applet tag. Pick out the code attribute in the source and determine if it contains any special parameters that you need to pass to the browser. In most cases, you should be able to copy and paste the code attribute and its corresponding parameters directly into the HTML field for your bookmark. Note, however, that if a parameter references a resource on the local Web server, you cannot copy and paste the reference into the bookmark because the system does not have access to the other Web server's local resources. When copying and pasting parameters from another source, always check the values of the parameters.

Related Documentation

- Required Attributes for Uploaded Java Applets

Use case: Creating a Citrix JICA 9.5 Java Applet Bookmark

This topic discusses how to enable access to a Citrix Metaframe server through the system using the 9.5 Java version of the Citrix ICA client (JICA).

> **i** NOTE: In addition to the method described here, you can also use Terminal Services resource profiles to host the Java versions of Citrix ICA clients.
>
> The system supports several mechanisms for intermediating traffic between a Citrix server and client, including the Terminal Services, JSAM, WSAM, VPN Tunneling, and hosted Java applets features.

To enable the Citrix JICA 9.5 client using the Java applet upload feature:

1. Import code-signing certificates.

2. Download **JICAcomponents.zip** from the citrix.com downloads page.

3. Create a hosted Java applet resource profile through the Users > Resource Profiles > Web page of the admin console. When defining the resource profile:

   1. Upload the archived Citrix container file.

2. When uploading the applet, select the **Uncompress jar/cab file** check box because the container file contains multiple jar and cab files.

3. Specify any Metaframe servers to which these applets may connect.

4. Assign the resource profile to the appropriate roles.

5. Generate the Web page for the bookmark in the resource profile's Bookmarks tab. The system automatically inserts all of the .jar files into the corresponding Web page. (JICA 95 supports only Sun JVM, so no cab files are present.) Then, specify parameters for the Citrix client using the following examples as a guide. (Note that the bookmark in the following example can contain references to the jar and cab files that are in the zip file.

JICA 9.5 Applet Example

```
<html>
<head>
<title>jica95 Applet</title>
<meta http-equiv="Content-Type" content="text/html; charset=iso-8859-1">
</head>
<!--
Notes:
1) << CODEBASE >> is a system value that will get replaced at the time the applet is launched.
Please do not modify this value.
2) Please modify the remaining values as needed.
3) Please make sure all attribute names/values are enclosed in double quotes.
-->
<body>
    <applet code="com.citrix.JICA"
            codebase="<< CODEBASE >>"
            archive="JICA-browseN.jar,JICA-cdmN.jar,JICA-clipboardN.jar,JICA-configN.jar,
  JICA-coreN.jar,JICA-printerN.jar,JICA-seamlessN.jar,JICA-sicaN.jar,
  JICA-zlcN.jar,JICAEngN.jar,cryptojN.jar,sslN.jar,JICA-audioN.jar"
            width="640" height="480"
            name="jica95" align="top">
        <param name="code" value="com.citrix.JICA">
        <param name="codebase" value="<< CODEBASE >>">
        <param name="archive" value="JICA-browseN.jar,JICA-cdmN.jar,JICA-clipboardN.jar,
  JICA-configN.jar,JICA-coreN.jar,JICA-printerN.jar,JICA-seamlessN.jar,
  JICA-sicaN.jar,JICA-zlcN.jar,JICAEngN.jar,cryptojN.jar,sslN.jar,JICA-audioN.jar">
        <param name="cabbase" value="">
        <param name="name" value="jica95">
        <param name="width" value="640">
        <param name="height" value="480">
        <param name="align" value="top">
        <!--
        Please specify additional params here after the comment.
        <param name="paramname" value="paramvalue">
        -->
        <param name="Address" value="__PLEASE_SPECIFY__">
        <param name="Username" value="<< user >>">
        <param name="password" value="<< password >>">
        <param name="EncryptionLevel" value="1">
        <param name="BrowserProtocol" value="HTTPonTCP">
    </applet>
</body>
</html>
```

JICA 8.x Applet Example

The following sample includes generated HTML code for the 8.x JICA client, which supported both Sun and MS JVMs:

```html
<html>
<head>
<title>CitrixJICA Applet.</title>
<meta http-equiv="Content-Type" content="text/html; charset=iso-8859-1">
</head>
<!--
Notes:
1) << CODEBASE >> is a system value that will get
replaced at the time the applet is launched.
Please do not modify this value.
2) Please modify the remaining values as needed.
3) Please make sure all attribute names/values are
enclosed in double quotes.
-->
<body>
    <applet code="com.citrix.JICA"
            codebase="<< CODEBASE >>"
            archive="JICAEngN.jar,JICA-sicaN.jar,cryptojN.jar,JICA-configN.jar,JICA-coreN.jar"
            width="640" height="480"
            name="CitrixJICA" align="top">
    <param name="code" value="com.citrix.JICA">
    <param name="codebase" value="<< CODEBASE >>">
    <param name="archive" value="JICAEngN.jar,JICA-sicaN.jar,cryptojN.jar,JICA-
configN.jar,JICA-coreN.jar">
    <param name="cabbase" value="cryptojM.cab,JICA-configM.cab,JICAEngM.cab,JICA-
sicaM.cab,JICA-coreM.cab">
    <param name="name" value="CitrixJICA">
    <param name="width" value="640">
    <param name="height" value="480">
    <param name="align" value="top">
    <!--
Please specify additional params here after the comment.
    <param name="paramname" value="paramvalue">
    -->
    <param name="Address" value="__PLEASE_SPECIFY__">
    <param name="Username" value="<< user >>">
    <param name="password" value="<< password >>">
    <param name="EncryptionLevel" value="1">
    <param name="BrowserProtocol" value="HTTPonTCP">
    </applet>
</body>
</html>
```

Related Documentation

- Accessing Java Applet Bookmarks

- Configuring Hosted Java Applet Resource Profile Bookmarks

CHAPTER 18 Citrix Templates

About Citrix Templates

The system supports several mechanisms for intermediating traffic between a Citrix server and client, including the Juniper Networks Citrix Terminal Services proxy, JSAM, WSAM, VPN Tunneling, and the hosted Java applets feature.

The Citrix Web template enables you to easily configure access to a Citrix server using the Juniper Networks Citrix Terminal Services proxy, JSAM, or WSAM. The Citrix Web template is a resource profile that controls access to Citrix applications and configures Citrix settings as necessary. Citrix Web templates significantly reduce your configuration time by consolidating configuration settings into one place and by prepopulating a variety of resource policy settings for you depending on the type of Citrix setup you select. You should use the Citrix Web template if you have the Citrix Web Interface already installed in your environment or if you are using a Web server to host your ICA files.

Because of their highly simplified configurations, templates are the ideal Citrix configuration method if you want to deliver ActiveX or Java applets from a third-party Web server through the system.

Citrix Web templates simplify your configuration by automatically detecting whether the Citrix Web client or the Citrix Java applet is being used and employing the appropriate access mechanism accordingly. For instance, if you have configured the Citrix Web Interface to deliver a Java client, the system automatically uses its Java rewriting engine to tunnel traffic. If you have configured the Citrix Web Interface to deliver an ActiveX client, the system uses its Citrix Terminal Services feature, JSAM, or WSAM (depending on the option you select) to tunnel traffic.

We strongly recommend using Citrix templates instead of the traditional role and resource policy configuration options available through the system.

ⓘ   NOTE: Pulse Secure does not support saving a Citrix application shortcut to the desktop through the system when the loopback IP address is running on the client. Double-clicking this shortcut returns an error as it does not use WSAM or JSAM.

Related Documentation

Comparing Access Mechanisms for Configuring Citrix

Connect Secure supports several mechanisms for intermediating traffic between a Citrix server and client, including the Citrix Terminal Services proxy, JSAM, WSAM, VPN Tunneling, and the hosted Java applets feature.

Table 78 describes key differences when accessing a Citrix Metaframe Server thought a Citrix Web Interface server. The descriptions in this table focus on configuring Citrix Terminal Services, JSAM, and WSAM through Web resource profile templates (Select Users > Resource Profiles > Web, click New Profile and select Citrix Web interface/JICA from the Type list.)

> **NOTE:** If you want to configure access to a Citrix Metaframe server through a Citrix Web Interface server, you must use Web resource profile templates. If you want to configure access to a Citrix Metaframe server without using a Citrix Web Interface server, you must use a standard Citrix Terminal Services or WSAM resource profile or role.

Table78: Accessing the Citrix Web Interface Server Using Web Resource Profile Templates

Requirement	Terminal Services	JSAM	WSAM
User experience	1. The user clicks a Citrix Web Interface bookmark in the Web Bookmarks section of the end user console. 2. The user is taken to the Citrix Web Interface (WI) sign-in page (assuming you do not configure FORM POST SSO). 3. Once the user signs into the WI portal (either manually or automatically through SSO), he is taken to the Citrix WI portal page, which contains the list of published applications in icon form. 4. When the user clicks the published application, the Juniper Networks Citrix Terminal Services (CTS) proxy launches and the ICA traffic is tunneled through the Juniper Networks CTS proxy.	1. The user launches JSAM. 2. The user clicks a Citrix Web Interface bookmark in the Web Bookmarks section of the end user console. 3. The user is taken to the Citrix Web Interface (WI) sign-in page (assuming you do not configure FORM POST SSO). 4. Once the user signs into the WI portal (either manually or automatically through SSO), he is taken to the Citrix WI portal page, which contains the list of published applications in icon form. 5. When the user clicks the published application, the ICA traffic is tunneled through JSAM.	1. The user launches WSAM. 2. The user clicks a Citrix Web Interface bookmark in the Web Bookmarks section of the end user console. 3. The user is taken to the Citrix Web Interface (WI) sign-in page (assuming you do not configure FORM POST SSO). 4. Once the user signs into the WI portal (either manually or automatically through SSO), he is taken to the Citrix WI portal page, which contains the list of published applications in icon form. 5. When the user clicks the published application, the ICA traffic is tunneled through WSAM.
Accessing published applications from Mac or Linux	Not supported on Mac and Linux.	Supported on Mac and Linux.	Not supported on Mac and Linux.
Configuring ports	Automatically monitor all traffic on port 1494 if session reliability is turned off on the server. The system monitors port 2598 if session reliability is turned on. You do not need to specify which ports to monitor or which applications to intermediate.	You must specify which ports to monitor. This enables you to access published applications that use ports other than 1494.	You do not need to specify which ports to monitor or which applications to intermediate. WSAM works in app mode and monitors all traffic coming from certain Citrix executables.

Administrator privileges	If a Citrix Web client is not installed on the user's desktop, administrator privileges are required. This is a limitation of the installation of the Citrix client. To install and run the Juniper Networks Citrix Terminal Services proxy client, administrator privileges are not required.	If a Citrix Web client is not installed on the user's desktop, administrator privileges are required. This is a limitation of the installation of the Citrix client. To run JSAM, administrator privileges are not required.	Requires administrator privileges to install WSAM.
Modifying host file	Does not require modification of the etc/hosts file.	Does not require modification of the etc/hosts file.	Does not require modification of the etc/hosts file.

Table 79 describes key differences when accessing a Citrix Metaframe Server without using a Citrix Web Interface server. The descriptions in this table focus on configuring Citrix Terminal Services, JSAM, and WSAM through standard resource profiles (Select Users > Resource Profiles > SAM or Terminal Services.)

Table79: Accessing Citrix Metaframe Server Without Using a Citrix Web Interface Server

Requirement	Terminal Services	JSAM	WSAM
User experience	The user launches the published application by clicking the bookmark or icon in the Terminal Services section of the end user console.	1. JSAM auto-launches when the user signs into the device or the user launches JSAM manually. 2. The user launches the published application using standard methods such as the Windows Start menu or a desktop icon.	1. WSAM auto-launches when the user signs into the device or the user launches WSAM manually. 2. The user launches the published application using standard methods such as the Windows Start menu or a desktop icon.
Accessing published applications from Mac or Linux	Macintosh and Linux users cannot access published applications from a Citrix Metaframe server.	Macintosh and Linux users can access published applications from a Citrix Metaframe server.	Macintosh and Linux users cannot access published applications from a Citrix Metaframe server.
Admin configuration	You can specify which ports the system intermediates. If you do not configure this information, the system automatically monitors ports 1494 and 2598.	You cannot configure Citrix as a standard application. Instead, you need to create a custom JSAM application, provide the server names of all Metaframe servers, and specify which ports to monitor. This enables you to use applications such as Citrix Secure Gateways (CSGs) and published applications that use ports other than 1494.	You must specify which ports and applications the system monitors. This enables you to use applications such as Citrix Secure Gateways (CSGs) and published applications that use ports other than 1494.
Administrator privileges	If a Citrix Web client is not installed on the user's desktop, administrator privileges are required. This is a limitation of the installation of the Citrix client. To install and run the Juniper Networks Citrix Terminal Services proxy client, administrator privileges are not required.	Requires administrator privileges to run JSAM because etc/hosts file modifications are required.	Requires administrator privileges to install WSAM.
Modifying host file	Does not require modification of the etc/hosts file.	Requires modification of the etc/hosts file.	Does not require modification of the etc/hosts file.

Related Documentation

- About Citrix Templates

- Creating Resource Profiles Using Citrix Web Applications

- Creating Resource Profiles for Citrix Clientless Access

Creating Resource Profiles Using Citrix Web Applications

The Citrix Web template enables you to easily configure Citrix access using the Juniper Networks Citrix Terminal Services proxy, JSAM, or WSAM.

To create a resource profile using the Citrix template:

1. Select **Users > Resource Profiles > Web** in the admin console.

2. Click **New Profile**.

3. Select Citrix Web Interface/JICA from the Type list.

4. Enter a unique name and optionally a description for the Citrix resource profile.

5. Enter the URL of the Web server that hosts your ICA files in the Web Interface (NFuse) URL field. Use the format: [protocol://]host[:port][/path]. For instance, enter the URL of an NFuse server, the Web interface for a Citrix Metaframe Presentation Server, or a Web server from which the system can download Citrix Java applets or Citrix cab files. (The system uses the specified URL to define the default bookmark for the Citrix resource profile.) You may enter a directory URL or a file URL.

6. Specify which type of Citrix implementation you are using in your environment by selecting one of the following options:

   - Java ICA Client with Web Interface (NFuse)—Select this option if you have deployed the Citrix Web Interface for MPS (that is, NFuse) to deliver Java ICA clients.

   - Java ICA Client without Web Interface (NFuse)—Select this option if you have deployed a generic Web server to deliver Java ICA clients.

   - Non-Java ICA Client with Web Interface (NFuse)—Select this option if you have deployed the Citrix Web Interface for MPS (that is, NFuse) to use any of the different clients (Java, ActiveX, local).

   - Non-Java ICA Client without Web Interface (NFuse)—(Read only) If you have deployed a non-Java ICA client without the Citrix Web Interface for MPS (that is, NFuse), you cannot create a Citrix resource profile through this template. Instead, click the client application profile link beneath this option. The link brings you to the Client Application Profiles page, where you can create a SAM resource profile.

7. From the Web Interface (NFuse) version list, select which Citrix version you are using. (The system uses this value to pre-populate the Forms POST SSO values in your single sign-on autopolicy.

8. Specify the Metaframe Servers to which you want to control access in the MetaFrame servers area. Then click **Add**. When specifying servers, you can enter wildcards or IP ranges.

   The system uses the values that you enter to automatically create a corresponding resource policy that enables access to the necessary resources:

   - If you select either **Java ICA Client with** or **without Web Interface**, the system creates a corresponding Java ACL resource policy that enables Java applets to connect to the specified Metaframe servers.

- If you select **Non-Java ICA Client with Web Interface**, and then you select **ICA client connects over WSAM or JSAM**, the system creates a corresponding SAM resource policy that enables users to access the specified Metaframe servers.

- If you select **Non-Java ICA Client with Web Interface**, and then you select **ICA client connects over CTS**, the system creates corresponding Terminal Services and Java resource policies that enable users to access the specified Metaframe servers.

9. (Java ICA clients only.) If you deployed Citrix using a Java ICA Client, select the **Sign applets with uploaded code-signing certificate(s)** check box to re-sign the specified resources using the certificate uploaded through the System > Configuration > Certificates > Code-signing Certificates page of the admin console.

   When you select this option, the system uses all of the "allow" values that you enter in the resource profile's Web access control autopolicy to automatically create a corresponding code-signing resource policy. Within this policy, the system uses the specified Web resources to create a list of trusted servers.

10. (Non-Java ICA clients only) If you have deployed Citrix using a non-Java ICA Client with a Web interface, you must use the Juniper Networks Citrix Terminal Services proxy, Secure Application Manager, or VPN Tunneling to secure traffic to your Metaframe servers instead of the Content Intermediation Engine.

    To secure traffic through the Juniper Citrix Terminal Services proxy or the Secure Application Manager, select one of the following options in the ICA Client Access section:

    - ICA client connects over CTS Client—Select this option to secure your Citrix traffic through the Citrix Terminal Services client (if your users are using Active X clients) or Java rewriting engine (if your users are using Java clients). (When you select this option, the system automatically enables the Terminal Services option on the Users > User Roles > *Select_Role* > General > Overview page of the admin console.)

      NOTE: If you are using a third-party Web server such as your company's Intranet server to deliver the ICA file, make sure the Content-Type of the HTTP Response header is application/x-ica. Only then does the system automatically intermediate the ICA file and launch its Citrix Terminal Services client to tunnel the traffic.

      NOTE: If you select this option, we recommend that you disable Citrix client downloads through the Citrix Web Interface. Otherwise, users could inadvertently start two different windows downloading two versions of the Citrix client simultaneously–one through the system (which automatically attempts to download the Citrix client if one is not present on the user's computer) and one through the Citrix Web Interface.

    - ICA client connects over WSAM—Select this option to secure traffic using WSAM. (When you select this option, the system automatically enables the Secure Application Manager option on the Users > User Roles > *Select_Role* > General > Overview page of the admin console.)

    - ICA client connects over JSAM—Select this option to secure traffic using JSAM. Then, configure the following options:

      - Number of Servers/Applications—Enter the lesser of the following two numbers: maximum number of Citrix servers in your environment or the maximum number of published applications that a user can open simultaneously. For instance, if your environment contains one server and five published applications, enter 1 in this field. Or, if your environment contains 20 servers and 10 published applications, enter 10 in this field. The maximum value this field accepts is 99.

- Citrix Ports—Specify the ports on which the Metaframe servers listen.

  When you select the ICA client connects over JSAM option, the system automatically enables the Secure Application Manager option on the Users > User Roles > *Select_Role* > General > Overview page of the admin console.

  > *i* NOTE: You cannot enable WSAM and JSAM for the same role. Therefore, if you try to create a Citrix resource profile that uses one of these access mechanisms (for instance, JSAM) and another profile associated with role already uses the other access mechanism (for instance, WSAM), the system does not enable the new access mechanism (JSAM) for the role. Also note that you can only use WSAM or JSAM to configure access to one Citrix application per user role.

11. (Non-Java ICA Client with Web Interface only.) If you want to allow users to access local resources such as printers and drives through their Citrix Web Interface sessions, select the **Configure access to local resources** check box. Then, select from the following options:

- Select **Connect printers** if you want to enable the user to print information from the terminal server to his local printer.

- Select **Connect drives** if you want to enable the user to copy information from the terminal server to his local client directories.

- Select **Connect COM Ports** if you want to enable communication between the terminal server and devices on the user's serial ports.

  > *i* NOTE: To control access to local resources exclusively through your Citrix Metaframe server settings, clear the Configure access to local resources check box. When you clear the option, the Metaframe server settings take effect. Or, if you want to selectively override Citrix Metaframe server settings for the bookmark, select the Configure access to local resources check box and then specify the local resources to which you want to enable or disable access. Note that if you enable access to a local resource through the system, you still must enable access to it through the Metaframe server as well.

  > When you enable local resources through the terminal server, each user can only access his own local resources. For instance, user 1 cannot see user 2's local directories.

12. Select the **Autopolicy: Web Access Control** check box to create a policy that allows or denies users access to the resource specified in the Web Interface (NFuse) URL field. (By default, the system automatically creates a policy for you that enables access to the resource and all of its subdirectories.)

13. If you selected one of the Web interface options above, update the SSO policy created by the Citrix template. Select the **Autopolicy: Single Sign-on** check box. (Single sign-on autopolicies configure the system to automatically pass data such as usernames and passwords to the Citrix application. The system automatically adds the most commonly used values to the single sign-on autopolicy based on the Citrix implementation you choose.)

When you select single sign-on, the WIClientInfo and WINGSession cookies are prepopulated automatically in addition to the POST Resource and URL.

Or, if you selected the non-Web interface option, you may optionally create your own single sign-on autopolicy.

14. Click **Save and Continue**.

15.   Select the roles in the Roles tab to which the Citrix resource profile applies and click **Add**.

The selected roles inherit the autopolicies and bookmarks created by the Citrix resource profile. If it is not already enabled, the system also automatically enables the Web option in the Users > User Roles > *Select_Role* > General > Overview page of the admin console and the Allow Java Applets option in the Users > User Roles > *Select_Role* > Web > Options page of the admin console for all of the roles you select.

Also enable the Terminal Services access feature under User Roles > *Select_Role* > General Overview. If the user role does not have this feature enabled, the Citrix ICA file is delivered as is (without being rewritten) and the Juniper Citrix component (CTS) will not start. In this case, the Citrix native client attempts to establish a connection with the back-end server directly (without going through the system) and will fail.

16.   Click **Save** Changes.

17.   (Optional.) In the Bookmarks tab, modify the default bookmark created by the system and/or create new ones.

By default, the system creates a bookmark to the Web interface (NFuse) URL defined in the Web Interface (NFuse) URL field and displays it to all users assigned to the role specified in the Roles tab.

Related Documentation

- About Citrix Templates

- Creating WSAM Client Application Resource Profiles

- Creating a JSAM Application Resource Profile

- Creating Resource Profiles for Citrix Clientless Access

Creating Resource Profiles for Citrix Clientless Access

If you have the Citrix StoreFront, you can create a Web template to allow users to access Citrix applications without the need for a Citrix client. Users must have one of the following browser versions (or later) to support HTML5 and Websockets:

- Internet Explorer 10

- Safari 6

- Google Chrome 23

- Mozilla Firefox 17

    i   NOTE: You can collect all the logs related to this feature using **hprewrite-server** as the process name.

To create a resource profile using the Citrix template:

1. Select **Users > Resource Profiles > Web** in the admin console.

2. Click **New Profile**.

3. Select **Citrix StoreFront (HTML Receiver)** from the Type list.

4. Enter a unique name and optionally a description for the Citrix resource profile.

5. Enter the URL of the Citrix StoreFront Web server in the Base URL field. Use the format: [protocol://]host[:port][/path]. The system uses the specified URL to define the default bookmark for the Citrix resource profile. You may enter a directory URL or a file URL.

6. Select the **Autopolicy: Web Access Control** check box to create a policy that allows or denies users access to a specific resource under the Base URL. Enter the full URL of the resource, select **Allow** or **Deny**, and click **Add**. By default, the system automatically creates a policy that enables access to the resource and all of its subdirectories.

7. Select the **Autopolicy: Single Sign-on** check box to automatically pass data such as usernames and passwords to the Citrix application. The system automatically adds the most commonly used values to the single sign-on autopolicy.

8. If you want to perform a form POST when a user makes a request to the resource specified in the Resource field, select the **POST the following data** check box and specify the following:

   1. In the Resource field, specify the application's sign-in page, such as: http://my.domain.com/public/login.cgi. Wildcard characters are not supported in this field.

      To automatically post values to a specific URL when an end user clicks on a system bookmark, the resource that you enter here must exactly match the URL that you specify in the Base URL field.

   2. In the Post URL field, specify the absolute URL where the application posts the user's credentials, such as: http://yourcompany.com/login.cgi. You can determine the appropriate URL using a TCP dump or by viewing the application's sign-in page source and searching for the POST parameter in the FORM tag.

   3. Select the **Deny direct login for this resource** check box if you do not want to allow users to manually enter their credentials in a sign-in page. Users may see a sign-in page if the form POST fails.)

   4. Select the **Allow multiple POSTs to this resource** check box if you want to send POST and cookie values to the resource multiple times if required. If you do not select this option, the system does not attempt single sign-on when a user requests the same resource more than once during the same session.

   5. Optionally specify the following for each item of user data you want to post and click **Add**:

      - **Label**—The name used to identify the data.

      - **Name**—The name used to identify the data in the Value field. The back-end application should expect this name.

      - **Value**—The value to post to the form for the specified Name. You can enter static data, a system variable, or system session variables containing username and password values.

      - **User modifiable?**—Select **Not modifiable** to prevent users from changing the information in the Value field. Select **User CAN change value** to allow users to specify data for a back-end application. Select **User MUST change value** if users must enter additional data to access a back-end application. If users can or must change the value, a field for data entry appears on the user's Advanced Preferences page. This field is labeled using the name in the Label field. If you enter a value in the Value field, this data appears in the field but is editable.

   6. To post header data to the specified URL when a user makes a request to a resource specified in the Resource field, select the **Send the following data as request headers** check box. Then:

      1. In the Resource section, specify the resources to which this policy applies.

      2. Optionally specify the header data to post by entering data in the following fields and clicking **Add**:

         - **Header name**—The text to send as header data.

- **Value**—The value for the specified header.

3. Click **Save and Continue**.

4. Select the roles in the Roles tab to which the Citrix resource profile applies and click **Add**.

   The selected roles inherit the autopolicies and bookmarks created by the Citrix resource profile. If it is not already enabled, the system also automatically enables the Web option in the Users > User Roles > *Select_Role* > General > Overview page of the admin console and the Allow Java Applets option in the Users > User Roles > *Select_Role* > Web > Options page of the admin console for all of the roles you select.

5. Click **Save** Changes.

6. (Optional.) Select the **Bookmarks** tab to modify the default bookmark created by the system and/or create new bookmarks. By default, the system creates a bookmark for the URL defined in the Base URL field and displays it to all users assigned to the role specified in the Roles tab.

Related Documentation

- About Citrix Templates

- Creating Resource Profiles Using Citrix Web Applications

CHAPTER 19 Lotus iNotes Templates

Creating Resource Profiles Using the Lotus iNotes Template

A Lotus iNotes template is a resource profile that controls access to the Web application and configures iNotes settings as necessary. Lotus iNotes templates significantly reduce your configuration time by consolidating settings into one place and by prepopulating a variety of resource policy settings for you depending on the type of setup you select.

The system supports intermediating traffic to Lotus iNotes through a Web rewriting resource profile template, JSAM, WSAM, and VPN Tunneling. This topic describes how to configure access using the Web rewriting template. The prepopulated values vary depending on the version of iNotes you select and are based on the most common deployment of the servers.

To create a resource profile using the Lotus iNotes template:

1. Select **Users > Resource Profiles > Web** in the admin console.

2. Click **New Profile**.

3. Select the Lotus Notes version from the Type list.

4. Enter a unique name and optionally a description for the Lotus Notes resource profile.

5. Enter the URL of the Lotus iNotes resource to which you want to control access in the Base URL box. Use the format: [protocol://]host[:port][/path]. The system uses the specified URL to define the default bookmark for the Lotus iNotes resource profile. You may enter a directory URL or a file URL.

6. Under iNotes setting, select **Allow caching on client** to let Web browsers store non-user data, such as Javascript and CSS files, on a user's machine. Select **Minimize caching on client** to allow the system to send a cache-control:no-store header or a cache-control:no-cache header based on the user's Web browser and content type. This is the same as smart caching.

   The Allow caching on client option caches content that the backend iNotes server typically caches. This caching option improves performance by using the cached content instead of retrieving the content from the server the next time the page displays. The Minimize caching on client option provides security by sending a cache-control:no-store header or a cache-control:no-cache header to either not store content or to re-validate the cached content each time it is requested. With both caching option, you can choose to either allow or prevent the uploading or downloading of attachments.

7. Select the **Prevent download of attachments** check box to prohibit users from downloading attachments to their systems. Select the **Prevent upload of attachments** check box (available only for Lotus iNotes 6.5 and Lotus iNotes 7) to prevent users from transmitting (uploading) attachments to the system.

8. Select the **Autopolicy: Web Access Control** check box to create a policy that allows or denies users access to the Web resource (and all of its subdirectories) listed in the Resource field.

   1. In the Resource box, specify the Web server or HTML page to which you want to control access using the format: [protocol://]host[:port][/path].

   2. From the Action list, select **Allow** to enable access to the specified resource or **Deny** to block access to the specified resource.

   3. Click **Add**.

4. Select the **Autopolicy: Caching** check box to specify the resources to which this policy applies in the Resource box.

   > **NOTE:** The correct caching resource policy must be configured to allow end users to open and save e-mail attachments of different document types in iNotes. For example, if the caching policy is set to Smart, end users cannot save .htm or .html attachments to disk.

5. Select the **Autopolicy: Web Compression** check box to create a policy that specify which types of Web data the system should and should not compress.

   1. In the Resources field, specify the resources to which this policy applies.

   2. Select one of the following options from the Action list:

      - Compress—Compresses the supported content types from the specified resource.

      - Do not compress—Do not compress the supported content types from the specified resource.

   3. Click **Add**.

4. Select the **Autopolicy: Single Sign-On** check box to pass data such as the username and password to the Lotus iNotes application.

5. Click **Save and Continue**.

6. Select the roles to which the Lotus iNotes resource profile applies in the Roles tab and click **Add**.

   The selected roles inherit the autopolicies and bookmarks created by the Lotus iNotes resource profile. If it is not already enabled, the system also automatically enables the Web option in the Users > User Roles > *Select Role* > General > Overview page of the admin console.

7. Click **Save Changes**.

8. (Optional.) In the Bookmarks tab, modify the default bookmark created by the system and/or create new ones.

Related Documentation

- About Hosted Java Applet Templates

- Creating WSAM Client Application Resource Profiles

- Creating a JSAM Application Resource Profile

- Creating Resource Profiles Using Citrix Web Applications

CHAPTER 20 Microsoft OWA Templates

- Creating Resource Profiles Using the Microsoft OWA Template

Creating Resource Profiles Using the Microsoft OWA Template

A Microsoft Outlook Web Access (OWA) template is a resource profile that controls access to the application and configures OWA settings as necessary. OWA templates significantly reduce your configuration time by consolidating configuration settings into one place and by prepopulating a variety of resource policy settings for you depending on the type of setup you select.

The system supports intermediating traffic to Microsoft OWA through a Web rewriting resource profile template, JSAM, WSAM, and VPN Tunneling. This topic describes how to configure access using the Web rewriting template. The prepopulated values vary depending on the version of OWA you select and are based on the most common deployment of the servers.

To create a resource profile using the Microsoft OWA template:

1. Select **Users > Resource Profiles > Web Applications/Pages** in the admin console.

2. Click **New Profile**.

3. Select your Microsoft OWA version from the Type list.

4. Enter a unique name and optionally a description for the Citrix resource profile.

5. Enter the URL of the OWA resource to which you want to control access In the Base URL box. Use the format: [protocol://]host[:port][/path]. The system uses the specified URL to define the default bookmark for the OWA resource profile. You may enter a directory URL or a file URL.

6. Under OWA settings select the following options,

    1. (OWA 2000 and OWA 2003.) Select **Allow caching on client** to let Web browsers store non-user data, such as Javascript and CSS files, on a user's machine.

       The Allow caching on client option caches content the backend OWA server typically caches. This caching option improves performance by using the cached content instead of retrieving the content from the server the next time the page displays.

    2. (OWA 2000 and OWA 2003.) Select **Minimize caching on client** to allow the system to send a cache-control:no-store header or a cache-control:no-cache header (do not store content or revalidate the cached content each time it is requested) based on the user's Web browser and content type. This is the same as smart caching.

    3. (OWA 2007.) Select **Managed Device** to cache files. If you configure a Form post SSO, the trusted parameter is set to 4. This indicates the end user's device is private.

    4. (OWA 2007.) Select **Unmanaged Device** to not cache files. If you configure a Form post SSO, the trusted parameter is set to 0. This indicates the end user's device is public.

       i   NOTE: If it is necessary to download an attachment, the file is cached even though you select Unmanaged Device.

---

5. Select **Prevent download of attachments** to prohibit users from downloading attachments to their systems.

6. Select **Prevent upload of attachments** to prevent users from transmitting (uploading) attachments to the system.

7. Under Autopolicy: Web Access Control, create a policy that allows or denies users access to the Web resource (and all of its subdirectories) listed in the Resource field.

    1. Specify the Web server or HTML page to which you want to control access in the Resource field. Use the format: [protocol://]host[:port][/path].

    2. Select **Allow** to enable access to the specified resource or Deny to block access to the specified resource from the Action list.

    3. Click **Add**.

4. Under Autopolicy: Caching, specify the resources to which this policy applies in the Resource box.

    > ⓘ NOTE: The correct caching resource policy must be configured to allow end users to open and save e-mail attachments of different document types in OWA. For example, if the caching policy is set to Smart, end users cannot save .htm or .html attachments to disk.

5. Under Autopolicy: Web Compression, create a policy that specifies which types of Web data the system should and should not compress.

    1. Specify the resources to which this policy applies in the Resources box.

    2. Select one of the following options from the Action list:

        - Compress—Compress the supported content types from the specified resource.

        - Do not compress—Do not compress the supported content types from the specified resource.

    3. Click **Add**.

4. Select the **Autopolicy: Single Sign-On** check box to pass data such as the username and password to the OWA application.

5. Click **Save and Continue**.

6. Select the roles to which the resource profile applies in the Roles tab and click **Add**.

    The selected roles inherit the autopolicies and bookmarks created by the Microsoft OWA resource profile. If it is not already enabled, the system also automatically enables the Web option in the Users > User Roles > Select _Role > General > Overview page of the admin console.

7. Click **Save Changes**.

8. (Optional.) Modify the default bookmark created by the system in the Bookmarks tab, and/or create new ones.

Related Documentation

- About Hosted Java Applet Templates

- Creating WSAM Client Application Resource Profiles

- Creating a JSAM Application Resource Profile

- Creating Resource Profiles Using Citrix Web Applications

CHAPTER 21 Microsoft Sharepoint Templates

Creating Resource Profiles Using the Microsoft Sharepoint Template

A Microsoft Sharepoint template is a resource profile that controls access to the application and configures Sharepoint settings as necessary. Microsoft Sharepoint templates significantly reduce your configuration time by consolidating configuration settings into one place and by pre-populating a variety of resource policy settings for you depending on the type of setup you select.

The system supports intermediating traffic to Microsoft Sharepoint through a Web rewriting resource profile template, JSAM, WSAM, and VPN Tunneling. This topic describes how to configure access using the Web rewriting template.

Sharepoint 2010 and Sharepoint 2013 are supported, with the caveats listed in Table 80.

Table80: Sharepoint 2010 and 2013 Caveats

Version	Caveats
Sharepoint 2010	- Office Web Apps through Windows Live is not supported. - OneNote document support is limited to only notebooks created to be stored on a local computer and then published on SharePoint 2010. - Office documents residing on SharePoint server cannot be opened with Microsoft Office when the server is accessed through the system. NOTE: In the current release, we support sending contact information from Sharepoint to your Outlook client through the Content Intermediation Engine (Web rewriting feature). Transferring the contact information to the backend Exchange server requires WSAM, JSAM, or VPN Tunneling. To import contact information into the Sharepoint server from your Outlook client, first export your contacts and then upload them to the Sharepoint server.
Sharepoint 2013	- Active Directory Federation Services and claims-based authentication are not supported in this release. - Integrated service with Exchange 2013 and Lync 2013 is not supported.

To create a resource profile using the Microsoft Sharepoint template:

1. Select **Users > Resource Profiles > Web** in the admin console.

2. Click **New Profile**.

3. Select **Microsoft Sharepoint** from the Type list.

4. Enter a unique name and optionally a description for the Sharepoint resource profile.

5. Enter the URL of the Sharepoint resource to which you want to control access in the Base URL field. Use the format: [protocol://]host[:port][/path]. The system uses the specified URL

to define the default bookmark for the Sharepoint resource profile. You may enter a directory URL or a file URL.

6. Under Sharepoint Settings, select **Allow in-line editing of documents within explorer view** to allow users to modify files displayed in the explorer view.

> ⓘ NOTE: This option is supported only if you enable persistent session (**User > User Roles >** *RoleName* > **General > Session Options**.)

   1. Enter the URL to the Explorer View page, and then click Add. Do not enter a value that resolves to non-Explorer View URLs (such as http://*:*). Doing so might cause Explorer View to not launch.

   2. Order the resources in your list, if appropriate, by selecting the check box next to an item and then using the up and down arrows to move it to the correct place in the list.

   3. Enter the number of minutes a persistent cookie resides on a user's computer before it expires in the Persistent cookie timeout box.

   > ⓘ NOTE: Do not confuse this timeout option with Max. Session Length, which determines the number of minutes an active nonadministrative user session may remain open before ending.

4. Select **Add Web ACL** if you have Sharepoint 2013 and the Office Web Apps are on a separate server. The cursor is moved to the Resource text box where you can enter the URL for the Office Web Apps server (see Step 8).

5. Under Autopolicy: Web Access Control, create a policy that allows or denies users access to the Web resource (and all of its subdirectories) listed in the Resource box.

   1. Specify the Web server or HTML page to which you want to control access in the Resource box. Use the format: [protocol://]host[:port][/path].

   2. Select **Allow** to enable access to the specified resource or Deny to block access to the specified resource from the Action list.

   3. Click **Add**.

4. (Optional.) Click **Show ALL autopolicy types** to create additional autopolicies that fine-tune access to the resource. Then, create the autopolicies.

5. Click **Save and Continue**.

6. Select the roles to which the resource profile applies in the Roles tab, and click **Add**.

   The selected roles inherit the autopolicies and bookmarks created by the Microsoft Sharepoint resource profile. If it is not already enabled, the system also automatically enables the Web option in the Users > User Roles > Select Role > General > Overview page of the admin console.

7. Click **Save Changes**.

8. (Optional.) Modify the default bookmark created by the system in the Bookmarks tab or create new ones.

Related Documentation

- About Hosted Java Applet Templates

- Creating WSAM Client Application Resource Profiles

- Creating a JSAM Application Resource Profile

- Creating Resource Profiles Using Citrix Web Applications

CHAPTER 22 Web Rewriting

Web Rewriting

The Web rewriting feature enables you to intermediate Web URLs through the Content Intermediation Engine. You can intermediate URLs on the World Wide Web or on your corporate Intranet.

When you intermediate standard Web content, you can create supplemental policies that "fine-tune" the access requirements and processing instructions for the intermediated content. You can create these supplemental policies through resource profiles (recommended) or resource policies.

Standard Web rewriting policy types include:

- Web access control—Web access policies control which Web resources users can access in order to connect to the Internet, intranet, or extranet.

- Single sign-on—Single sign-on policies enable you to automatically pass user credentials to a Web application. You can configure single sign-on policies to intercept basic authentication and NTLM challenges or post the credentials and headers that you specify to the Web application.

- Caching—Caching policies control which Web content the system caches on a user's machine.

- Java—Java policies control to which servers and ports Java applets can connect. These policies also specify trusted servers for which the system resigns content.

- Rewriting—Rewriting policies specify resources that the system should not intermediate, minimally intermediation, or only intermediate selectively.

- Web compression—Web compression policies specify which types of Web data the system should and should not compress.

- Web proxy—(Resource policies only) Web proxy resource policies specify Web proxy servers for which the system should intermediate content. Note that the system intermediates both forward and backwards proxies, but only enables single sign-on to trusted proxies.

- Launch JSAM—(Resource policies only) Launch JSAM policies specify URLs for which the system automatically launches J-SAM on the client. This feature is useful if you enable applications that require J-SAM but do not want to require users to run J-SAM unnecessarily.

- Protocol—(Resource policies only) Protocol resource policies enable or disable HTTP 1.1 protocol support on the system.

- Options— (Resource policies only) You can enable IP based matching for hostnames as well as case-sensitive matching for path and query strings in Web resources through resource policy options.

Web rewriting is a standard feature on Connect Secure devices.

Related Documentation

- Task Summary: Configuring the Web Rewriting Feature

- Writing a Web Access Resource Policy

- Defining a Single Sign-On Autopolicy

- Defining a Caching Autopolicy

- Writing a Java Access Control Resource Policy

- Defining a Rewriting Autopolicy

- Defining a Web Compression Autopolicy

- Writing a Web Proxy Resource Policy

- Automatically Launching JSAM

- Writing An HTTP 1.1 Protocol Resource Policy

- Defining Resource Policies: General Options

Silverlight Support

The system supports Silverlight for rewriting for:

- Sharepoint 2010 – only the default Silverlight pages on Sharepoint 2010

- OWA 2010 - including file attachments in e-mails

The system does not support custom XAP packages (custom Silverlight applications that a user can upload to a Sharepoint site).

No configuration steps are required for Silverlight support.

To support Silverlight on Sharepoint 2010, a system-generated rewriting resource policy is added for ActiveX with the following details:

Classid= DFEAF541-F3E1-4C24-ACAC-99C30715084A
Parameter= initParams:mediaSource,previewImageSource
Action= Rewrite URL and response (Static HTML only)

This policy is loaded during both upgrades and new installations. Do not remove this resource policy.

The colon (:) in the above Parameter field means the object tag contains the **initParams** parameter followed by comma separated name-value pairs.

For example, the above policy works for the following tag:

```
<param name="initParams"
value="mediaSource=/silverlight/media/Wildlife.wmv,previewImageSource=/silverli
ght/image/VideoPreview.png" />
```

Task Summary: Configuring the Web Rewriting Feature

i NOTE: When intermediating content through the content intermediation engine, we recommend that the GMT time on both Connect Secure and the backend Web application server be the same. This prevents any premature expiration of cookies if the Connect Secure system time is later than the Web application server time.

To configure the Web rewriting feature:

1. Create resource profiles that enable access to web sites, create supporting autopolicies (such as single sign-on and Java access control policies) as necessary, include bookmarks that link to the web fsites, and assign the policies and bookmarks to user roles using settings in the Web Applications Resource Profiles page (Users > Resource Profiles > Web) of the admin console.

   We recommend that you use resource profiles to configure Web rewriting (as described above). However, if you do not want to use resource profiles, you can configure Web rewriting using role and resource policy settings in the following pages of the admin console instead:

   1. Create resource policies that enable access to web sites using settings in the Users > Resource Policies> Web > Web ACL page of the admin console.

   2. As necessary, create supporting resource policies (such as single sign-on and Java access control policies) using settings in the Users > Resource Policies> Select Policy Type pages of the admin console.

   3. Determine which user roles may access the web sites that you want to intermediate, and then enable Web access for those roles through the Users > User Roles > Select Role > General > Overview page of the admin console.

   4. Create bookmarks to your web sites using settings in the Users > User Roles > Select Role > Web > Bookmarks page of the admin console.

   5. As necessary, enable Web general options that correspond to the types of Web content you are intermediating (such as Java) using settings in the Users > User Roles > Select Role > Web > Options page of the admin console.

2. After enabling access to Web applications or sites using Web rewriting resource profiles or roles and resource policies, you can modify general role and resource options in the following pages of the admin console:

   1. (Optional) Set additional Web browsing options (such as allowing users to create their own bookmarks or enabling hostname masking) Users > User Roles > Select Role > Web > Options page of the admin console.

      i NOTE: Even if you enable hostname masking, links corresponding to protocols not rewritten by Web rewriting are not obfuscated. For example, ftp://xyz.pulsesecure.net and file://fileshare.pulsesecure.net/filename are not obfuscated. By not obfuscating the hostname, users can still access these resources.

   2. (Optional) Set additional Web options for individual resources (such as enabling Web rewriting to match IP addresses to hostnames) using settings in the Users > Resource Policies> Web > Options page of the admin console.

> **NOTE:** Certain Web rewriting features (such as passthrough proxy and SSO to NTLM resources) require additional configuration. For more information, see the appropriate configuration instructions.

Related Documentation

- Writing a Web Access Resource Policy
- Defining a Single Sign-On Autopolicy
- Defining a Caching Autopolicy
- Writing a Java Access Control Resource Policy
- Defining a Rewriting Autopolicy
- Defining a Web Compression Autopolicy
- Writing a Web Proxy Resource Policy
- Automatically Launching JSAM
- Writing An HTTP 1.1 Protocol Resource Policy
- Defining Resource Policies: General Options

Remote SSO Overview

The Remote Single Sign-On (SSO) feature enables you to specify the URL sign-in page of an application to which you want the system to post a user's credentials, minimizing the need for users to re-enter their credentials when accessing multiple back-end applications. You may also specify additional forms values and custom headers (including cookies) to post to an application's sign-in form.

Remote SSO configuration consists of specifying Web resource policies:

- Form POST policy—This type of Remote SSO policy specifies the sign-in page URL of an application to which you want to post system data and the data to post. This data can include the user's primary or secondary username and password as well as system data stored by system variables. You can also specify whether or not users can modify this information.

- Headers/Cookies policy—This type of Remote SSO policy specifies resources, such as customized applications, to which you can send custom headers and cookies.

If a user's system credentials differ from those required by the back-end application, the user can alternatively access the application:

- By signing in manually—The user can quickly access the back-end application by entering his credentials manually into the application's sign-in page. The user may also permanently store his credentials and other required information in the system through the Preferences page as described below, but is not required to enter information in this page.

- Specifying the required credentials on Connect Secure —The user must provide the system with his correct application credentials by setting them through the Preferences page. Once set, the user must sign out and sign back in to save his credentials. Then, the next time the user clicks the Remote SSO bookmark to sign in to the application, the system sends the updated credentials.

ℹ️ NOTE: Use the Remote SSO feature to pass data to applications with static POST actions in their HTML forms. It is not practical to use Remote SSO with applications that employ frequently changing URL POST actions, time-based expirations, or POST actions that are generated at the time the form is generated.

Related Documentation

- Using Custom Expressions in Rule Configuration

- Multiple Sign-In Credentials Execution

- Writing a Remote SSO Form POST Resource Policy

- Writing a Remote SSO Headers/Cookies Resource Policy

Passthrough Proxy Overview

The passthrough proxy feature enables you to specify Web applications for which the system performs minimal intermediation. Unlike traditional reverse proxy functionality, which also rewrites only selective parts of a server response but requires network changes as well as complex configuration, this feature only requires that you specify application servers and the way in which the system receives client requests to those application servers:

- Via a Connect Secure port—When specifying an application for the passthrough proxy to intermediate, you specify a port on which the system listens for client requests to the application server. When the system receives a client request for the application server, it forwards the request to the specified application server port. When you choose this option, you must open traffic to the specified system port on your corporate firewall.

- Via virtual hostname—When specifying an application for the passthrough proxy to intermediate, you specify an alias for the application server hostname. You need to add an entry for this alias in your external DNS server that resolves to the system. When the system receives a client request for the alias, it forwards the request to the port you specify for the application server.

  This option is useful if your company has restrictive policies about opening firewall ports to either internal servers or servers in the DMZ. When using this option, we recommend that each hostname alias contains the same domain substring as your hostname and that you upload a wild card server certificate to the system in the format: *.domain.com.

  For example, if your system is iveserver.yourcompany.com, then a hostname alias should be in the format appserver.yourcompany.com and the wild card certificate format would be *.yourcompany.com. If you do not use a wild card certificate, then a client's browser issues a certificate name check warning when a user browses to an application server, because the application server hostname alias does not match the certificate domain name. This behavior does not prevent a user from accessing the application server, however.

  ℹ️ NOTE: When you configure passthrough proxy to work in virtual hostname mode, users must use the hostname that you specify through the System > Network > Overview page of the admin console when signing into the device. They cannot access the use passthrough proxy feature if they sign into the device using its IP address.

Just as with the Content Intermediation Engine, the passthrough proxy option offers increased security relative to the Secure Application Manager, because when enabled for an application, the system allows the client to send only Layer 7 traffic directed to fixed application ports to the

enterprise network. Use this option to enable the system to support applications with components that are incompatible with the Content Intermediation Engine, such as Java applets in Oracle e-business suite applications or applets that run in an unsupported Java Virtual Machine (JVM).

Note the following:

- Passthrough proxy URLs must be hostnames. Paths of hostnames are not supported.

- Pulse Secure strongly recommends that you not mix passthrough proxy Port mode and passthrough proxy Host mode.

- The passthrough proxy option works only for applications that listen on fixed ports and where the client does not make direct socket connections.

- To use passthrough proxy with Oracle E-Business applications, you must install a real certificate on the system and you must configure Oracle Forms to use the Forms Listener Servlet mode.

- The following advanced features of the framed toolbar are not available in passthrough proxy: bookmark current page, display the original URL, display the favorite bookmarks.

Related Documentation

- Creating a Passthrough Proxy Resource Policy

Creating a Custom Web Application Resource Profile

A custom Web application resource profile is a resource profile that controls access to a Web application, Web server, or HTML page.

To create a custom Web application resource profile:

1. In the admin console, select **Users > Resource Profiles > Web**.

2. Click **New Profile**.

3. From the Type list, choose **Custom**.

4. Enter a unique name and optionally a description for the resource profile.

5. In the Base URL field, enter the URL of the Web application or page for which you want to control access using the format: [protocol://]host[:port][/path]. (The system uses the specified URL to define the default bookmark for the resource profile.)

6. In the Autopolicy: Web Access Control section, create a policy that allows or denies users access to the resource specified in the Base URL field. (By default, the system automatically creates a policy for you that enables access to the Web resource and all of its sub-directories.)

7. (Optional) Click **Show ALL autopolicy types** to create additional autopolicies that fine-tune access to the resource. Then, create the autopolicies using instructions in the following sections:

8. Click **Save and Continue**.

9. In the Roles tab, select the roles to which the resource profile applies and click **Add**.

   The selected roles inherit the autopolicies and bookmarks created by the resource profile. If it is not already enabled, the system also automatically enables the Web option in the Users

> User Roles > Select Role > General > Overview page of the admin console for all of the roles you select.

10.   Click **Save Changes**.

11.   (Optional) In the Bookmarks tab, modify the default bookmark created by the system and/or create new ones. (By default, the system creates a bookmark to the base URL defined in the Base URL field and displays it to all users assigned to the role specified in the Roles tab.)

Defining Base URLs

When creating a Web resource profile, you must use the following format when defining base URLs:

[protocol://]host[:port][/path]

Within this format, the components are:

- Protocol (required)—Possible values: http:// and https://. Note that you cannot use special characters within the protocol.

- Host (required)—Possible values:

  - DNS Hostname—For example: www.pulsesecure.net

  - IP address—You must enter the IP address in the format: a.b.c.d. For example: 10.11.149.2. You cannot use special characters in the IP address.

- Ports (optional)—You must use the delimiter ":" when specifying a port. For example: 10.11.149.2/255.255.255.0:*

- Path (optional)—When specifying a path for a base URL, the system does not allow special characters. If you specify a path, you must use the "/" delimiter. For example, http://www.pulsesecure.net/sales.

Defining Web Resources

When creating a Web resource profile, you must use the following format when defining resources for autopolicies:

[protocol://]host[:ports][/path]

Within this format, the four components are:

- Protocol (required)—possible values: http:// and https://. Note that you cannot use special characters within the protocol.

- Host (required)—possible values:

  - DNS Hostname—For example: www.pulsesecure.net

    Table 81 lists the special characters allowed in the hostname.

Table81: DNS Hostname Special Characters

*	Matches ALL characters.
%	Matches any character except dot (.)

?	Matches exactly one character

- IP address/Netmask—You must enter the IP address in the format: a.b.c.d

  You may use one of two formats for the netmask:

  - Prefix: High order bits

  - IP: a.b.c.d

  For example: 10.11.149.2/24 or 10.11.149.2/255.255.255.0

  You cannot use special characters in the IP address or netmask.

- Ports (optional)—You must use the delimiter ":" when specifying a port. For example: 10.11.149.2/255.255.255.0:* Table 82 lists the possible port values.

Table82: Possible Port Values

*	Matches ALL ports; you cannot use any other special characters
port[,port]*	A comma-delimited list of single ports. Valid port numbers are [1-65535].
[port1]-[port2]	A range of ports, from port1 to port2, inclusive.

NOTE: You can mix port lists and port ranges, such as: 80,443,8080-8090

If the port is missing, then the default port 80 is assigned for http, 443 for https.

- Path (optional)—When specifying a path for a Web access control autopolicy, you may use a * character, meaning ALL paths match. (The system does not support any other special characters.) If you specify a path, you must use the "/" delimiter. For example:

  - http://www.pulsesecure.net/sales

  - http://www.pulsesecure.net:80/*

  - https://www.pulsesecure.net:443/intranet/*

Related Documentation

- Resource Profiles

- About Hosted Java Applet Templates

- About Citrix Templates

- Creating Resource Profiles Using the Lotus iNotes Template

- Creating Resource Profiles Using the Microsoft OWA Template

- Creating Resource Profiles Using the Microsoft Sharepoint Template

- Defining a Web Access Control Autopolicy

- Defining a Single Sign-On Autopolicy

- Defining a Rewriting Autopolicy

- Defining a Web Compression Autopolicy

- Defining a Java Access Control Autopolicy

- Defining a Caching Autopolicy

Defining a Web Access Control Autopolicy

Web access policies control which Web resources users can access in order to connect to the Internet, intranet, or extranet. When defining a custom Web resource profile, you must enable a corresponding Web access control autopolicy that enables access to the profile's primary resource. The system simplifies the process for you by automatically creating an autopolicy that allows access to the Web resource and all of its sub-directories.

If necessary, you may choose to modify this default autopolicy or create supplementary Web access control autopolicies that control access to additional resources. For instance, your IT department may use one server to store Web pages for your company intranet (http://intranetserver.com) and another server to store the images that the Web pages reference (http://imagesserver.com). In this case, you can create two Web access control autopolicies that enable access to both servers so that your users can access both your Web pages and the corresponding images.

To create a new Web access control autopolicy:

1. Create a custom Web application resource profile.

2. If available, click the **Show ALL autopolicy types** button to display the autopolicy configuration options.

3. If it is not already enabled, select the **Autopolicy: Web Access Control** check box.

4. In the Resource field, specify the Web server or HTML page to which you want to control access using the format: [protocol://]host[:ports][/path].

5. From the Action list, choose **Allow** to enable access to the specified resource or Deny to block access to the specified resource.

6. Click **Add**.

7. Click **Save Changes**.

Related Documentation

- Creating a Custom Web Application Resource Profile

- Configuring General Role Options

Defining a Single Sign-On Autopolicy

Single sign-on policies enable you to automatically pass user credentials to the Web application specified in your policy. Single sign-on autopolicies also intermediate the data that you pass.

To create a single sign-on (SSO) autopolicy:

1. Create a Web resource profile.

2. If available, click the **Show ALL autopolicy types** button to display the autopolicy configuration options.

3. Select the **Autopolicy: Single Sign-On** check box.

4. Select a single sign-on method and configure the corresponding SSO options:

NOTE: SSO options require you to select credentials. If you have not already done so, define the credentials using the Resource Policies > Web > General page prior to defining your SSO autopolicy.

- **Disable SSO**—Disables single sign-on.

- **Basic Auth**—Enables the system to intermediate the challenge/response sequence during basic authentication and use the credentials it collects to sign into a protected resource within the same Intranet zone. This option does not apply to Citrix resource profiles.

- **NTLM**—Enables the system to intermediate the challenge/response sequence during NTLM authentication and use the credentials it collects to sign into a protected resource within the same Intranet zone. This option does not apply to Citrix resource profiles.

  NOTE: Web rewriting and file browsing both support NTLM v1 and NTLM v2.

- **Kerberos**—Enables the system to intermediate the challenge/response sequence during Kerberos authentication and use the credentials it collects to sign into a protected resource within the same Intranet zone.

- **Constrained Delegation**—Enables authentication of users by Kerberos after their identity has been verified using a non-Kerberos authentication method. For example, suppose a user authenticates with RADIUS and enters their passcode (typically PIN and tokencode). When accessing a service, the user may be challenged again because the PIN is not recognized. With constrained delegation, the administrator sets up passwords for constrained delegation users. The users do not need to know this password. When accessing the same HTTP service, the system now fetches the ticket on behalf of the user without challenging the user.

- **Remote SSO**—Enables the system to post the data that you specify (including usernames, passwords, and system data stored by variables) to Web applications. This option also enables you specify custom headers and cookies to post to Web applications.

5. Click **Save Changes**.

Specifying Basic Authentication, NTLM or Kerberos SSO Autopolicy Options

To configure basic authentication, NTLM or Kerberos SSO autopolicy options:

1. Create an SSO autopolicy and choose Basic Auth, NTLM or Kerberos.

2. In the Resource field, specify the resources to which this policy applies.

   When entering a resource in this field, note that:

   - If you want to automatically post values to a specific URL when an end-user clicks on a system bookmark, the resource that you enter here must exactly match the URL that you specify in the Base URL field of the resource profile.

   - If you want to automatically submit user credentials to other web sites within the same Intranet zone, the hostname that you enter here must end in the DNS suffix configured in the System > Network > Overview page of the admin console.

3. Select the credentials to use. If this pull-down menu is blank, no credentials are defined in the SSO General tab.

4. (NTLM only) Select the **Fallback to NTLM V1** option to fallback to NTLM V1 if NTLM V2 fails. If you do not select this option, the system falls back only to NTLM V2. An intermediation page appears if SSO fails.

5. (Kerberos only) Select the **Fallback to NTLM V2 only** option to fallback only to NTLM V2 if kerberos fails. If you do not select this option, a Kerberos intermediation page appears if Kerberos SSO fails.

6. (Constrained delegation only) Select the **Fallback to Kerberos** option fallback to Kerberos if constrained delegation fails. If you do not select this option, an error page appears if SSO fails.

Specifying Remote SSO Autopolicy Options

To configure remote SSO autopolicy options:

1. Create an SSO autopolicy through a custom Web resource profile and choose **Remote SSO**.

2. If you want to perform a form POST when a user makes a request to the resource specified in the Resource field, select the POST the following data check box. Then:

    1. In the Resource field, specify the application's sign-in page, such as: http://my.domain.com/public/login.cgi. Widlcard characters are not supported in this field.

       If you want to automatically post values to a specific URL when an end user clicks on a system bookmark, the resource that you enter here must exactly match the URL that you specify in the Base URL or Web Interface (NFuse) URL field of the resource profile.

    2. In the Post URL field, specify the absolute URL where the application posts the user's credentials, such as: http://yourcompany.com/login.cgi. You can determine the appropriate URL using a TCP dump or by viewing the application's sign-in page source and searching for the POST parameter in the FORM tag.

    3. Optionally specify the user data you want to post and user modification permissions.

       To specify user data to post, enter data in the following fields and click **Add**:

       - **Name**—The name to identify the data of the Value field. (The back-end application should expect this name.)

       - **Value**—The value to post to the form for the specified Name. You can enter static data, a system variable, or system session variables containing username and password values.

       - **User modifiable? setting**—Set to **Not modifiable** if you do not want the user to be able to change the information in the Value field. Set to **User CAN change value** if you want the user to have the option of specifying data for a back-end application. Set to **User MUST change value** if users must enter additional data in order to access a back-end application. If you choose either of the latter settings, a field for data entry appears on the user's Advanced Preferences page. This field is labeled using the data you enter in the User label field. If you enter a value in the Value field, this data appears in the field but is editable.

    4. Select the **Deny direct login for this resource** check box if you do not want to allow users to manually enter their credentials in a sign-in page. (Users may see a sign-in page if the form POST fails.)

    5. Select the **Allow multiple POSTs to this resource** check box if you want to send POST and cookie values to the resource multiple times if required. If you do not select this

option, the system does not attempt single sign-on when a user requests the same resource more than once during the same session.

6. If you want to post header data to the specified URL when a user makes a request to a resource specified in the Resource field, select the **Send the following data as request headers** check box. Then:

   1. In the Resource section, specify the resources to which this policy applies.

   2. Optionally specify the header data to post by entering data in the following fields and clicking **Add**:

      - **Header name**—The text to send as header data.

      - **Value**—The value for the specified header.

   3. Click **Save Changes**.

Related Documentation

- Creating a Custom Web Application Resource Profile

- Remote SSO Overview

Defining a Caching Autopolicy

Caching policies control which Web content the system caches on a user's machine.

To create a Web caching autopolicy:

1. Create a custom Web application resource profile.

2. If available, click the **Show ALL autopolicy types** to display the autopolicy configuration options.

3. Select the **Autopolicy: Caching** check box.

4. In the Resource field, specify the resources to which this policy applies.

5. In the Action field, select one of the following options:

   - **Smart**—Select this option to send a cache-control:no-store header or a cache-control:no-cache header based on the user's Web browser and content type.

     When you select this option, the system makes media files and zip files work properly by removing their origin server's cache-control headers. For example, the following logic searches for "msie" or "windows-media-player" in user-agent headers in order to remove cache or cache-control:no-store response headers and make the files cacheable:

     ```
     (if content type has "audio/x-pn-realaudio" OR
      if content type begins with "video/" OR
      if content type begins with "audio/" OR
      if content type is "application/octet-stream" and the file extension
     begins with "rm" or "ram"
     )
     ```

     If the system finds "msie" or "windows-media-player" in the user-agent header and any of the following apply:

   - Request is for Flash, .xls, .pps, .ppt files

   - Content-type is application/, text/rtf, text/xml, model/

- Origin server sends a content-disposition header

then the system sends the cache-control:no-store header and removes the origin server's cache-control header.

In all other cases, the system adds the pragma:no-cache or cache-control:no-store response headers.

> **NOTE:** Citrix .ica and QuickPlace files get some special treatment. Citrix .ica files get cache-control:private only when smart caching is enabled. QuickPlace files that do not match a specified rule files (which takes precedence) get CCNS and cache-control:private.
>
> Also note that if you select this option, enable GZIP compression, and try to access a text file attachment using Domino Web Access 6.5 through Internet Explorer, you cannot open the attachment. To enable text attachments, you must either install the Internet Explorer 323308 patch or enable the No Store option.

- **No-Store**—Select this option to deliver attachments to Internet Explorer without saving them to the disk. (The browser temporarily writes files to the disk, but immediately removes them once it has opened the file in the browser.) When you select this option, the system removes the origin server's cache-control header and adds a cache-control:no-store response header if the user-agent string sent by the browser contains "msie" or "windows-media-player."

  This option might slow browsing by causing repeated content fetches, which can cause performance issues on very slow connections.

- **No-Cache**—Select this option to prevent the user's browser from caching files to the disk. When you select this option, the system adds the standard HTTP pragma:no-cache header and cache-control:no-cache (CCNC) header (HTTP 1.1) to response files. Also, the system does not forward the origin server's caching headers, such as age, date, etag, last-modified, expires.

  When no-cache headers are present on certain types of attachments (PDF, PPT, streaming files), Internet Explorer does not properly render the documents because the rendering process requires the browser to temporarily writes these files to cache.

- **Unchanged**—Select this option to forward the origin server's caching headers as is.

  When using Citrix published applications through the Web interface, the Web interface server may send a Cache-Control:no-cache in the response header of the .ica file. Because the caching header is not removed when using the Unchanged setting, .ica files are not downloaded to the client PC. To resolve this, use the Smart caching option.

6. Click **Add**.

7. Click **Save Changes**.

Related Documentation

- Creating a Custom Web Application Resource Profile

Defining a Java Access Control Autopolicy

A Java access control autopolicy defines the list of servers and ports to which Java applets can connect. This autopolicy also specifies which resources the system signs using the code-signing certificate that you upload.

When you enable Java access control using this autopolicy, the system automatically enables the Allow Java applets option on the Users > User Roles > Select Role > Web > Options page of the admin console.

To create a Java access control autopolicy:

1. Create a custom Web application resource profile.

2. Click **Show ALL autopolicy types**.

3. Select the **Autopolicy: Java Access Control** check box.

4. In the Resource field, specify the server resources to which this policy applies using the format: host:[ports]. (By default, the system populates this field with the server specified in your resource profile's base URL.)

5. Select one of the following options from the Action list:

   - **Allow socket access**—To enable Java applets to connect to the servers (and optionally ports) in the Resource list.

   - **Deny socket access**—To prevent Java applets from connecting to the servers (and optionally ports) in the Resource list.

6. Click **Add**.

7. Select the **Sign applets with code-signing certificate** check box to resign the specified resources using the certificate uploaded through the System > Configuration > Certificates > Code-signing Certificates page of the admin console. (The system uses the imported certificate to sign the server resources that you specify in the Resources field.)

8. Click **Save Changes**.

Defining a Server to Which Java Applets Can Connect

When defining servers to which Java applets can connect, you must use the following format:

host[:ports]

Within this format, the two components are:

- Host (required)—Possible values:

  - DNS Hostname—For example: www. www.pulsesecure.net

    You may use the following special characters allowed in the hostname:

*	Matches ALL characters.
%	Matches any character except dot (.)
?	Matches exactly one character

  - IP address/Netmask—You must enter the IP address in the format: a.b.c.d.

You may use one of two formats for the netmask:

Prefix: High order bits

IP: a.b.c.d

For example: 10.11.149.2/24 or 10.11.149.2/255.255.255.0 You cannot use special characters in the IP address or netmask.

- Ports—You must use the delimiter ":" when specifying a port. For example: 10.11.149.2/255.255.255.0:* Table 83 lists the possible port values.

Table83: Possible Port Values

*	Matches ALL ports; you cannot use any other special characters
port[,port]*	A comma-delimited list of single ports. Valid port numbers are [1-65535].
[port1]-[port2]	A range of ports, from port1 to port2, inclusive.

> **NOTE:** You can mix port lists and port ranges, such as: 80,443,8080-8090.

Related Documentation

- About Hosted Java Applet Templates
- Creating a Custom Web Application Resource Profile

Defining a Rewriting Autopolicy

By default, the system intermediates all user requests to Web hosts—unless you have configured it to serve requests to certain hosts using a different mechanism, such as the Secure Application Manager. Rewriting autopolicies enable you to "fine-tune" the default options by changing which mechanisms to rewrite Web data and defining resources that you want to minimally rewrite or not rewrite at all.

To create a rewriting autopolicy:

1. Create a custom Web application resource profile.

2. Click **Show ALL autopolicy types**.

3. Select the **Autopolicy: Rewriting Options** check box.

4. Select one of the following options:

   - **Passthrough Proxy**—Select this option to specify Web applications for which the Content Intermediation Engine performs minimal intermediation..

   - **No rewriting (use WSAM)**—Select this option to intermediate content using WSAM instead of the Content Intermediation Engine. Then, specify the application server for which you want to intermediate content. (At minimum, you need to click **Add** in order to intermediate content to and from the server that the system extracts from the Web access control policy.) .

   - **No rewriting (use JSAM)**—Select this option to intermediate content using JSAM instead of the Content Intermediation Engine. Then, specify the application server for which you want to intermediate content.(At minimum, you need to click **Add** in order to intermediate

content to and from the server that the system extracts from the Web access control policy.)

- **No rewriting**—Select this option to automatically create a selective rewriting policy for the autopolicy's URL, thereby configuring the system to not intermediate any content to and from the resource. For example, you may choose this option if you do not want the system to intermediate traffic from web sites that reside outside of the corporate network, such as yahoo.com. If you select this option, you do not have to configure any additional rewriting settings.

Specifying Passthrough Proxy Autopolicy Options

To configure passthrough proxy autopolicy options:

1. Create an rewriting autopolicy and select **Passthrough Proxy**.

2. Choose the way in which you want to enable the passthrough proxy feature:

   - **Use virtual hostname**—If you choose this option, specify a hostname alias for the application server. When the system receives a client request for the application server hostname alias, it forwards the request to the specified application server port in the Base URL field.

   - **Use Secure Gateway port**—If you choose this option, specify a unique port in the range 11000-11099. The system listens for client requests to the application server on the specified port and forwards any requests to the application server port specified in the Base URL field.

   The corresponding URL for the resource profile must specify the application server hostname and the port used to access the application internally. You cannot enter a path for the base URL.

   In order to make Sharepoint work successfully through the system, you must select the Override automatic cookie handling check box in Internet Explorer under Tools Internet options > Privacy > Advanced Privacy Settings if the following conditions true:

   - You select the **Use virtual hostname** option during Pass Through Proxy configuration.

   - The virtual hostname that you specify in your Sharepoint configuration is different from the hostname that you configure through the system setup (that is, if the domains are different).

   - You enable persistent cookies through the Users > User Roles > Select Role > General > Session Options page of the admin console.

3. Select the **Rewrite XML** check box if you want to rewrite URLs contained within XML content. If this option is disabled, the system passes the XML content "as is" to the server.

4. Select the **Rewrite external links** check box if you want to rewrite all the URLs presented to the proxy. If this option is disabled, the system rewrites only those URLs where the hostname is configured as part of the passthrough proxy policy.

5. Select the **Block cookies from being sent to the browser** check box if you want to block cookies destined for the client's browser. The system stores the cookies locally and sends them to applications whenever they are requested.

6. Select the **Host-Header forwarding** check box if you want to pass the hostname as part of the host header instead of the actual host identifier.

The Host-Header forwarding option is only valid in passthrough proxy Virtual hostname mode.

7. Click **Save Changes**.

8. If you select:

- Use virtual hostname, you must also:

  - Add an entry for each application server hostname alias in your external DNS that resolves to the system.

  - Upload a wildcard server certificate to the system (recommended).

  - Define the system name and hostname in the Network Identity section of the System > Network > Internal Port tab.

- To use the system port, you must also open traffic to port you specified for the application server in your corporate firewall.

  If your application listens on multiple ports, configure each application port as a separate passthrough proxy entry with a separate port. If you intend to access the server using different hostnames or IP addresses, configure each of those options separately; in this case, you can use the same port.

Specifying WSAM Rewriting Autopolicy Options

To configure WSAM rewriting autopolicy options:

1. Create an rewriting autopolicy and select **No rewriting (use WSAM)**.

2. In the Destination field, specify resources for which WSAM secures client/server traffic between the client and the system. By default, the system extracts the correct server from the Web access control policy. You may choose to use this server as-is, modify it, and/or add new servers to the list.

   When specifying a server, specify the hostname (the wild cards '*' or '?' are accepted) or an IP/netmask pair. Specify multiple ports for a host as separate entries.

3. Click **Add**.

4. Click **Save Changes**.

When you intermediate through WSAM using this autopolicy, the system automatically enables the Secure Application Manager option on the Users > User Roles > Select Role > General > Overview page of the admin console.

Specifying JSAM Rewriting Autopolicy Options

To configure JSAM rewriting autopolicy options:

1. Create an rewriting autopolicy and select **No rewriting (use JSAM)**.

2. In the Server Name field, enter the DNS name of the application server or the server IP address.

3. In the Server Port field, enter the port on which the remote server listens for client connections.

   For example, to forward Telnet traffic from a remote machine, specify port 23 for both the client port (on which JSAM listens) and the server port (on which the Telnet server listens).

ⓘ NOTE: To enable drive mapping to this resource, enter 139 as the server port.

4. In the Client Loopback IP field, provide a static loopback address. If you do not provide a static IP loopback address, the system assigns an IP loopback address dynamically.

5. In the Client Port field, enter the port on which JSAM should listen for client application connections.

   Typically, the local port value is the same value as the server port; the local port value usually only differs for Linux or Macintosh users who want to add applications for port forwarding that use ports under 1024.

   ⓘ NOTE: To enable drive mapping to this resource, enter 139 as the server port.

   You may configure more than one application on a single port, such as app1.mycompany.com, app2.mycompany.com, app3.mycompany.com. Either you assign a static loopback address or the system assigns a dynamic loopback address (127.0.1.10, 127.0.1.11, 127.0.1.12) to each application. JSAM then listens on these multiple loopback addresses on the specified port. For example, when there is traffic on 127.0.1.12 on the specified port, the system forwards the traffic to the app3.mycompany.com destination host.

6. Select **Launch JSAM** to automatically start JSAM when the system encounters the Base URL.

7. Click **Add**.

8. Click **Save Application** or **Save + New**.

Related Documentation

- Web Rewriting

- Passthrough Proxy Overview

- Task Summary: Configuring WSAM

- Task Summary: Configuring JSAM

Defining a Web Compression Autopolicy

Web compression autopolicies specify which types of Web data the system should and should not compress. For example, since javascript does not work when compressed, you might use this feature to specify that the system should not compress javascript data going to and from an e-mail server by entering the following resource: http://owa. pulsesecure.net.net/*.js.

ⓘ NOTE: In order to properly compress data, you must enable compression at the system level as well as creating compression autopolicies. To enable compression, use settings in the Maintenance > System > Options page of the admin console.

To create a Web compression autopolicy:

1. Create a custom Web application resource profile.

2. If available, click the **Show ALL autopolicy types** button to display the autopolicy configuration options.

3. Select the **Autopolicy: Web compression** check box.

4. In the Resource field, specify the resources to which this policy applies.

5. Select one of the following options from the Action list:

- **Compress**—Compress the supported content types from the specified resource.

- **Do not compress**—Do not compress the supported content types from the specified resource.

6. Click **Add**.

7. Click **Save Changes**.

Related Documentation

- About Compression

- Creating a Custom Web Application Resource Profile

Defining Web Resource Profile Bookmarks

When you create a Web resource profile, the system automatically creates a bookmark that links to the primary URL or domain that you specified in the resource profile. The system enables you to modify this bookmark as well as create additional bookmarks within the same domain.

For example, you may create a resource profile that controls access to your company intranet. Within the profile, you may specify:

- Resource profile name: Your Intranet

- Primary resource: http://intranet.com

- Web access control autopolicy: Allow access to http://intranet.com:80/*

- Roles: Sales, Engineering

When you create this policy, the system automatically creates a bookmark called "Your Intranet" enabling access to http://intranet.com and displays the bookmark to members of the Sales and Engineering roles.

You may then choose to create the following additional bookmarks to associate with the resource profile:

- "Sales Intranet" bookmark: Creates a link to the http://intranet.com/sales page and displays the link to members of the Sales role.

- "Engineering Intranet" bookmark: Creates a link to the http://intranet.com/engineering page and displays the link to members of the Engineering role.

When configuring bookmarks, note that:

- You can only assign bookmarks to roles that you have already associated with the resource profile—not all of the roles defined on the system. To change the list of roles associated with the resource profile, use settings in its Roles tab.

- Bookmarks simply control which links to display to users—not which resources the users can access. For instance, in the example used above, a member of the Sales role would not see a link to the Engineering Intranet page, but he could access it by entering http://intranet.com/engineering his Web browser's address bar.

- You cannot create bookmarks that link to additional URLs and domains defined through Web access control autopolicies.

You can use two different methods to create Web bookmarks:

- Create bookmarks through existing resource profiles (recommended)—When you select this method, the system automatically populates the bookmark with key parameters (such as the Web interface (NFuse) URL) using settings from the resource profile. Additionally, while you are creating the associated resource profile, the system guides you through the process of creating any required policies to enable access to the bookmark.

- Create standard bookmarks—When you select this option, you must manually enter all bookmark parameters during configuration. Additionally, you must enable access to the Web feature and create resource policies that enable access to the web sites defined in the bookmark.

Creating Bookmarks Through Existing Resource Profiles

To configure Web resource profile bookmarks:

1. If you want to create a resource profile bookmark through the standard resource profiles page:

    1. In the admin console, select **Users > Resource Profiles > Web > *Resource Profile Name* > Bookmarks**.

    2. Click the appropriate link in the Bookmark column if you want to modify an existing bookmark. Or, click **New Bookmark** to create an additional bookmark.

    Alternatively, if you want to create a resource profile bookmark through the user roles page:

    1. In the admin console, select **Users > User Roles > *Role Name* > Web > Bookmarks**.

    2. Click **New Bookmark**.

    3. From the Type list, choose **Pick a Web Resource Profile**. (The system does not display this option if you have not already created a Web resource profile.)

    4. Select an existing resource profile.

    5. Click **OK**. (If you have not already associated the selected role with the resource profile, the system automatically makes the association for you. The system also enables any access control policies for the role that are required by the resource profile.)

    6. If this role is not already associated with the selected resource profile, the system displays an informational message. If you see this message, click **Save Changes** to add this role to the resource profile's list of roles and to update the profile's autopolicies as required. Then, repeat the previous steps to create the bookmark.

        When you create a resource profile bookmark through the user roles page (instead of the standard resource profiles page), the system only associates the generated bookmark with the selected role. The system does not assign the bookmark to all of the roles associated with the selected resource profile.

    7. Optionally change the name and description of the bookmark. (By default, the system populates names the bookmark using the resource profile name.)

    8. In the URL field, add a suffix to the URL if you want to create links to sub-sections of the domain defined in the primary resource profile.

Make sure to enter a unique URL in this field. If you create two bookmarks with the same URL, the system deletes one of the bookmarks from the end-user view. You will still be able to see both bookmarks, however, in the administrator console.

9. Under Options, select the **Bookmark opens in new window** check box if want to enable the system to automatically open the Web resource in a new browser window. Next, select:

- **Do not display browser address bar**—Select this option to remove the address bar from the browser window. This feature forces all Web traffic through the system by precluding users in the specified role from typing a new URL in the address bar, which circumvents the system.

- **Do not display browser toolbar**—Select this option to remove the menu and toolbar from the browser. This feature removes all menus, browsing buttons, and bookmarks from the browser window so that the user browses only through the system.

10. If you are configuring the bookmark through the resource profile pages, under Roles, specify the roles to which you want to display the bookmark:

- **ALL selected roles**—Select this option to display the bookmark to all of the roles associated with the resource profile.

- **Subset of selected roles**—Select this option to display the bookmark to a subset of the roles associated with the resource profile. Then select roles from the ALL Selected Roles list and click **Add** to move them to the Subset of selected roles list.

11. Click **Save Changes**.

Creating Standard Web Bookmarks

Information in this section is provided for backwards compatibility. We recommend that you configure access to Web URLs and servers through resource profiles instead, since they provide a simpler, more unified configuration method.

Use the Bookmarks tab to create bookmarks that appear on the welcome page for users mapped to this role. You can create two types of bookmarks through this page:

- Web URL bookmarks—These bookmarks link the user to Web URLs on the World Wide Web or on your corporate Intranet. When you create Web bookmarks, you can insert the user's username in the URL path to provide single sign-on access to back-end Web applications. For Web bookmark configuration instructions, see the instructions that follow.

- Java applet bookmarks—These bookmarks link the user to a Java applets that you upload through the Users > Resource Profiles > Web > Hosted Java Applets page of the admin console.

When you create either of these bookmark types, the corresponding links appear on the welcome page for users mapped to this role.

To create a bookmark to a Web resource:

1. In the admin console, choose **Users > User Roles > Role > Web > Bookmarks**.

2. Click **New Bookmark**.

3. Select **Standard**.

4. Enter a name and description for the bookmark (optional). This information displays on the home page instead of the URL.

5. Enter the URL to bookmark. If you want to insert the user's username, enter <username> at the appropriate place in the URL.

   Make sure to enter a unique URL in this field. If you create two bookmarks with the same URL, the system deletes one of the bookmarks from the end-user view. You will still be able to see both bookmarks, however, in the administrator console.

6. Under Auto-allow, click **Auto-allow Bookmark** to automatically create a corresponding Web access resource policy. Note that this functionality applies only to role bookmarks and not bookmarks created by users. Next, select:

   - **Only this URL** to allow users to access only the URL.

   - **Everything under this URL** to allow the user to access any path under the URL.

     You may not see the Auto-allow option if you are using a new installation or if an administrator hides the option.

7. Under Display options, click **Open bookmark in a new window** to automatically open the Web resource in a new browser window. Note that this functionality applies only to role bookmarks and not bookmarks created by users. Next, select:

   - **Do not display the URL address bar** if you want to remove the address bar from the browser window. This feature forces all Web traffic through the system by precluding users in the specified role from typing a new URL in the address bar, which circumvents the system.

   - **Do not display the menu and the toolbar** to remove the menu and toolbar from the browser. This feature removes all menus, browsing buttons, and bookmarks from the browser window so that the user browses only through the system.

8. Click **Save Changes** or **Save + New** to add another.

Related Documentation

- Defining Resource Profile Bookmarks

- Using Custom Expressions in Rule Configuration

Specifying Web Browsing Options

the system enables you to configure a wide-variety of Web browsing options for a user role.

To configure the Web browsing options for a role:

1. Select **Users > User Roles > *RoleName* > Web > Options**.

2. Complete the configuration as described in Table 84.

3. Click **Save Changes**.

Table84: Web Browsing Options for a Role

Settings	Guidelines
User can type URLs in the browse bar	(Default) Select this option to enable users to enter URLs on the welcome page and browse to Internet sites.

User can add bookmarks	(Default) Select this option to enable users to create personal web bookmarks on the system welcome page.
Mask hostnames while browsing	Select this option to obscure the target resources in the URLs to which the users browse. When you select this option, the system masks IP addresses and hostnames in the user's:
	• Web browser address bar (when the user navigates to a page)
	• Web browser status bar (when a user hovers over a hyperlink)
	• HTML source files (when the user chooses to View Source)
	The hostname encoding feature (also called hostname obfuscation or URL obfuscation) prevents casual observers from noting the URL of an internal resource by obscuring the target server within the URL without masking the full path name, target file, or port number. For example, if a user navigates to www.msn.com without selective rewriting or hostname encoding enabled, the system displays an unobscured URL in his Web browser's address bar:
	http://www.msn.com/
	If you then enable selective rewriting, the system might display the following URL:
	https://mycompanyserver.com/,DanaInfo=www.msn.com,SSO=U+
	If you then enable hostname encoding, and the same user navigates to the same site, he sees a URL in which the hostname (www.msn.com) is obscured:
	https://i5.asglab.pulsesecure.net/,DanaInfo=.awxyCqxtGkxw,SSO=U+
	Hostname encoding uses a lightweight reversible algorithm so that users can bookmark encoded URLs. (The system can translate the encoded URL and resolve it back to the original URL.) For compatibility, previously created bookmarks to unmasked URLs continue to work when hostname encoding is enabled.
	NOTE:
	• If you enable selective rewriting and hostname encoding, the system only obscures the hostnames and IP addresses of those servers that you have chosen to rewrite using the selective rewrite feature.
	• Links not rewritten by the system are not obscured. For example, the rewriter does not intermediate ftp, rtsp, mms and mailto links and therefore the hostnames in these links are not masked. This is required to pass security audits.
	• If you enable the framed toolbar and hostname encoding, the system does not obscure hostnames that the user enters in the framed toolbar's browse field.
	• The system does not obscure hostnames and IP addresses in log entries, including hostname encoding log entries.
Advanced options	
Allow Java applets	(Default) Select this option to enable users to browse to Web pages containing client-side Java applets. The system appears to the application server as a browser over SSL. The system transparently handles any HTTP requests and TCP connections initiated by a Java applet and handles signed Java applets.
	If you enable this feature, users can launch Java applets and run applications that are implemented as client-side Java applets, such as the Virtual Computing (VNC) Java client, Citrix NFuse Java client, WRQ Reflections Web client, and Lotus WebMail.

Allow Flash content	(Default) Select this option to enable the system to intermediate Flash content through its Content Intermediation Engine. Note that the system provides limited support for ActionScript 2.0 and Flash Remoting, and does not support XMLSocket connections. The Content Intermediation Engine supports Flash versions 5, 6, 7 and 8, including dynamic rewriting of internal Web links during an access request. We support the rewriting of Actionscript in Flash. The calls in Actionscript that are supported are: load, send, sendAndLoad, loadVariables, loadMovie, loadVariablesNum, loadMovieNum, loadClip, loadSound, apply, connect on classes of XML, Sound, MovieClip, NetConnection, and MovieClipLoader. The eval equivalent of Actionscript is not supported. Therefore we recommend that the above function calls not be embedded in an Actionscript string object. Note, Flash applications that use the XMLSocket object or Flash Remoting are not supported. For more information, see the *Content Intermediation Engine Best Practices Guide*.
Persistent cookies	(Default) Select this option to enable users to customize their browsing experiences by enabling them to keep persistent cookies. By default, the system flushes Web cookies that are stored during a user session. A user can delete cookies through the Advanced Preferences page if you enable this option.
Unrewritten pages open in new window	Select this option to configure the system to open content in a new browser window when a user access an unrewritten Web page. Opening content in a new windows can help remind users that they still have a secure session. When a user request is made to a resource to which this option applies, the system displays a page that contains a link to the requested resource and directs the users to click on the link. This link opens the resource in a new browser window and the page from which the request originates continues to display in the system. If you uncheck this box, users might not realize that their session is still active and that to return to the system, they need to use the browser's Back button. Users must return to the system to sign out. If they simply close the browser window, their sessions remain active until the session time limit expires.
Allow browsing untrusted SSL Web servers	(Default) Select this option to allow access to untrusted web sites through the system. Untrusted web sites are those whose server certificates are not installed, expired, or revoked through the System > Configuration > Certificates > Trusted Servers CAs tab of the admin console. NOTE: If a web page has internal references to files within a SCRIPT tag and these files are hosted on different HTTPS servers that have SSL certificates not trusted by the system, the web page does not render correctly. In these cases, the Warn users about the certificate problems option must be disabled. • **Warn users about the certificate problems**. (Default) Select this option to warn users about the certificate problems option and the user accesses non-HTML content (such as images, js, and css) served from a different SSL server than the HTML page, the page containing the links may not display correctly. You can avoid this problem either by deselecting this option or by uploading a valid production SSL certificate on the servers that serve the non-HTML content. If enabled, display a warning to the user when he first accesses an untrusted web site telling him why the site's certificate is untrusted and allowing him to either continue or cancel. If the user chooses to continue after viewing the warning, the system does not display any more warnings for that site during the current session. NOTE: This option is not applicable for auth-only URLs (for example, ActiveSync) and Secure Mail URLs. • **Allow users to bypass warnings on a server-by-server basis**. Select this option to allow the user to suppress all further warnings for an untrusted web site. If a user chooses this option, he never sees a warning for this site again, provided that he accesses it from the current device or cluster. If you choose to allow users to access untrusted web sites without seeing a warning, the system still logs a message to the user access log whenever a user navigates to an untrusted site. Also note that if a user chooses to suppress warnings, he can clear the persistent settings of the untrusted web sites using the Delete Passwords option in the System > Preferences > Advanced tab in the end user console.
Rewrite file:// URLs	Select this option to to rewrite file:// URLs so that they are routed through the system's file browsing CGI.
Rewrite links in PDF files	Select this option to rewrite hyperlinks in PDFs.

HTTP Connection Timeout

HTTP Connection Timeout	Specify the duration to wait for a response from an HTTP server before timing out and closing the connection. Use values from 30 to 1800 seconds (default is 240).
	Higher timeout values might exhaust system resources if applications do not close connections properly or take too long to close the connections. Unless an application requires a higher timeout value, we recommend accepting the default value.
WebSocket Connection Timeout	Specify the duration to wait for data transfer between the client and server. Use values from 30 to 1800 seconds (default is 900).
	WebSocket is a web technology that provides bidirectional, full-duplex communication channels over a single TCP connection. This provides a mechanism for browser-based applications that need two-way communication with servers that do not rely on opening multiple HTTP connections. Communication is done over the regular TCP port numbers 80 or 443.
	Currently, only the following web resource policies support WebSocket:
	• Web ACL Access
	• Passthrough Proxy
	• Options
	The WebSocket URL that starts with **ws://** or **wss://** is not allowed in any of the web resource profile pages, web resource policies or web bookmark pages.
	The following Web options under Roles accept WebSocket requests:
	• Mask hostnames while browsing
	• Persistent cookies
	• Allow browsing untrusted SSL web sites
ActiveSyncLongLived Connection Timeout	Specify the duration of a long-lived request used to synchronize an iOS device with a Microsoft Exchange server (Secure Mail must be enabled for the role). When the request expires, the device issues a new request. Use values from 30 to 7200 seconds. Microsoft recommends using 1800 seconds (the default).

Related Documentation

- Web Rewriting

- Defining Resource Profile Bookmarks

- Using Custom Expressions in Rule Configuration

- Creating a Hosted Java Applet Resource Profile

Resource Policy Overview

When you enable the Web access feature for a role, you need to create resource policies that specify which resources a user can access, whether or not to rewrite the content requested by the user, and caching, applet, or single sign-on requirements. For every Web request, the system first evaluates the rewriting policies you configure. If the user's request is to a resource specified as "don't rewrite" due to either a selective rewriting or passthrough proxy resource policy, then forward the user's request to the appropriate back-end resource. Otherwise, the system continues to evaluate those resource policies corresponding to the request, such as Java resource policies for a request to fetch a Java applet. After matching a user's request to a resource listed in a relevant policy, the system performs the action specified for the resource.

You can create resource policies through the standard interface (as described in this section) or through resource profiles (recommended method).

When writing a Web resource policy, you need to supply key information:

- Resources—A resource policy must specify one or more resources to which the policy applies. When writing a Web policy, you need to specify Web servers or specific URLs, as explained in the section that follows.

- Roles—A resource policy must specify the roles to which it applies. When a user makes a request, the system determines what policies apply to the role and then evaluates those policies that correspond to the request.

- Actions—Each type of resource policy performs a certain action, which is either to allow or deny a resource or to perform or not perform some function, such as rewrite content, re-sign an applet, or post Web data. You can also write detailed rules that apply more conditions to a user request.

The system platform's engine that evaluates resource policies requires that the resources listed in a policy's Resources list follow a canonical format.

Canonical Format

This section outlines special considerations you must consider when specifying a Web resource using the canonical format.

[protocol://]host[:ports][/path]

The four components are:

- Protocol (optional)—Possible values: http and https (case-insensitive)

  If the protocol is missing, then both http and https are assumed. If a protocol is specified, then the delimiter "://" is required. No special characters are allowed.

- Host (required)—Possible values:

  - DNS Hostname—For example: www.pulsesecure.net

    Allowed special characters are described in Table 9.

Table85: DNS Hostname Special Characters

*	Matches ALL characters
%	Matches any character except dot (.)
?	Matches exactly one character

  - IP address/Netmask—The IP address needs to be in the format: a.b.c.d

    The netmask can be in one of two formats:

    - Prefix: High order bits

    - IP: a.b.c.d

    For example: 10.11.149.2/24 or 10.11.149.2/255.255.255.0

    No special characters are allowed.

- Ports—You must specify a port when specifying IP/netmask as a resource. The port is optional when specifying a DNS hostname. If a port is specified, then the delimiter ":" is required. For example: 10.11.149.2/255.255.255.0:* Table 86 lists the possible port values.

Table86: Possible Port Values

*	Matches ALL ports; no other special characters are allowed
port[,port]*	A comma-delimited list of single ports. Valid port numbers are [1-65535].
[port1]-[port2]	A range of ports, from port1 to port2, inclusive.

> *i* NOTE: You can mix port lists and port ranges, such as: 80,443,8080-8090

If the port is missing, then the default port 80 is assigned for http, 443 for https.

- Path (optional)—If the path is missing, then star (*) is assumed, meaning ALL paths match. If a path is specified, then the delimiter "/" is required. No other special characters are supported. For example:

  - http://www.pulsesecure.net:80/*

  - https://www.pulsesecure.net:443/intranet/*

  - *.yahoo.com:80,443/*

  - %.danastreet.net:80/share/users/<username>/*

Writing a Web Access Resource Policy

Web access resource policies control which Web resources users can access in order to connect to the Internet, intranet, or extranet. You can deny or allow access to Web resources by URL or IP range. For URLs, you can use the "*" and "?" wildcards to efficiently specify multiple hostnames and paths. For resources that you specify by hostname, you can also choose either HTTP, HTTPS, or both protocols.

To write a Web Access resource policy:

1. In the admin console, choose **Users > Resource Policies > Web > Web ACL**.

2. On the Web Access Policies page, click **New Policy**.

3. On the New Policy page, enter a name to label this policy and optionally a description.

4. In the Resources section, specify the resources to which this policy applies.

5. In the Roles section, specify:

   - **Policy applies to ALL roles**—To apply this policy to all users.

   - **Policy applies to SELECTED roles**—To apply this policy only to users who are mapped to roles in the Selected roles list. Make sure to add roles to this list from the Available roles list.

   - **Policy applies to all roles OTHER THAN those selected below** —To apply this policy to all users except for those who map to the roles in the Selected roles list. Make sure to add roles to this list from the Available roles list.

6. In the Action section, specify:

   - **Allow access**—To grant access to the resources specified in the Resources list.

- **Deny access**—To deny access to the resources specified in the Resources list.

- **Use Detailed Rules**—To specify one or more detailed rules for this policy.

7. Click **Save Changes**.

8. On the Web Access Policies page, order the policies according to how you want to evaluate them. Keep in mind that once the system matches the resource requested by the user to a resource in a policy's (or a detailed rule's) Resource list, it performs the specified action and stops processing policies.

Related Documentation

- Defining a Resource Profile

- Creating Detailed Rules for Resource Policies

Defining Single Sign-On Policies

Single sign-on policies enable you to automatically pass user credentials to the Web application specified in your policy. You can configure single sign-on policies to intercept basic authentication, Kerberos and NTLM challenges and display an intermediate sign-in page to collect credentials for the Web resource. Or, you can post the credentials and headers that you specify to the Web application.

Related Documentation

- About Basic, NTLM and Kerberos Resources

- Writing the Basic, NTLM and Kerberos Resources

- Writing a Basic Authentication, NTLM or Kerberos Intermediation Resource Policy

- Writing a Remote SSO Form POST Resource Policy

- Writing a Remote SSO Headers/Cookies Resource Policy

About Basic, NTLM and Kerberos Resources

Use the SSO > General tab to set up the basic, NTLM and Kerberos credentials. The credentials you define here are used when defining Web resource profiles with SSO autopolicies and Web resource policies.

The following outlines the basic ideas behind the handling of SSO:

- The system will do Kerberos if challenged with Negotiate header, NTLM if challenged with NTLM header and Basic Auth if challenged with Basic.

- If the system receives multiple challenges, the order of preference is:

  - Kerberos

  - NTLM

  - Basic

- The system will first try constrained delegation if the service is configured in a service list.

- Policy configurations override any settings in the SSO > General tab.

- Disabling SSO or disabling all sections in the General tab prevents single sign-on. However, the system will continue to intermediate and display an intermediation page to the end user.

- Basic authentication intermediation can be explicitly turned off in a policy. For kerberos and NTLM, the system will always intermediate.

- Depending on the SSO used, the intermediation page will show different fields for the end user to complete:

  - Basic authentication intermediation page displays username and password fields

  - NTLM intermediation page displays username, password and domain fields

  - Kerberos intermediation page displays username, password and realm fields

- For constrained delegation, you must defined a policy and specify roles. Entering data in the General tab only is not sufficient.

- If no policies are configured for single sign-on, the system uses the default system credentials.

- If credentials are defined, the order of preference is:

  - System credentials

  - Variable credentials

  - Fixed or static credentials

- For fixed or static credentials, you must defined a policy and specify roles. Entering data in the General tab only is not sufficient.

- If there is a policy match, the credential and protocol of the policy is used. If the policy fails to authenticate, the fallback mechanism defined in the policy is used. If the policy protocol does not match the protocol of the challenge, the logic defined in the General tab is used.

- When upgrading a device or performing a new install, the default SSO policy of BasicAuthNoSSO is preserved. Even if all sections of the General tab are enabled, SSO will not be enabled until the BasicAuthNoSSO policy is deleted.

Related Documentation

- Writing the Basic, NTLM and Kerberos Resources

- Defining Single Sign-On Policies

Writing the Basic, NTLM and Kerberos Resources

To set up the basic, NTLM and Kerberos resources:

1. In the admin console, select **Users > Resource Policies > Web**.

2. If your administrator view is not already configured to show SSO policies, make the following modifications:

   1. Click the **Customize** button in the upper right corner of the page.

   2. Select the **SSO** check box.

   3. Select the **General** check box below the SSO check box.

   4. Click **OK**.

5. Select the **SSO > General** tab.

6. Select **Enable kerberos** to enable Kerberos SSO. You can then define the type of intermediation: constrained delegation or Connect Secure. If you do not define any

intermediation types, the system attempts to figure out the realm from the hostname and performs SSO using the system credentials.

For realm intermediation, enter the following and click **Add**:

- **Realm**—Enter the Kerberos realm name. For example, KERBER.NET. The system uses KERBER.NET to obtain the list of Key Distribution Centers (KDCs).

- **Site Name**—(optional) Enter the Active Directory site names. Use this field to have the system contact the KDC at a specific site. For example, if site name is Sunnyvale and realm is KERBER.NET, then the system uses Sunnyvale•KERBER.NET to get a list of KDCs. Note that the Active Directory must have the sites defined and DNS should be configured to return the KDCs in the site.

- **Pattern List**—Enter the hostnames mapped to the Kerberos realm. You can enter wildcard characters, such as *.y.com, *.kerber.net, or *.*. Note the following:

  - Make sure that realms to not have hostnames matching a subset of the patterns defined for another realm.

  - You do not need to define a pattern if all servers follow the mirrored DNS namespace convention. The system determines the realm from the hostname.

  - All disjoined hostname patterns must be defined.

  - You can use * as the default realm. Do not list more than one * when defining multiple realms.

- **KDC**—Enter the hostname or IP address of the Key Distribution Centers if DNS is unavailable or if you want the system to contact a specific KDC for tickets. If you enter a KDC, the system does not use DNS to obtain the list of KDCs based on the values entered in the Site Name and Realm fields.

For constrained delegation intermediation, enter the following and click **Add**:

- **Label**—Enter a name to uniquely identify this row. No external mapping is made to the label value.

- **Realm**—Select the realm to use. The drop-down list is populated by values in the Realm Definition table.

- **Principal Account**—Enter the constrained delegation account to use to get constrained delegation tickets on behalf of the user.

- **Password**—Enter the constrained delegation account password.

- **Service List**—Select the service list to use. Click Edit to define and upload service lists. The list should be an exact match with the service list in Active Directory if you want to perform constrained delegation for all the services. Hostnames must be an exact match.

For more information about constrained delegation, see http://msdn.microsoft.com/en-us/library/aa480585.aspx.

For system intermediation, enter the following and click **Add**:

- **Label**—Enter a name to uniquely identify this row. No external mapping is made to the label value.

- **Realm**—Select the realm to use. The drop-down list is populated by values in the Realm Definition table.

- **Credential Type**—Select one of the following credential types:

  - **System credentials**—Use the set of user credentials, such as primary and secondary authorization credentials, stored on the device. If you select this option, you do not need to enter values in the Username and Password fields.

  - **Variable**—Allow tokens such as <username> and <password> to be used in the username and Variable Password fields.

  - **Static**—Use the username and password exactly as they are entered in the username and password fields.

- **Username and Password**—Enter the account username and password. If you select Variable as the credential type, you can enter the username token here. For example, <username>.

- **Variable Password**—If you select Variable as the credential type, enter the password token here. For example, <password>.

- **Fallback to NTLM V2**—Select this option to fallback to NTLM V2 if Kerberos fails. If you do not select this option and Kerberos SSO fails, an intermediation page appears.

7. Select **Enable NTLM** to enable NTLM SSO. If you do not enter any configuration information, the system attempts to figure out the domain from the hostname and performs SSO using the system credentials.

   > NOTE: Do not edit or delete the default system credential.

- **Label**—Enter a name to uniquely identify this row. No external mapping is made to the label value.

- **Domain**—Enter the Active Directory domain name here.

- **Credential Type**—Select one of the following credential types:

  - **System credentials**—Use the set of user credentials, such as primary and secondary authorization credentials, stored on the device. If you select this option, you do not need to enter values in the Username and Password fields.

  - **Variable**—Allow tokens such as <username> and <password> to be used in the Username and Variable Password fields.

  - **Static**—Use the username and password exactly as they are entered in the username and password fields.

- **Username and Password**—Enter the account username and password. If you select Variable as the credential type, you can enter the username token here. For example, <username>.

- **Variable Password**—If you select Variable as the credential type, enter the password token here. For example, <password>.

- **Fallback to NTLM V1**—Select this option to fallback to NTLM V1 if SSO fails. If you do not select this option and SSO fails, only NTLM V2 is attempted. An intermediation page appears if NTLM V2 fails.

8. Select **Enable Basic Authentication** to enable basic authentication SSO. If you select this option but do not set up any configuration data, the system will attempt SSO using system credentials.

> **_i_** NOTE: Do not edit or delete the default system credential.

- **Label**—Enter a name to uniquely identify this row. No external mapping is made to the label value.

- **Credential Type**—Select one of the following credential types:

  - **System credentials**—Use the set of user credentials, such as primary and secondary authorization credentials, stored on the device. If you select this option, you do not need to enter values in the Username and Password fields.

  - **Variable**—Allow tokens such as <username> and <password> to be used in the Username and Variable Password fields.

  - **Static**—Use the username and password exactly as they are entered in the username and password fields.

- **Username and Password**—Enter the account username and password. If you select Variable as the credential type, you can enter the username token here. For example, <username>.

- **Variable Password**—If you select Variable as the credential type, enter the password token here. For example, <password>.

- **Pattern List**—Enter the hostnames mapped to the Kerberos realm. You can enter wildcard characters, such as *.y.com, *.kerber.net, or *.*. Note the following:

  - Make sure that realms to not have hostnames matching a subset of the patterns defined for another realm.

  - You do not need to define a patter if all servers follow the mirrored DNS namespace convention. The system determines the realm from the hostname.

  - All disjoined hostname patterns must be defined.

  - You can use * as the default realm. Do not list more than one * when defining multiple realms.

  - You can use * as the default domain. Do not list more than one * when defining multiple domains.

Related Documentation

- Using Custom Expressions in Rule Configuration

- About Single Sign-On

Writing a Basic Authentication, NTLM or Kerberos Intermediation Resource Policy

Basic Authentication, NTLM or Kerberos Intermediation resource policies enable you to control NTLM and Kerberos intermediation on the system. If a user accesses a Web resource that sends a basic authentication challenge, the system can intercept the challenge, display an intermediate sign-in page to collect credentials for the Web resource, and then rewrite the credentials along with the entire challenge/response sequence.

The initial HTTP request generated for an NTLM protected server should be for a request that results in HTML content. If SSO is not enabled or if the SSO credentials fail, the system

responds with an HTML page to gather user credentials. If the browser is expecting non-HTML content, the browser rejects the response and the navigation to the resource fails.

With the Kerberos Intermediation resource policy, backend web applications protected by Kerberos are accessible to end users. For example, a user logs in to a device using Active Directory as the authentication server and the authentication protocol is Kerberos. When the user browses to a Kerberos-protected server, the user is single-signed on to the backend server and is not prompted for credentials. Or, if a user logs in to a device using an authentication protocol other than Kerberos and then browses to a Kerberos-protected server. Depending on the settings in Kerberos Intermediation resource policy and the configured Kerberos authentication server, the user will either be authenticated by the rewriter or the user will be prompted to enter a username and password.

To write a Basic Authentication, NTLM or Kerberos Intermediation resource policy:

1. In the admin console, select **Users > Resource Policies > Web**.

2. If your administrator view is not already configured to show SSO policies, make the following modifications:

    1. Click the **Customize** button in the upper right corner of the page.

    2. Select the **SSO** check box.

    3. Select the **Kerberos/Basic Auth/NTLM** check box below the SSO check box.

    4. Click **OK**.

5. Select the **SSO > Kerberos/NTLM/BasicAuth** tab.

6. Click **New Policy**.

7. Enter a name to label this policy (required) and a description of the policy (optional).

8. In the Resources section, specify the resources to which this policy applies.

    If you want to automatically post values to a specific URL when an end user clicks on a bookmark, the resource that you enter here must exactly match the URL that you specify in the Users > User Roles > Role > Web > Bookmarks page of the admin console.

9. In the Roles section, specify:

    - **Policy applies to ALL roles**—To apply this policy to all users.

    - **Policy applies to SELECTED roles**—To apply this policy only to users who are mapped to roles in the Selected roles list. Make sure to add roles to this list from the Available roles list.

    - **Policy applies to all roles OTHER THAN those selected below**—To apply this policy to all users except for those who map to the roles in the Selected roles list. Make sure to add roles to this list from the Available roles list.

10. In the Action section, specify:

    - **Disable SSO**—Disables automatic SSO authentication for this user role and, instead, prompts the user for sign-in credentials.

    - **Basic**—This option uses the Basic Authentication Intermediation method to control SSO behavior.

- **Enable Intermediation**—Select the credentials to use. If this pull-down menu is blank, no basic authentication SSO settings are defined in the SSO General tab.

- **Disable Intermediation**—When you select this option, the system does not intermediate the challenge/response sequence.

  The system always intermediates requests to Web proxies that require basic authentication, even if you select Disable Intermediation.

  Although you are given an option to disable basic authentication intermediation, we do not recommend this option, as it is a very insecure authentication method and, in some cases, can transmit user credentials over the network in clear (unencrypted) text.

- **NTLM**—This option specifies that the system use the Microsoft NTLM Intermediation method to control SSO behavior.

  - Select the credentials to use. If this pull-down menu is blank, no NTLM SSO settings are defined in the SSO General tab.

  - Select the **Fallback to NTLM V1** option to try both NTLM V1 and NTLM V2. If you do not select this option, the system falls back only to NTLM V2. An intermediation page appear if SSO fails.

- **Kerberos**—This option specifies that the system use the Kerberos Intermediation method to control SSO behavior.

  - Select the credentials to use. If this pull-down menu is blank, no kerberos SSO settings are defined in the SSO General tab

  - Select the **Fallback to NTLM V2** option to fallback only to NTLM V2 if kerberos fails. If you do not select this option, a Kerberos intermediation page appears if Kerberos SSO fails.

- **Constrained Delegation**—This option specifies that the system use the constrained delegation intermediation method to control SSO behavior.

  - Select the credentials to use. If this pull-down menu is blank, no constrained delegation SSO settings are defined in the SSO General tab.

  - Select the **Fallback to Kerberos** option to fallback to Kerberos if constrained delegation fails. If you select this option, an intermediation page appears if constrained delegation fails. If you do not select this option and constrained delegation fails, an error page appears.

- **Use Detailed Rules**—To specify one or more detailed rules for this policy.

11. Click **Save Changes**.

12. On the Basic Auth, NTLM and Kerberos policies page, order the policies according to how you want to evaluate them. Keep in mind that once the system matches the resource requested by the user to a resource in a policy's (or a detailed rule's) Resource list, it performs the specified action and stops processing policies.

Check the activity events listed in the user log if you encounter any problems.

Related Documentation

- Specifying Resources for a Resource Policy

- Resource Policy Components

- Writing the Basic, NTLM and Kerberos Resources

Writing a Remote SSO Form POST Resource Policy

Remote SSO Form POST resource policies specify Web applications to which the system posts data. This data can include a user's username and password, as well as system data stored by system variables.

To write a remote SSO Form POST resource policy:

1. In the admin console, navigate to **Users > Resource Policies > Web**.

2. If your administrator view is not already configured to show SSO policies, make the following modifications:

   1. Click the **Customize** button in the upper right corner of the page.

   2. Select the **SSO** check box.

   3. Select the **Form Post** check box below the SSO check box.

   4. Click **OK**.

5. Select the **SSO> Form Post** tab.

6. On the Form POST Policies page, click **New Policy**.

7. Enter a name to label this policy (required) and a description of the policy (optional).

8. In the Resources section, specify the application's sign-in page, such as: http://yourcompany.com.

   If you want to automatically post values to a specific URL when an end user clicks on a bookmark, the resource that you enter here must exactly match the URL that you specify in the Users > User Roles > Role > Web > Bookmarks page of the admin console.

9. In the Roles section, specify:

   - **Policy applies to ALL roles**—To apply this policy to all users.

   - **Policy applies to SELECTED roles**—To apply this policy only to users who are mapped to roles in the Selected roles list. Make sure to add roles to this list from the Available roles list.

   - **Policy applies to all roles OTHER THAN those selected below**—To apply this policy to all users except for those who map to the roles in the Selected roles list. Make sure to add roles to this list from the Available roles list.

10. In the Action section, specify:

    - **Perform the POST defined below**—Perform a form POST with the user data specified in the POST details section to the specified URL when a user makes a request to a resource specified in the Resources list.

    - **Do NOT perform the POST defined below**—Do not perform a form POST with the user data specified in the POST details section.

    - **Use Detailed Rules**—Select this option to specify one or more detailed rules for this policy.

11. In the POST details section:

    - In the POST to URL field, specify the absolute URL where the application posts the user's credentials, such as: http://yourcompany.com/login.cgi. You can determine the

appropriate URL using a TCP dump or by viewing the application's sign-in page source and searching for the POST parameter in the FORM tag. (Wildcard characters are not supported in this field.)

- Check **Deny direct login** for this resource if you do not want users to be able to access the URL directly.

- Select the **Allow multiple POSTs to this resource** check box if you want to send POST and cookie values to the resource multiple times if required. If you do not select this option, the system does not attempt single sign-on when a user requests the same resource more than once during the same session.

- Specify the user data to post and user modification permission:

  - **User label**—The label that appears on a user's Preferences page. This field is required if you either enable or require users to modify data to post to back-end applications.

  - **Name**—The name to identify the data of the Value field. (The back-end application should expect this name.)

  - **Value**—The value to post to the form for the specified Name. You can enter static data, a system variable, or system session variables containing username and password values.

  - **User modifiable?** setting—Set to Not modifiable if you do not want the user to be able to change the information in the Value field. Set to User CAN change value if you want the user to have the option of specifying data for a back-end application. Set to User MUST change value if users must enter additional data in order to access a back-end application. If you choose either of the latter settings, a field for data entry appears on the user's Advanced Preferences page. This field is labeled using the data you enter in the User label field. If you enter a value in the Value field, this data appears in the field but is editable.

12. Click **Save Changes**.

13. On the Form POST Policies page, order the policies according to how you want to evaluate them. Keep in mind that once the system matches the resource requested by the user to a resource in a policy's (or a detailed rule's) Resource list, it performs the specified action and stops processing policies.

Related Documentation

- Remote SSO Overview

- Specifying Resources for a Resource Policy

- Writing a Detailed Rule for Resource Policies

- Using Custom Expressions in Rule Configuration

- Multiple Sign-In Credentials Execution

Writing a Remote SSO Headers/Cookies Resource Policy

Remote SSO Headers/Cookies resource policies specify customized Web applications to which the system posts custom headers and cookies.

When creating a Headers/Cookies policy, note that the system does not parse or "understand" the headers that you enter in this section. For instance, if you add an Accept-Encoding: gzip or

Accept-Encoding:deflate header, it does not mean that the system can handle gzip content or deflated content.

To write a remote SSO Headers/Cookies resource policy:

1. In the admin console, select **Users > Resource Policies > Web**.

2. If your administrator view is not already configured to show SSO policies, make the following modifications:

   1. Click the **Customize** button in the upper right corner of the page.

   2. Select the **SSO** check box.

   3. Select the **Headers/Cookies** check box below the SSO check box.

   4. Click **OK**.

5. Select the **SSO > Headers/Cookies** tab.

6. On the Headers/Cookies Policies page, click **New Policy**.

7. Enter a name to label this policy (required) and a description of the policy (optional).

8. In the Resources section, specify the resources to which this policy applies.

9. In the Roles section, specify:

   - **Policy applies to ALL roles**—To apply this policy to all users.

   - **Policy applies to SELECTED roles**—To apply this policy only to users who are mapped to roles in the Selected roles list. Make sure to add roles to this list from the Available roles list.

   - **Policy applies to all roles OTHER THAN those selected below**—To apply this policy to all users except for those who map to the roles in the Selected roles list. Make sure to add roles to this list from the Available roles list.

10. In the Action section, specify:

    - **Append headers as defined below**—Post the user data specified in the POST details section to the specified URL when a user makes a request to a resource specified in the Resources list.

    - **Do NOT append headers as defined below**—Do not post the user data specified in the POST details section to the specified URL when a user makes a request to a resource specified in the Resources list.

    - **Use Detailed Rules**—Select this option to specify one or more detailed rules for this policy.

11. In the Headers and values section, specify the:

    - **Header name**—The text for to send as header data.

    - **Value**—The value for the specified header.

        > *i* NOTE: If you need to forward a cookie to a backend server, you must set the Header Name field to "Cookie" and the Value field to "CookieName=CookieValue".

12. Click **Save Changes**.

---

13. On the Headers/Cookies Policies page, order the policies according to how you want to evaluate them. Keep in mind that once the system matches the resource requested by the user to a resource in a policy's (or a detailed rule's) Resource list, it performs the specified action and stops processing policies.

Related Documentation

- Remote SSO Overview
- Specifying Resources for a Resource Policy
- Writing a Detailed Rule for Resource Policies
- Resource Policies

Writing a Web Caching Resource Policy

To write a Web Caching resource policy:

1. In the admin console, select **Users > Resource Policies > Web**.

2. If your administrator view is not already configured to show caching policies, make the following modifications:

    1. Click the **Customize** button in the upper right corner of the page.

    2. Select the **Caching** check box.

    3. Select the **Policies** check box below the Caching check box.

    4. Click **OK**.

3. Select the **Caching > Policies** tab.

4. On the Web Caching Policies page, click **New Policy**.

5. Enter a name to label this policy (required) and a description of the policy (optional).

6. In the Resources section, specify the resources to which this policy applies.

7. In the Roles section, specify:

    - **Policy applies to ALL roles**—To apply this policy to all users.

    - **Policy applies to SELECTED roles**—To apply this policy only to users who are mapped to roles in the Selected roles list. Make sure to add roles to this list from the Available roles list.

    - **Policy applies to all roles OTHER THAN those selected below**—To apply this policy to all users except for those who map to the roles in the Selected roles list. Make sure to add roles to this list from the Available roles list.

8. In the Action section, select one of the following options:

    - **Smart Caching (send headers appropriate for content and browser)** —Select this option to send a cache-control:no-store header or a cache-control:no-cache header based on the user's Web browser and content type.

        When you select this option, the system makes media files and zip files work properly by removing their origin server's cache-control headers. For example, the following logic searches for "msie" or "windows-media-player" in user-agent headers in order to remove cache or cache-control:no-store response headers and make the files cacheable:

```
(if content type has "audio/x-pn-realaudio" OR
 if content type begins with "video/" OR
 if content type begins with "audio/" OR
 if content type is "application/octet-stream" and the file extension
begins with "rm" or "ram"
)
```

If the system finds "msie" or "windows-media-player" in the user-agent header and any of the following apply:

- Request is for Flash, .xls, .pps, .ppt files

- Content-type is application/, text/rtf, text/xml, model/

- Origin server sends a content-disposition header

then the system sends the cache-control:no-store header and removes the origin server's cache-control header.

In all other cases, the system adds the pragma:no-cache or cache-control:no-store response headers.

Citrix .ica and QuickPlace files get some special treatment. Citrix .ica files are always cacheable and get cache-control-private as well. QuickPlace files that do not match a specified rule files (which takes precedence) get CCNS and cache-control:private.

Also note that if you select this option, enable GZIP compression, and try to access a text file attachment using Domino Web Access 6.5 through Internet Explorer, you cannot open the attachment. To enable text attachments, you must either install the Internet Explorer 323308 patch or enable the Don't Cache (send "Cache Control: No Store") option.

- **Don't Cache (send "Cache Control: No Store")**—Select this option to deliver attachments to Internet Explorer without saving them to the disk. (The browser temporarily writes files to the disk, but immediately removes them once it has opened the file in the browser.) When you select this option, the system removes the origin server's cache-control header and adds a cache-control:no-store response header if the user-agent string sent by the browser contains "msie" or "windows-media-player."

This option might slow browsing by causing repeated content fetches, which can cause performance issues on very slow connections. Alternatively, you can specify a policy that allows certain kinds of content to be cached, such as images that do not exceed a specified size limit.

- **Don't Cache (send "Pragma: No Cache")**—Select this option to prevent the user's browser from caching files to the disk. When you select this option, the system adds the standard HTTP pragma:no-cache header and cache-control:no-cache (CCNC) header (HTTP 1.1) to response files. Also, the system does not forward the origin server's caching headers, such as age, date, etag, last-modified, expires.

When no-cache headers are present on certain types of attachments (PDF, PPT, streaming files), Internet Explorer does not properly render the documents because the rendering process requires the browser to temporarily writes these files to cache.

- **Unchanged (do not add/modify caching headers)**—Select this option to not add the pragma:no-cache or cache-control:no-store response headers and forwards the origin server's caching headers.

- **Remove Cache-Control: No-Cache|No-Store**—Select this option to help "cache" files sent by web applications in an HTTPS environment. This option removes the Cache

Control:No Cache and Pragma:no-cache headers. Removing these headers is necessary to allow the successful download of certain file types. These headers work fine in an HTTP environment, but fail in an HTTPS environment where the associated pages become uncacheable, preventing the user's web browser from downloading the pages.

Use this option when you want the end user to have the ability to download and open a file that will be opened by another third party application. For example, zip files and wav files are stored on disk and opened by another application.

- **Use Detailed Rules**—To specify one or more detailed rules for this policy.

11.　　Click **Save Changes**.

12.　　On the Web Caching Policies page, order the policies according to how you want to evaluate them. Keep in mind that once the system matches the resource requested by the user to a resource in a policy's (or a detailed rule's) Resource list, it performs the specified action and stops processing policies.

Related Documentation

- Writing a Detailed Rule for Resource Policies

- Defining Resource Policies: General Options

About OWA and Lotus Notes Caching Resource Policies

Table 87 and Table 88 include examples of some of the content types that the system supports with the Outlook Web Access (OWA) and Lotus iNotes applications. Additionally, it specifies the cache control directives that you must implement in Microsoft Internet Explorer in order to support opening and saving the specified content types.

Note that for performance reasons, we recommend creating caching policies for everything in the iNotes directory.

Table87: OWA Caching Resource Policies

Attachment type	To open the attachment, use:	To save the attachment, use:
zip	Cache	Smart caching
ppt	Smart caching	Smart caching
doc	Smart caching	Smart caching
xls	Smart caching	Smart caching
pdf	Smart caching	Smart caching
txt	Cache	Cache control: No store
html	Smart caching	Cache control: No store

Table88: iNotes Caching Resource Policies

Attachment type	To open the attachment, use:	To save the attachment, use:
zip	Cache control: No store	Cache control: No store

ppt	Cache control: No store	Cache control: No store
doc	Smart caching	Smart caching
xls	Cache control: No store	Cache control: No store
pdf	Cache control: No store	Cache control: No store
txt	Cache control: No store	Cache control: No store
html	Cache control: No store	Cache control: No store
other file types	Cache control: No store	Cache control: No store

Related Documentation

- Standard Application Support: MS Outlook

- Standard Application Support: Lotus Notes

Specifying General Caching Options

You can use caching options to specify the maximum image file size that is cached on a client. If the content-type header from the origin server begins with "image/" and the content-length header specifies a size less than the maximum size configured for this option, then the system passes along the origin server's caching headers. Otherwise, the system treats the request as though caching is disabled.

To specify caching options:

1. In the admin console, choose **Users > Resource Policies > Web**.

2. If your administrator view is not already configured to show caching policies, make the following modifications:

    1. Click the **Customize** button in the upper right corner of the page.

    2. Select the **Caching** check box.

    3. Select the **Options** check box below the Caching check box.

    4. Click **OK**.

5. Select the **Caching > Options** tab.

6. On the Caching Options page, specify a maximum allowable image size in the Clients should cache all images less than field.

7. On the Caching Options page, specify a maximum allowable image size in the Clients should cache all images less than field.

Writing a Java Access Control Resource Policy

Java access control resource policies control to which servers and ports Java applets can connect.

To write a Java access control resource policy:

1. In the admin console, select **Users > Resource Policies > Web**.

2. If your administrator view is not already configured to show Java policies, make the following modifications:

1. Click the **Customize** button in the upper right corner of the page.

2. Select the **Java** check box.

3. Select the **Access Control** check box below the Java check box.

4. Click **OK**.

5. Select the **Java > Access Control** tab.

6. On the Java Access Policies page, click **New Policy**.

7. Enter a name to label this policy (required) and a description of the policy (optional).

8. In the Resources section, specify the resources to which this policy applies.

9. In the Roles section, specify:

   - **Policy applies to ALL roles**—To apply this policy to all users.

   - **Policy applies to SELECTED roles**—To apply this policy only to users who are mapped to roles in the Selected roles list. Make sure to add roles to this list from the Available roles list.

   - **Policy applies to all roles OTHER THAN those selected below** —To apply this policy to all users except for those who map to the roles in the Selected roles list. Make sure to add roles to this list from the Available roles list.

10. In the Action section, specify:

    - **Allow socket access**—To enable Java applets to connect to the servers (and optionally ports) in the Resources list.

    - **Deny socket access**—To prevent Java applets from connecting to the servers (and optionally ports) in the Resources list.

    - **Use Detailed Rules**—To specify one or more detailed rules for this policy.

11. Click **Save Changes**.

12. On the Java Access Policies page, order the policies according to how you want to evaluate them. Keep in mind that once the system matches the resource requested by the user to a resource in a policy's (or a detailed rule's) Resource list, it performs the specified action and stops processing policies.

13. (Optional) To improve the performance of your Java applications:

    1. Select **Enable Java instrumentation caching** on the Maintenance > System > Options page of the admin console. This option can improve the performance of downloading Java applications.

    2. After you finish configuring the system, cache your Java applet and access it as end user. This action eliminates the performance hit that occurs through the intermediation engine when the first end user accesses the applet.

Related Documentation

- About Hosted Java Applet Templates

- Specifying Resources for a Resource Policy

- Defining Resource Policies: General Options

Writing a Java Code Signing Resource Policy

Java code signing resource policies specify how the system rewrites Java applets. By default, when the system intermediates a signed Java applet, it re-signs the applet with its own certificate, which is not chained to a standard root certificate. When a user requests an applet that performs potentially high-risk tasks, such as accessing network servers, the user's browser displays a security warning that the root is not a trusted root. To forestall this warning, you can import a code-signing certificate that the system uses to re-sign applets that it intermediates.

When configuring Java code signing resource policies, enter the servers from which you trust applets. You can enter a server IP address or domain name. The system only re-signs applets served by a trusted server. If a user requests an applet from server not on the list, the system does not use the imported production certificates to sign the applet, which means the user is prompted by the browser with a security warning. For Sun JVM users, the system additionally checks that the root CA of the original applet certificate is on its list of trusted root certificate authorities.

To write a Java code signing resource policy:

1. In the admin console, choose **Users > Resource Policies > Web**.

2. If your administrator view is not already configured to show java policies, make the following modifications:

   1. Click the **Customize** button in the upper right corner of the page.

   2. Select the **Java** check box.

   3. Select the **Code-Signing** check box below the Java check box.

   4. Click **OK**.

5. Select the **Java > Code-Signing** tab.

6. On the Java Signing Policies page, click **New Policy**.

7. Enter a name to label this policy (required) and description of the policy (optional).

8. In the Resources section, specify the resources to which this policy applies.

9. In the Roles section, specify:

   - **Policy applies to ALL roles**—To apply this policy to all users.

   - **Policy applies to SELECTED roles**—To apply this policy only to users who are mapped to roles in the Selected roles list. Make sure to add roles to this list from the Available roles list.

   - **Policy applies to all roles OTHER THAN those selected below**—To apply this policy to all users except for those who map to the roles in the Selected roles list. Make sure to add roles to this list from the Available roles list.

10. In the Action section, specify:

    - **Resign applets using applet certificate**—The uploaded code-signing certificate will be used to sign the Java applets intermediated by the system.

    - **Resign applets using default certificate**—The system re-signs the applet with its own self-signed code signing certificate that is not chained to a standard root certificate.

    - **Use Detailed Rules**—To specify one or more detailed rules for this policy.

11. Click **Save Changes**.

12. On the Java Signing Policies page, order the policies according to how you want to evaluate them. Keep in mind that once the system matches the resource requested by the user to a resource in a policy's (or a detailed rule's) Resource list, it performs the specified action and stops processing policies.

Related Documentation

- Using Code-Signing CAs
- Specifying Resources for a Resource Policy
- Defining Resource Policies: General Options
- Writing a Detailed Rule for Resource Policies

Creating a Selective Rewriting Resource Policy

Selective rewriting resource policies enable you to define a list of hosts for which you want to intermediate content as well as exceptions to this list. By default, the system intermediates all user requests to Web hosts—unless you have configured the system to serve requests to certain hosts using a different mechanism, such as the Secure Application Manager.

Create a selective rewriting policy if you do not want the system to intermediate traffic from web sites that reside outside of the corporate network, such as yahoo.com, or if you do not want the system to intermediate traffic for client/server applications you have deployed as Web resources, such as Microsoft OWA (Outlook Web Access).

To write a selective rewriting resource policy:

1. In the admin console, choose **Users > Resource Policies > Web**.

2. If your administrator view is not already configured to show rewriting policies, make the following modifications:

   1. Click the **Customize** button in the upper right corner of the page.

   2. Select the **Rewriting** check box.

   3. Select the **Selective Rewriting** check box below the Rewriting check box.

   4. Click **OK**.

5. Select the **Rewriting > Selective Rewriting** tab.

6. On the Web Rewriting Policies page, click **New Policy**.

7. Enter a name to label this policy (required) and a description of the policy (optional).

8. In the Resources section, specify the resources to which this policy applies.

9. In the Roles section, specify:

   - **Policy applies to ALL roles**—To apply this policy to all users.

   - **Policy applies to SELECTED roles**—To apply this policy only to users who are mapped to roles in the Selected roles list. Make sure to add roles to this list from the Available roles list.

- **Policy applies to all roles OTHER THAN those selected below**—To apply this policy to all users except for those who map to the roles in the Selected roles list. Make sure to add roles to this list from the Available roles list.

10.   In the Action section, specify:

- **Rewrite content**—The system intermediates all Web content from the resources specified in the Resources list.

- **Rewrite content as...**—The system intermediates all Web content from the resources specified in the Resources list and rewrites the content as if it were the file type specified in the drop-down list. The available options are:

  - **HTML**—Rewrite content as Hypertext Markup Language (HTML)

  - **XML**—Rewrite content as Extensible Markup Language (XML)

  - **Javascript**—Rewrite content as Java scripting language

  - **VBScript**—Rewrite content as Virtual Basic scripting language

  - **CSS**—Rewrite content as Cascading Style Sheets

  - **XSLT**—Rewrite content as XML Style Sheets

  - **Flash**—Rewrite content as Shockwave Flash

  - **DTD**—Rewrite content as Document Type Definitions (DTD)

  - **HTC**—Rewrite content as HTML component

- **Don't rewrite content: Redirect to target Web server**—The system does not intermediate Web content from the resources specified in the Resources list and automatically redirects the request to the target Web server. This is the default option for all rewrite resource policies that you create. If you select this option, you might want to specify that the system open the unrewritten pages in a new window.

  > NOTE: Do not select this option if the specified content needs to access resources inside your corporate network. For instance, if you specify that the system should not rewrite a particular file, and that file calls another file within your network, the user will see an error.

- **Don't rewrite content: Do not redirect to target Web server**—The system retrieves the content from the original Web server, but does not modify it. This is useful in cases where users may not be able to reach the original server, thus disabling redirection. (For example, if the Web server is not accessible from the public internet because it resides behind a firewall.)

  The **Don't rewrite content: Do not redirect to target Web server** option allows users to download data from network resources via the system, but bypasses the rewriting engine in the process. We recommend you use this feature only when rewriting signed Java applets—not other content types. For other content types such as HTML and Javascript, use the **Don't rewrite content: Redirect to target Web server** option to download an applet via the system, thus enabling direct connections to network resources.

- **Optimize as long lived resource (no rewrite)**—The system maintains a health-check connection for VMware Horizon View HTML5, which persists until the client is disconnected. Specify the following in the Resources text box:

  :/*system/wts,system/heartbeat*

- **Use Detailed Rules**—To specify one or more detailed rules for this policy.

11. Click **Save Changes**.

12. On the Web Rewriting Policies page, order the policies according to how you want to evaluate them. Keep in mind that once the system matches the resource requested by the user to a resource in a policy's (or a detailed rule's) Resource list, it performs the specified action and stops processing policies.

Related Documentation

- Specifying Resources for a Resource Policy

- Defining Resource Policies: General Options

- Writing a Detailed Rule for Resource Policies

Creating a Passthrough Proxy Resource Policy

Passthrough proxy resource policies specify Web applications for which the system performs minimal intermediation. To create a passthrough proxy resource policy, you need to specify two things:

- Which Web application to intermediate with the passthrough proxy

- How the system listens for client requests to the application server

To write a passthrough proxy resource policy:

1. In the admin console, choose **Users > Resource Policies > Web**.

2. If your administrator view is not already configured to show rewriting policies, make the following modifications:

   1. Click the **Customize** button in the upper right corner of the page.

   2. Select the **Rewriting** check box.

   3. Select the **Passthrough Proxy** check box below the Rewriting check box.

   4. Click **OK**.

5. Select the **Rewriting > Passthrough Proxy** tab.

6. On the Passthrough Proxy Policies page, click **New Application**.

7. Enter a name to label this policy (required) and a description of the policy (optional).

8. In the URL field, specify the application server hostname and the port used to access the application internally. Note that you cannot enter a path in this field.

9. Choose the way in which you want to enable the passthrough proxy feature:

   - **Use virtual hostname**—If you choose this option, specify a hostname alias for the application server. When the system receives a client request for the application server hostname alias, it forwards the request to the specified application server port in the URL field.

     If you choose this option, you must also define the name and hostname in the Network Identity section of the System > Network > Internal Port tab. In order to make Sharepoint work successfully through the system, you must select the Override automatic cookie handling check box in Internet Explorer under Tools Internet options > Privacy > Advanced Privacy Settings if the following conditions true:

- You select the **Use virtual hostname** option during Pass Through Proxy configuration.

- The virtual hostname that you specify in your Sharepoint configuration is different from the hostname that you configure through the system setup (that is, if the domains are different).

- You enable persistent cookies through the Users > User Roles > Select Role > General > Session Options page of the admin console.

- **Use SA port**—If you choose this option, specify a unique port in the range 11000-11099. The system listens for client requests to the application server on the specified port and forwards any requests to the application server port specified in the URL field.

10.   In the Action section, specify the method to use to intermediate traffic:

- **Rewrite XML**—If you select this option, the system rewrites URLs contained within XML content. If you disable this option, the system passes the XML content "as is" to the server.

- **Rewrite external links**—If you select this option, the system rewrites all URLs. If you disable this option, the system rewrites only those URLs that contain a hostname specified in the passthrough proxy policy.

- **Block cookies from being sent to the browser**—If you select this option, the system blocks cookies destined for the client's browser. The system stores the cookies locally and sends them to applications whenever they are requested.

- **Host-Header forwarding**—If you select this option, the system passes the hostname as part of the host header instead of the actual host identifier.

  The Host-Header forwarding option is only valid in passthrough proxy Virtual Host mode.

11.   Click **Save Changes**.

12.   On the Pass-through Proxy Policies page, order the policies according to how you want to evaluate them. Keep in mind that once the system matches the application requested by the user to an application specified in a policy's (or a detailed rule's) Resource list, it performs the specified action and stops processing policies.

13.   If you select:

- **Use virtual hostname**, you must also:

  1. Add an entry for each application server hostname alias in your external DNS that resolves to the system.

  2. Upload a wildcard server certificate to the system (recommended).

- **Use SA port**, open traffic to the port you specified for the application server in your corporate firewall.

  If your application listens on multiple ports, configure each application port as a separate passthrough proxy entry with a separate port. If you intend to access the server using different hostnames or IP addresses, configure each of those options separately; in this case, you can use the same port.

External passthrough proxy links that are embedded in a passthrough proxy page may not work. For example, if the bar.company.com page contains a link to foo.company.com and foo.company.com is configured as a host-mode passthrough proxy application, the link to

foo.company.com fails. To avoid this, use port-mode passthrough proxy for passthrough proxy links embedded in passthrough proxy applications.

Related Documentation

- Passthrough Proxy Overview

- Using Virtual Ports

Creating a Custom Header Resource Policy

By default, the rewriting engine only sends selected custom headers to browsers (clients) and backend servers. You can use custom header resource policies, however, to allow or deny custom headers for specific resources.

Note that custom header resource policies do not control standard HTTP headers such as Content-Type.

To write a custom header resource policy:

1. In the admin console, choose **Users > Resource Policies > Web**.

2. If your administrator view is not already configured to show rewriting policies, make the following modifications:

   1. Click the **Customize** button in the upper right corner of the page.

   2. Select the **Rewriting** check box.

   3. Select the **Custom Headers** check box below the Rewriting check box.

   4. Click **OK**.

5. Select the **Rewriting > Custom Headers** tab.

6. On the Custom Header Policies page, click **New Policy**.

7. Enter a name to label this policy (required) and a description of the policy (optional).

8. In the Resources section, specify the resources to which this policy applies.

9. In the Roles section, specify:

   - **Policy applies to ALL roles**—To apply this policy to all users.

   - **Policy applies to SELECTED roles**—To apply this policy only to users who are mapped to roles in the Selected roles list. Make sure to add roles to this list from the Available roles list.

   - **Policy applies to all roles OTHER THAN those selected below**—To apply this policy to all users except for those who map to the roles in the Selected roles list. Make sure to add roles to this list from the Available roles list.

10. In the Action section, specify:

    - **Allow Custom Headers**—Select this option to prevent the system from blocking the headers to browsers (clients) and backend servers.

    - **Deny Custom Headers**—Select this option to use the default custom header behavior on the system. When you select this option, the system blocks custom headers for added security.

    - **Use Detailed Rules**—To specify one or more detailed rules for this policy.

11. Click **Save Changes**.

12. On the Web Rewriting Policies page, order the policies according to how you want to evaluate them. Keep in mind that once the system matches the resource requested by the user to a resource in a policy's (or a detailed rule's) Resource list, it performs the specified action and stops processing policies.

Related Documentation

- Defining Resource Policies: General Options
- Specifying Resources for a Resource Policy
- Writing a Detailed Rule for Resource Policies

Creating an ActiveX Parameter Resource Policy

When the system rewrites a Web page, it does not rewrite the ActiveX controls that are embedded in the Web page. However, you can create resource policies specifying that the system should rewrite the URL and hostname parameters that are passed by the Web page to the Active X controls. To configure these resource policies, you must obtain the following information:

- Class ID—Web pages generally use a class ID to embed an ActiveX control. A class ID is a unique, constant string that uniquely identifies an ActiveX control.

  You can determine what an ActiveX object's class ID is using Internet Explorer 6: Select **Tools > Internet Options**, click **Settings**, and then click **View Objects**. Select the ActiveX object, right-click, and select **Properties**. The ActiveX object's ID is highlighted.

- Language—Web pages can use either static or dynamic HTML (that is, by using JavaScript) to embed an Active X control. When a Web page uses static HTML, the system can rewrite the specified ActiveX parameters on the system itself while it intermediates traffic, since all of the required information passes between the user's browser and the application's Web server. When a Web page uses dynamic HTML to embed an ActiveX control, however, the page frequently pulls information from the client and then generates HTML to embed the ActiveX control. Therefore, the system needs to run script in the user's browser in order to obtain the information it needs to rewrite the specified ActiveX parameters.

- Parameter type—When configuring the system to rewrite a parameter, you must determine whether the parameter is a URL or hostname. The system does not support any other parameter types.

- Parameter name—You must specify the name of the parameter that you want to rewrite. You can find the parameters by searching for the param tag within an object tag. For example, you might find a flash movie embedded in a page using the following code:

```
<object classid="clsid:D27CDB6E-AE6D-11cf-96B8-444553540000" > <param
name="movie" value="mymovie.swf" />
<param name="quality" value="high" />
</object>
```

  When configuring the corresponding resource policy, you should enter movie in the Parameter name field because movie refers to the URL requires rewriting. Frequently, pages contain multiple param tags, but not all of them require rewriting. In this example, the quality parameter does not require rewriting.

To write an ActiveX parameter rewriting resource policy:

1. In the admin console, choose **Users > Resource Policies > Web**.

2. If your administrator view is not already configured to show rewriting policies, make the following modifications:

   1. Click the **Customize** button in the upper right corner of the page.

   2. Select the **Rewriting** check box.

   3. Select the **ActiveX Parameter Rewriting** check box below the Rewriting check box.

   4. Click **OK**.

5. Select the **Rewriting > ActiveX Parameter Rewriting** tab.

6. On the ActiveX Parameter Rewriting Policies page, click **New Policy**.

7. Enter class ID of the ActiveX control that you want to control with the policy (required) and description of the policy (optional).

8. In the Parameters section, specify the ActiveX parameters that you want to control with the policy and the corresponding actions. Possible actions include:

   - **Rewrite URL and response (Static HTML only)**—Rewrite the specified URL parameter on the system. The system also rewrites any response from the Web server requesting the URL. Note that you should select this option if the Web page embeds the ActiveX control using only static HTML.

   - **Rewrite URL and response (Static and dynamic HTML)**—Rewrite the specified URL on the client in addition to rewriting on the system. The system also rewrites any response from the Web server requesting the URL. Note that you should select this option if the Web page embeds the ActiveX control using dynamic HTML.

   - **Rewrite URL (Static HTML only)**—Rewrite the specified URL parameter on the system. Note that you should select this option if the Web page embeds the ActiveX control using only static HTML.

   - **Rewrite URL (Static and dynamic HTML)**—Rewrite the specified URL on the client in addition to rewriting on the system. Note that you should select this option if the Web page embeds the ActiveX control using dynamic HTML.

   - **Rewrite hostname (Static HTML only)**—Rewrite the specified hostname parameter on the system. Note that you should select this option if the Web page embeds the ActiveX control using only static HTML.

   - **Rewrite hostname (Static and dynamic HTML)**—Rewrite the specified hostname on the client in addition to rewriting on the system. Note that you should select this option if the Web page embeds the ActiveX control using dynamic HTML.

   - **Do not rewrite**—Do not rewrite any of the ActiveX component's parameters.

9. Click **Save Changes**.

Related Documentation

- Resource Policies

Restoring the Default ActiveX Resource Policies

The system comes with several predefined resource policies for rewriting the parameters of commonly used ActiveX objects. If you choose to delete any of these policies and then want to restore them later, you can recreate them using the information in Table 89 as a guideline.

Table89: Predefined Resource Policies

Description	Class ID	Parameter	Action
Citrix NFuse xginen_EmbeddedApp object	238f6f83-b8b4-11cf-8771-00a024541ee3	ICAFile	Rewrite URL and response (Static HTML only)
OrgPlus OrgViewer	DCB98BE9-88EE-4AD0-9790-2B169E8D5BBB	URL	Rewrite URL and response (Static HTML only)
Quickplace	05D96F71-87C6-11D3-9BE4-00902742D6E0	GeneralURL General_ServerName	Rewrite URL and response (Static and dynamic HTML) Rewrite hostname (Static and dynamic HTML)
iNotes Discussion	5BDBA960-6534-11D3-97C7-00500422B550	FullURL	Rewrite URL and response (Static and dynamic HTML)
B20D9D6A-0DEC-4d76-9BEF-175896006B4A	B20D9D6A-0DEC-4d76-9BEF-175896006B4A	ServerURL Error URL	Rewrite URL and response (Static and dynamic HTML) Rewrite hostname (Static and dynamic HTML)
Citrix NFuse Elite	2E687AA8-B276-4910-BBFB-4E412F685379	ServerURL	Rewrite URL and response (Static HTML only)
WebPhotos LEAD	00120000-B1BA-11CE-ABC6-F5B2E79D9E3F	BitmapDataPath	Rewrite URL and response (Static and dynamic HTML)
Shockwave Flash	D27CDB6E-AE6D-11cf-96B8-444553540000	Src Movie	Rewrite URL and response (Static and dynamic HTML) Rewrite URL and response (Static and dynamic HTML)
iNotes Blue	3BFFE033-BF43-11d5-A271-00A024A51325	General_URL General_ServerName	Rewrite URL and response (Static and dynamic HTML) Rewrite hostname (Static and dynamic HTML)
Tabular Data Control	333C7BC4-460F-11D0-BC04-0080C7055A83	DataURL	Rewrite URL (Static HTML only)
Windows Media Player	6BF52A52-394A-11D3-B153-00C04F79FAA6	URL	Rewrite URL and response (Static HTML only)
FlowPartPlace	4A266B8B-2BB9-47db-9B0E-6226AF6E46FC	URL	Rewrite URL and response (Static HTML only)
HTML Help	adb880a6-d8ff-11cf-9377-00aa003b7a11	Item1	Rewrite URL and response (Static and dynamic HTML)

MS Media Player	22d6f312-b0f6-11d0-94ab-0080c74c7e95	FileName	Rewrite URL and response (Static HTML only)
CSV Files Handler	333c7bc4-460f-11d0-bc04-0080c7055a83	DataURL	Rewrite URL and response (Static HTML only)
Special ActiveX control for Microsoft OWA	D801B381-B81D-47a7-8EC4-EFC111666AC0	mailboxUrl	Rewrite URL and response (Static HTML only)
FlowPartPlace1	639325C9-76C7-4d6c-9B4A-523BAA5B30A8	Url	Rewrite URL and response (Static HTML only)
scriptx print control	5445be81-b796-11d2-b931-002018654e2e	Path	Rewrite URL and response (Static HTML only)
94F40343-2CFD-42A1-A774-4E7E48217AD4	94F40343-2CFD-42A1-A774-4E7E48217AD4	HomeViewURL	Rewrite URL and response (Static HTML only)
Microsoft License Manager	5220cb21-c88d-11cf-b347-00aa00a28331	LPKPath	Rewrite URL and response (Static HTML only)
Domino 7 beta 2 UploadControl	E008A543-CEFB-4559-912F-C27C2B89F13B	General_URL General_ServerName	Rewrite URL and response (Static and dynamic HTML) Rewrite hostname (Static and dynamic HTML)
iNotes	1E2941E3-8E63-11D4-9D5A-00902742D6E0	General_URL General_ServerName	Rewrite URL and response (Static and dynamic HTML) Rewrite hostname (Static and dynamic HTML)
ActiveCGM	F5D98C43-DB16-11CF-8ECA-0000C0FD59C7	FileName	Rewrite URL and response (Static HTML only)
00130000-B1BA-11CE-ABC6-F5B2E79D9E3F	00130000-B1BA-11CE-ABC6-F5B2E79D9E3F	BitmapDataPath	Rewrite URL and response (Static and dynamic HTML)

Creating Rewriting Filters

Only use the Rewriting Filters tab when instructed to do so by the Pulse Secure Support team.

Writing a Web Compression Resource Policy

The system comes pre-equipped with one Web compression policy (*:*/*) which compresses all applicable Web data. You can enable this policy through the Users > Resource Policies > Web > Compression pages of the admin console.

To write a Web compression resource policy:

1. In the admin console, choose **Users > Resource Policies > Web**.

2. If your administrator view is not already configured to show compression policies, make the following modifications:

    1. Click the **Customize** button in the upper right corner of the page.

    2. Select the **Compression** check box.

    3. Click **OK**.

4. Select the **Compression** tab.

5. On the Web Compression Policies page, click **New Policy**.

6. Enter a name to label this policy (required) and a description of the policy (optional).

7. In the Resources section, specify the URLs to which this policy applies.

8. In the Roles section, specify:

   - **Policy applies to ALL roles**—To apply this policy to all users.

   - **Policy applies to SELECTED roles**—To apply this policy only to users who are mapped to roles in the Selected roles list. Make sure to add roles to this list from the Available roles list.

   - **Policy applies to all roles OTHER THAN those selected below**—To apply this policy to all users except for those who map to the roles in the Selected roles list. Make sure to add roles to this list from the Available roles list.

9. In the Action section, specify:

   - **Compress**—Compress the supported content types from the specified resource.

   - **Do not compress**—Do not compress the supported content types from the specified resource.

   - **Use Detailed Rules**—Select this option to specify one or more detailed rules for this policy.

10. Click **Save Changes**.

Related Documentation

- Specifying Resources for a Resource Policy

- Defining Resource Policies: General Options

- Writing a Detailed Rule for Resource Policies

Defining an OWA Compression Resource Policy

Due to caching issues with OWA, the system comes with the following built-in resource policies specifying that it should not compress Javascript or CSS files that are routed through OWA:

1. Do Not Compress *:*/exchWeb/controls/*.css (all roles)

2. Do Not Compress *:*/exchWeb/controls/*.js (all roles)

3. Do Not Compress *:*/exchWeb/*/controls/*.css (all roles)

4. Do Not Compress *:*/exchWeb/*/controls/*.js (all roles)

In the last two policies, a wildcard (*) is included in the path to account for different OWA build versions.

Pulse Secure recommends that you do not change the compression resource policies for OWA unless absolutely necessary.

Related Documentation

- About OWA and Lotus Notes Caching Resource Policies

Writing a Web Proxy Resource Policy

Web proxy resource policies specify Web proxy servers for which the system should intermediate content. Note that the system intermediates both forward and backwards proxies, but only enables single sign-on to a proxy when you use these tabs to configure the proxy and thereby specify that you trust it.

To write a Web proxy resource policy:

1. In the admin console, choose **Users > Resource Policies > Web**.

2. If your administrator view is not already configured to show Web proxy policies, make the following modifications:

   1. Click the **Customize** button in the upper right corner of the page.

   2. Select the **Web Proxy** check box.

   3. Select the **Policies** check box below the Web Proxy check box.

   4. Click **OK**.

5. Select the **Web Proxy > Policies** tab.

6. On the Web Proxy Policies page, click **New Policy**.

7. Enter a name to label this policy (required) and a description of the policy (optional).

8. In the Resources section, specify the resources to which this policy applies.

9. In the Roles section, specify:

   - **Policy applies to ALL roles**—To apply this policy to all users.

   - **Policy applies to SELECTED roles**—To apply this policy only to users who are mapped to roles in the Selected roles list. Make sure to add roles to this list from the Available roles list.

   - **Policy applies to all roles OTHER THAN those selected below**—To apply this policy to all users except for those who map to the roles in the Selected roles list. Make sure to add roles to this list from the Available roles list.

10. In the Action section, specify:

    - **Access Web resources directly**—Intermediate the user's request to a back-end server and the server's response to the user for requests made to a resource specified in the Resources list.

    - **Access Web resources through a Web proxy**—Specify a Web proxy server in the drop-down list that you have defined in the Users > Resource Policies > Web > Web Proxy > Servers tab.

    - **Use Detailed Rules**—To specify one or more detailed rules for this policy.

11. Click **Save Changes**.

12. On the Web Proxy Policies page, order the policies according to how you want to evaluate them. Keep in mind that once the system matches the resource requested by the user to a resource in a policy's (or a detailed rule's) Resource list, it performs the specified action and stops processing policies.

Related Documentation

Specifying Web Proxy Servers

You can direct all Web requests to a Web proxy rather than using the system to connect directly to Web servers. This feature can be useful if your network security policy requires this configuration or if you want to use a caching Web proxy to improve performance.

To specify servers for Web proxy resource policies:

1. In the admin console, choose **Users > Resource Policies > Web**.

2. If your administrator view is not already configured to show Web proxy policies, make the following modifications:

   1. Click the **Customize** button in the upper right corner of the page.

   2. Select the **Web Proxy** check box.

   3. Select the **Servers** check box below the Web Proxy check box.

   4. Click **OK**.

5. Select the **Web Proxy > Servers** tab.

6. Under Web Proxy Servers, enter the name or IP address of the Web proxy server and the port number at which the proxy server listens, and then click Add.

7. Repeat this step to specify additional Web proxy servers.

Related Documentation

- Writing a Web Proxy Resource Policy

Writing An HTTP 1.1 Protocol Resource Policy

Protocol resource policies enable or disable HTTP 1.1 protocol support between the system and backend servers. The system supports chunked Transfer-Encoding, gzip and deflate Content-Encoding, connection persistence, and caching headers such as If-Modified-Since, If-None-Match, If-Unmodified-Since and If-Match. The system supports range requests with partial content when you select the **Don't rewrite content: Do not redirect to target web server selective rewrite** option.

For a detailed description of the HTTP 1.1 protocol, refer to the Hyptertext Transfer Protocol -- HTTP 1.1 specification from the World Wide Web Consortium.

The system only communicates with network servers using HTTP 1.1 if the client also communicates using HTTP 1.1. If the client uses HTTP 1.0, the system communicates with backend servers using HTTP 1.0, regardless of whether or not HTTP 1.1 is enabled.

If you want to use HTTP 1.1 for a specific resource, enable HTTP 1.1 for that policy and ensure that the new policy appears above the default in the list of configured policies. You should add the HTTP 1.1 policy to the top of the policy list because the policy evaluation engine evaluates policies from top to bottom, stopping when it encounters a match.

The system comes with a default policy that disables HTTP 1.1 for all resources. If you want to use HTTP 1.1 for all resources, either redefine the "*:*/*" policy or create a new policy enabling HTTP 1.1 and move it to the top of your policy list. If you delete this default policy (and any other policies that disable HTTP 1.1), the system uses HTTP 1.0 for all resources

To write an HTTP 1.1 protocol resource policy:

1. In the admin console, choose **Users > Resource Policies > Web**.

2. If your administrator view is not already configured to show protocol policies, make the following modifications:

    1. Click the **Customize** button in the upper right corner of the page.

    2. Select the **Protocol** check box.

    3. Click **OK**.

4. Select the **Protocol** tab.

5. On the Web Protocol Policies page, click **New Policy**.

6. Enter a name to label this policy (required) and a description of the policy (optional).

7. In the Resources section, specify the URLs to which this policy applies.

8. In the Roles section, specify:

    - **Policy applies to ALL roles**—To apply this policy to all users.

    - **Policy applies to SELECTED roles**—To apply this policy only to users who are mapped to roles in the Selected roles list. Make sure to add roles to this list from the Available roles list.

    - **Policy applies to all roles OTHER THAN those selected below**—To apply this policy to all users except for those who map to the roles in the Selected roles list. Make sure to add roles to this list from the Available roles list.

9. In the Action section, specify:

    - **Disable HTTP 1.1**—Automatically communicate with backend servers via the HTTP 1.0 protocol.

    - **Enable HTTP 1.1**—Automatically communicate with backend servers using the HTTP 1.1 protocol as long as the client also communicates using the HTTP 1.1 protocol.

    - **Use Detailed Rules**—Select this option to specify one or more detailed rules for this policy.

10. Click **Save Changes**.

Related Documentation

- Resource Policy Evaluation

- Defining Resource Policies: General Options

- Writing a Detailed Rule for Resource Policies

Creating a Cross Domain Access Policy

The XMLHttpRequest object allows scripts to perform HTTP client functionality, such as submitting form data or loading data from a server. Today's web browsers impose a security

restriction on the use of XMLHttpRequest. You are not allowed to make XMLHttpRequests to any server except the server where your web page came from. For example, if both your web application and the data required for that application come from the same web server, then there is no restriction. But, if your web application is on one server and you make a request to a different server, the browser prevents the connection from opening. It is possible to bypass this security, however.

You can create a resource profile that determines whether or not to impose this restriction and to what level. By default, this restriction is bypassed and cross domain access is allowed.

To create a cross domain access policy:

1. In the admin console, choose **Users > Resource Policies > Web**.

2. If your administrator view is not already configured to show cross-domain policies, make the following modifications:

    1. Click the **Customize** button in the upper right corner of the page.

    2. Select the **Rewriting** check box.

    3. Select the **Cross Domain Access** check box below the Rewriting check box.

    4. Click **OK**.

5. Select the **Rewriting > Cross Domain Access** tab.

6. On the Cross Domain Access page, click **New Policy**.

7. Enter a name to label this policy (required) and a description of the policy (optional).

8. In the Resources section, specify the URLs to which this policy applies.

9. In the Roles section, specify:

    - **Policy applies to ALL roles**—To apply this policy to all users.

    - **Policy applies to SELECTED roles**—To apply this policy only to users who are mapped to roles in the Selected roles list. Make sure to add roles to this list from the Available roles list.

    - **Policy applies to all roles OTHER THAN those selected below**—To apply this policy to all users except for those who map to the roles in the Selected roles list. Make sure to add roles to this list from the Available roles list.

10. In the Action section, specify:

    - **Allow Cross Domain Access**—To not impose any restriction and allow cross domain access.

    - **Deny XMLHttpRequest Cross Domain Access only**—To deny cross domain access if the XMLHttpRequest object is used in the call.

    - **Deny all Cross Domain Access**—To deny cross domain access regardless of whether or not the XMLHttpRequest object is used in the call.

    - **Use Detailed Rules**—To specify one or more detailed rules for this policy.

11. Click **Save Changes**.

Related Documentation

- Defining Resource Policies: General Options

- Writing a Detailed Rule for Resource Policies

Defining Resource Policies: General Options

When you enable the Web resource policy options described in this section, the system compiles a list of hostnames specified in the Resources field of each Web resource policy. The system then applies the enabled options to this comprehensive list of hostnames.

To specify Web resource options:

1. In the admin console, navigate to **Users > Resource Policies > Web**.

2. If your administrator view is not already configured to show Web options, make the following modifications:

   1. Click the **Customize** button in the upper right corner of the page.

   2. Select the **Options** check box.

   3. Click **OK**.

4. Select the **Options** tab.

5. Select **IP based matching for Hostname based policy resources** if you want the system to look up IP address corresponding to each hostname specified in a Web resource policy. When a user tries to access a server by specifying an IP address rather than the hostname, the system compares the IP to its cached list of IP addresses to determine if a hostname matches an IP. If there is a match, then the system accepts the match as a policy match and applies the action specified for the resource policy.

   i NOTE: This option does not apply to hostnames that include wildcards and parameters.

6. Select **Case sensitive matching for the Path and Query string components in Web resources** if you want to require users to enter a case-sensitive URL to a resource. For example, use this option when passing username or password data in a URL.

7. Click **Save Changes**.

Related Documentation

- Resource Policies
- Resource Policy Components

Managing Resource Policies: Customizing UI Views

You can control which Web resource policy configuration pages the system displays so that you only have to view those pages that you actually use. Or, if you have a new system installation, you can use these settings to display additional pages (since the system only displays the most commonly used resource policy pages to new users).

To control which Web resource policy configuration pages to display:

1. In the admin guide, choose **Users > Resource Policies > Web >** *Policy Type.*

2. Click the **Customize View** button in the upper right corner of the console.

3. In the Customize View dialog box, specify which Web resource policies you want to display in the admin console. You may manually select individual check boxes, click **All Pages** to display all Web resource policy configuration pages, or click **Common Pages** to display the most commonly used Web resource policy configuration pages. Note you cannot hide the Web Access Policies page.

4. Click **OK**.

Related Documentation

CHAPTER 23 File Rewriting

File Rewriting Overview

A file resource profile controls access to resources on Windows server shares or UNIX servers.

File rewriting is a standard feature on all Connect Secure devices.
When creating a file resource profile, you must use the following formats when defining a resource policy's primary resource as well as its autopolicy resources.

Windows resources:

\\server[\share[\path]]

UNIX resources:

server[/path]

Within these formats, the three components are:

- Server (required)—Possible values:

  - Hostname—You may use the system variable <username> when defining the hostname.

  - IP address—The IP address needs to be in the format: a.b.c.d

  The leading two back slashes are required for Windows, non-Nfs resources.

- Share (required, Windows only)—The system variable <username> is allowed. Note that when the the system tries to connect to a Windows file share, it connects to ports 445 and 139.

- Path (optional)—Special characters allowed include:

---

*	Matches any character. Note that you cannot use the * wildcard character when defining a resource profile's primary resource (that is, the Server/share field for Windows resources or the Server field for UNIX resources).
%	Matches any character except slash (/)
?	Matches exactly one character

Valid Windows resources include:

```
\\pulsesecure.net\dana
\\10.11.0.10\share\web
\\10.11.254.227\public\test.doc
```

Valid UNIX resources include:

```
\\pulsesecure.net\dana
10.11.0.10/share/web
10.11.254.227/public/test.doc
```

Related Documentation

- Using Custom Expressions in Rule Configuration

Creating a File Rewriting Resource Profile

To create a file rewriting resource profile:

1. In the admin console, choose **Users > Resource Profiles > Files**.

2. Click **New Profile**.

3. From the Type list, select **Windows** or **Unix**.

4. Enter a unique name and optionally a description for the resource profile. (This name becomes the default bookmark's name.)

5. Enter the resource to which you want to control access. Note that the format of the resource varies depending on which type of resource profile you are creating:

   - **Windows**—Enter the server name or IP address, share name, and optionally the path that you want to control access to in the Server/share field. When entering the resource, use the format: \\server[\share[\path]].

- **Unix**—Enter the server name or IP address and optionally the path that you want to control access to in the Server field. When entering the resource, use the format: server[/path]

6. In the **Autopolicy: Windows File Access Control** section or the **Autopolicy: UNIX Access Control** section, create a policy that allows or denies users access to the resource specified the previous step. At minimum, you need to click **Add** in order to use the access control policy that is automatically created for you. This policy allows access to the specified directory and all of its sub-directories.

7. (Optional) Click **Show ALL autopolicy types** to create additional autopolicies that fine-tune access to the resource. Then, create the autopolicies.

8. Click **Save and Continue**.

9. In the Roles tab, select the roles to which the resource profile applies and click Add.

   The selected roles inherit the autopolicies and bookmarks created by the resource profile. If it is not already enabled, the system automatically enables the **Files, Windows** option or the **Files, UNIX/NFS** option in the Users > User Roles > Select Role > General > Overview page of the admin console for all of the roles you select.

10. Click **Save Changes**.

11. (Optional) In the Bookmarks tab, modify the default bookmark and/or create new ones. (By default, the system creates a bookmark to the resource defined in the Windows or UNIX field and displays it to all users assigned to the role specified in the Roles tab.)

Related Documentation

- Defining a Resource Profile
- Creating a File Access Control Autopolicy
- Defining a Single Sign-On Autopolicy

Creating a File Access Control Autopolicy

File access control policies specify resources on your file servers that users may access. When defining a file resource profile, you must create a corresponding access control autopolicy that enables access to the profile's primary resource. The system simplifies the process for you by automatically creating an autopolicy that allows access to the directory specified in the Server/share field (Windows) or the Server field (UNIX) and all of its sub-directories. To enable this autopolicy, you simply need to select it and click Add.

If necessary, you may choose to modify this default autopolicy or create supplementary file access control autopolicies that allow or deny access to additional resources.

To create a new file access control autopolicy:

1. Create a file resource profile.

2. If it is not already enabled, select the **Autopolicy: Windows File Access Control** check box or the **Autopolicy: Unix Access Control** check box.

3. In the Resource field, specify the resource to which this policy applies using the format: \\server[\share[\path]] for Windows resources and \\server[\path] for UNIX resources.

4. From the Action list, select one of the following options:

   - **Allow**—Select this option to enable access to the specified resource.

   - **Read-only**—Select this option to allow users to view but not edit the specified resource.

   - **Deny**—Select this option to block access to the specified resource.

5. Click **Add**.

6. Click **Save Changes**.

Related Documentation

- Creating a File Rewriting Resource Profile

Creating a File Compression Autopolicy

Compression autopolicies specify which types of file data to compress when you enable GZIP compression through the Maintenance > System > Options page of the admin console.

To create a file compression autopolicy:

1. Create a file resource profile.

2. Click **Show ALL autopolicy types**.

3. Select the **Autopolicy: Windows File Compression** check box or the **Autopolicy: Unix File Compression** check box.

4. In the Resource field, specify the resource to which this policy applies using the format: \\server[\share[\path]] for Windows resources and \\server[\path] for UNIX resources.

5. In the Action field, select one of the following options:

   - **Compress**—Select this option to compress data from the specified resource.

   - **Do not compress**—Select this option to disable compression for the specified resource.

6. Click **Add**.

Related Documentation

- About Compression

Creating a Single Sign-On Autopolicy (Windows Only)

Single sign-on (SSO) autopolicies configure the system to automatically submit credentials to a Windows share or directory so that the user does not have to reenter his credentials.

To create a Windows SSO autopolicy:

1. Create a Windows file resource profile.

2. Click **Show ALL autopolicy types**.

3. Select the **Autopolicy: Windows Server Single Sign-On** check box.

4. In the Resource field, specify the resource to which this policy applies using the format: \\server[\share[\path]].

5. Select one of the following options:

- **Use predefined credentials**—Select this option if you want to specify credentials to pass to the Windows share or directory. Then:

  - In the Username field, enter variable (such as <username> or a static username (such as administrator) to submit to the Windows share or directory. When entering a variable, you may also include a domain. For example, yourcompany.net\<username>.

  - Enter a variable (such as <password> in the Variable Password field or enter a static password in the Variable field. Note that the system masks the password you enter here with asterisks.

  When entering static credentials, note that the file browsing server maintains the connections open to a server share, however, so connecting to a different folder on the same share using a different account may not work reliably.

  If the specified credentials fail, the system may submit alternative credentials.

- **Disable SSO**—Select this option if you do not want the system to automatically submit credentials to the specified Windows share or directory.

6. Click **Save Changes**.

Related Documentation

- About Single Sign-On

- File Rewriting Overview

- Multiple Sign-In Credentials Execution

Configuring File Resource Profile Bookmarks

When you create a file resource profile, the system automatically creates a bookmark that links to the primary resource that you specified in the resource profile. The system enables you to modify this bookmark as well as create additional bookmarks within the same domain.

When configuring bookmarks, note that:

- You can only assign bookmarks to roles that you have already associated with the resource profile—not all of the roles defined in the system. To change the list of roles associated with the resource profile, use settings in its Roles tab.

- Bookmarks simply control which links to display to users—not which resources the users can access. For instance, if you enable access to a Windows directory but do not create a bookmark to that directory, users can access the directory through Windows Explorer.

- You cannot create bookmarks that link to additional servers defined through file access control autopolicies.

- If you use a bookmark to reference a file shortcut, note that the system only displays bookmarks with shortcuts to files or folders on a network share such as \\server5\share\users\jdoe\file.txt. However, the system does not display bookmarks with shortcuts to local directories such as C:\users\jdoe\file.txt.

To configure file resource profile bookmarks:

1. If you want to create a resource profile bookmark through the standard resource profiles page:

   1. Navigate to the **Users > Resource Profiles > Files >** *Resource Profile Name* **> Bookmarks** page in the admin console.

   2. Click the appropriate link in the Bookmark column if you want to modify an existing bookmark. Or, click **New Bookmark** to create an additional bookmark.

   Alternatively, if you want to create a resource profile bookmark through the user roles page:

   1. Navigate to the **Users > User Roles >** *Role Name* **> Files > Windows Bookmarks|Unix Bookmarks** page in the admin console.

   2. Click **New Bookmark**.

   3. From the Type list, choose **File Resource Profile**. This option appears only if you have already created a file resource profile.

   4. Select an existing resource profile.

   5. Click **OK**. If you have not already associated the selected role with the resource profile, the system automatically makes the association for you. The system also enables any access control policies for the role that are required by the resource profile.

   6. If this role is not already associated with the selected resource profile, the system displays an informational message. If you see this message, click **Save Changes** to add this role to the resource profile's list of roles and to update the profile's autopolicies as required. Then, repeat the previous steps to create the bookmark.

      > NOTE: When you create a resource profile bookmark through the user roles page (instead of the standard resource profiles page), the system only associates the generated bookmark with the selected role. The system does not assign the bookmark to all of the roles associated with the selected resource profile.

   7. Optionally change the name and description of the bookmark. (By default, the system populates names the bookmark using the resource profile name.)

   8. In the File Browsing Path field, add a suffix to the resource if you want to create links to sub-directories of the resource defined in the primary resource profile.

      Make sure to enter a unique server and path in this field. If you create two bookmarks that contain the same concatenated server and path string, the system deletes one of the bookmarks from the end-user view. You will still be able to see both bookmarks, however, in the administrator console.

   9. In the Appearance section, choose one of the following options:

      - **Appear as bookmark on homepage and in file browsing**—Select this option if you want the bookmark to appear both on a user's welcome page and when browsing network files.

- **Appear in file browsing only**—Select this option if you want the bookmark to appear only when users are browsing network files.

10. If you are configuring the bookmark through the resource profile pages, under Roles, specify the roles to which you want to display the bookmark:

- **ALL selected roles**—Select this option to display the bookmark to all of the roles associated with the resource profile.

- **Subset of selected roles**—Select this option to display the bookmark to a subset of the roles associated with the resource profile. Then select roles from the ALL Selected Roles list and click **Add** to move them to the Subset of selected roles list.

11. Click **Save Changes**.

Related Documentation

- Using Custom Expressions in Rule Configuration

- Creating Windows File Bookmarks

Creating Windows File Bookmarks

You can use two different methods to create Windows file bookmarks:

- Create bookmarks through existing resource profiles (recommended)—When you select this method, the system automatically populates the bookmark with key parameters (such as the primary server and share) using settings from the resource profile. Additionally, while you are creating the associated resource profile, the system guides you through the process of creating any required policies to enable access to the bookmark.

- Create standard bookmarks—When you select this option, you must manually enter all bookmark parameters during configuration. Additionally, you must enable access to the file browsing at the role level and create resource policies that enable access to the servers defined in the bookmark.

You can create Windows bookmarks that appear on the welcome page for users mapped to this role. You can insert the user's username in the URL path to provide quick access to the user's network directories.

When users are browsing files on a DFS server, the DFS server uses the site configuration data stored in Active Directory to return DFS referrals to the system in the right order. Referrals to closer servers are put higher in the list than referrals to servers that are farther away. Clients try referrals in the order in which they are received. If a request comes from a client which resides in a subnet which is not in this list, the server will not know where the client is coming from and will return the list of referrals to the customer in an arbitrary order. This could potentially cause the DFS requests from the system (acting as the client in this case) to access a server much farther away. In turn, this could cause serious delays, especially if the system attempts to access a server which is unreachable from the subnet which the system resides in. If the system is installed on a subnet which is not in the DFS server's list, the DFS administrator may use the "Active Directory Sites and Services" tool on the domain controller to add the system's subnet to the appropriate site.

Related Documentation

- Configuring File Resource Profile Bookmarks

Creating Advanced Bookmarks to Windows Resources

Information in this topic is provided for backwards compatibility. We recommend that you configure access to Windows shares and directories through resource profiles instead, since they provide a simpler, more unified configuration method.

To create a bookmark to a Windows resource:

1. In the admin console, choose **Users > User Roles >** *Role Name* **> Files > Windows Bookmarks**.

2. Click **New Bookmark** and then browse to or enter the server and share name. Specify a path to further restrict access. If you want to insert the user's username, enter <username> at the appropriate place in the path. For information about additional system variables and attributes that you can include in the bookmark. If you specify a name and description for the bookmark, this information displays on the home page instead of the server/share.

   You may not bookmark a Windows server. You must specify both the server and share name.

   Make sure to enter a unique server and path in this field. If you create two bookmarks that contain the same concatenated server and path string, the system deletes one of the bookmarks from the end-user view. You will still be able to see both bookmarks, however, in the administrator console.

3. For Appearance, choose either:

   - **Appear as bookmark on homepage and in file browsing** if you want the bookmark to appear both on a user's welcome page and when browsing network files.

   - **Appear in file browsing only** if you want the bookmark to appear only when browsing network files.

4. For Access, click **Enable auto-allow access to this bookmark** if you want the system to automatically create a corresponding Windows Access resource policy. Note that this functionality applies only to role bookmarks and not bookmarks created by users. Next, select:

   - **Read-write access** to enable users to save files on the server. Note that users cannot upload files greater than 500 MB to the server.

   - **Include sub-folders** to enable users to view files in directories below the specified bookmark path.

     NOTE: You may not see the Auto-allow option if you are using a new installation or if an administrator hides the option.

5. Click **Save Changes** or **Save + New** to add another.

Related Documentation

- Creating a File Rewriting Resource Profile

- Using Custom Expressions in Rule Configuration

Creating Windows Bookmarks that Map to LDAP Servers

To create a bookmark that automatically maps to a user's LDAP home directory:

1. Create an LDAP server instance.

2. Add the LDAP attribute homeDirectory to the Server Catalog.

3. Configure a realm and bind LDAP as the authentication server.

4. Configure role-mapping rules, as needed.

5. Create a Windows bookmark. During configuration, specify <userAttr.homeDirectory> in the bookmark.

6. Click **Save Changes**.

Related Documentation

- Using an LDAP Server

- Creating Windows File Bookmarks

Defining General Windows File Browsing Options

To specify general Windows file browsing options:

1. In the admin console, choose **Users > User Roles >** *Role Name* **> Files > Options**.

2. Under Windows Network Files, specify which options to enable for users:

   - **User can browse network file shares**—If enabled, users can view and create bookmarks to resources on available Windows file shares.

   - **User can add bookmarks**—If enabled, users can view and create bookmarks to resources on available Windows file shares.

3. Click **Save Changes**.

Related Documentation

Writing a File Resource Policy

When you enable the File access feature for a role, you need to create resource policies that specify which Windows and UNIX/NFS resources a user may access, as well as the encoding to use when communicating with Windows and NFS file shares. When a user makes a file request, the system evaluates the resource policies corresponding to the request, such as Windows access resource policies for a request to fetch an MS Word document (.doc file). After matching a user's request to a resource listed in a relevant policy, the system performs the action specified for the resource.

You can create resource policies through the standard interface (as described in this section) or through resource profiles (recommended method).

When writing a File resource policy, you need to supply key information:

- Resources—A resource policy must specify one or more resources to which the policy applies. When writing a File policy, you need to specify File servers or specific shares.

- Roles—A resource policy must specify the roles to which it applies. When a user makes a request, the system determines what policies apply to the role and then evaluates those policies that correspond to the request.

- Actions—Each type of resource policy performs a certain action, which is either to allow or deny a resource or to perform or not perform some function, such as allow a user to write to a directory. You can also write detailed rules that apply more conditions to a user request.

The system engine that evaluates resource policies requires that the resources listed in a policy's Resources list follow a canonical format.

Windows File Resources Canonical Format

Information in this section is provided for backwards compatibility. We recommend that you configure access to Windows file servers through resource profiles instead, since they provide a simpler, more unified configuration method.

When writing a resource policy for a Windows file resource, you need to understand the following canonical format.

\\server[\share[\path]]

The three components are:

- Server (required)—Possible values:

  - Hostname—The system variable <username> may be used.

  - IP address—The IP address needs to be in the format: a.b.c.d

- Share (optional)—If the share is missing, then star (*) is assumed, meaning ALL paths match. The system variable <username> is allowed.

- Path (optional)—Special characters allowed include:

*	Matches any character
%	Matches any character except slash (/)
?	Matches exactly one character

If the path is missing, then slash (/) is assumed, meaning only top-level folders are matched. For example:

```
\\%.danastreet.net\share\<username>\*
\\pulsesecure.net\dana\*
\\10.11.0.10\share\web\*
\\10.11.254.227\public\%.doc
```

Related Documentation

- Writing a Detailed Rule for Resource Policies

- Using Custom Expressions in Rule Configuration

Writing a Windows Access Resource Policy

Information in this topic is provided for backwards compatibility. We recommend that you configure access to Windows file servers through resource profiles instead, since they provide a simpler, more unified configuration method.

To write a Windows access resource policy:

1. In the admin console, choose **Users > Resource Policies > Files > Access > Windows**.

2. On the Windows File Access Policies page, click **New Policy**.

3. Enter a name to label this policy (required) and a description of the policy. (optional)

4. In the Resources section, specify the resources to which this policy applies.

5. In the Roles section, specify:

   - **Policy applies to ALL roles**—To apply this policy to all users.

   - **Policy applies to SELECTED roles**—To apply this policy only to users who are mapped to roles in the Selected roles list. Make sure to add roles to this list from the Available roles list.

   - **Policy applies to all roles OTHER THAN those selected below**—To apply this policy to all users except for those who map to the roles in the Selected roles list. Make sure to add roles to this list from the Available roles list.

6. In the Action section, specify:

   - **Allow access**—To grant access to the resources specified in the Resources list. Check **Read-only** to prevent users from saving files on the server.

   - **Deny access**—To deny access to the resources specified in the Resources list.

   - Use Detailed Rules—To specify one or more detailed rules for this policy.

7. Click **Save Changes**.

8. On the Windows File Access Policies page, order the policies according to how you want to evaluate them. Keep in mind that once the system matches the resource requested by the user to a resource in a policy's (or a detailed rule's) Resource list, it performs the specified action and stops processing policies.

If you want to write a File resource policy that enables you to specify credentials to submit to a file server when a user request matches a resource in the Resource list, you can use the following procedure to do so. You can also configure the system to prompt users for credentials.

Related Documentation

- Creating a File Rewriting Resource Profile

- Writing a Detailed Rule for Resource Policies

Writing a Windows SSO Resource Policy

Information in this topic is provided for backwards compatibility. We recommend that you configure access to Windows file servers through resource profiles instead, since they provide a simpler, more unified configuration method.

To write a Windows credentials resource policy:

1. In the admin console, choose **Users > Resource Policies > Files > SSO > Windows**.

2. On the Windows Credentials Policies page, click **New Policy**.

3. Enter a name to label this policy (required) and a description of the policy. (optional)

4. In the Resources section, specify the resources to which this policy applies.

5. In the Roles section, specify:

   - **Policy applies to ALL roles**—To apply this policy to all users.

   - **Policy applies to SELECTED roles**—To apply this policy only to users who are mapped to roles in the Selected roles list. Make sure to add roles to this list from the Available roles list.

   - **Policy applies to all roles OTHER THAN those selected below**—To apply this policy to all users except for those who map to the roles in the Selected roles list. Make sure to add roles to this list from the Available roles list.

6. In the Action section, specify the action to take when a resource requires credentials:

   - **Use System Credentials**—If the system has stored credentials for the specified user and resource in its cache, it submits the stored credentials. If the stored credentials fail or if no stored credentials exist for that user, the system prompts for new credentials and stores the new credentials.

   - **Use Specific Credentials**—You specify static credentials that the system submits to resources. The file browsing server maintains the connections open to a server\share so connecting to a different folder on the same share using a different account may not work reliably. If the specified credentials fail, the system may submit alternative credentials. Note that the system masks the password you enter here with asterisks.

   - **Prompt for user credentials**—The system intermediates the share challenge by presenting an authentication challenge the first time a user attempts to access the share. The user enters the credentials and the credentials are stored in the system. If the credentials later fail, the system again prompts the user for their credentials.

   - **Use Detailed Rules**—To specify one or more detailed rules for this policy.

7. Click **Save Changes**.

8. On the Windows File Access Policies page, order the policies according to how you want to evaluate them. Keep in mind that once the system matches the resource requested by the user to a resource in a policy's (or a detailed rule's) Resource list, it performs the specified action and stops processing policies.

Related Documentation

- Creating a File Rewriting Resource Profile

- About Single Sign-On

- Writing a Detailed Rule for Resource Policies

Writing a Windows Compression Resource Policy

Information in this section is provided for backwards compatibility. We recommend that you configure compression through resource profiles instead, since they provide a simpler, more unified configuration method.

Compression policies specify which types of file data to compress when you enable GZIP compression through the Maintenance > System > Options page of the admin console.

The system comes pre-equipped with two file compression policies (*:*/*) which compress all applicable file data. You may enable these policies through the Resource Policies > Files > Compression pages of the admin console.

To write a Windows file compression resource policy:

1. In the admin console, choose **Resource Policies > Files > Compression**.

2. Select the **Windows** tab.

3. Click **New Policy**.

4. Enter a name to label this policy (required) and a description of the policy. (optional)

5. In the Resources section, specify the resources to which this policy applies.

6. In the Roles section, specify:

   - **Policy applies to ALL roles**—To apply this policy to all users.

   - **Policy applies to SELECTED roles**—To apply this policy only to users who are mapped to roles in the Selected roles list. Make sure to add roles to this list from the Available roles list.

   - **Policy applies to all roles OTHER THAN those selected below**—To apply this policy to all users except for those who map to the roles in the Selected roles list. Make sure to add roles to this list from the Available roles list.

7. In the Action section, specify:

   - **Compress**—Compress the supported content types from the specified resource.

   - **Do not compress**—Do not compress the supported content types from the specified resource.

   - **Use Detailed Rules**—Select this option to specify one or more detailed rules for this policy.

8. Click **Save Changes**.

Related Documentation

- About Compression

- Creating a File Compression Autopolicy

- Specifying Resources for a Resource Policy

Defining General File Writing Options

You can specify File resource options that apply to your File resource policies. When you enable a File resource policy option, the system compiles a list of hostnames specified in the Resources field of each File resource policy. The system then applies the enabled options to this comprehensive list of hostnames.

To specify resource options for Windows file servers:

1. In the admin console, choose **Users > Resource Policies > Files > Options**.

2. Select:

    - **IP based matching for Hostname based policy resources**—The system looks up the IP address corresponding to each hostname specified in a File resource policy. When a user tries to access a server by specifying an IP address rather than the hostname, the system compares the IP to its cached list of IP addresses to determine if a hostname matches an IP. If there is a match, then the system accepts the match as a policy match and applies the action specified for the resource policy.

        This option does not apply to hostnames that include wildcards and parameters.

    - **Case sensitive matching for the Path component in File resources**—Require users to enter a case-sensitive path component.

    - **Encoding**—Select the encoding to use when communicating with Windows and NFS file shares.

    - **Use NTLM v1, NTLM v1 will be used for all NTLM negotiations**—Select this option to use only NTLM V1 for file share authentication.

    - **Use NTLM v2, NTLM v2 will be used for all NTLM negotiations**—Select this option to use only NTLM V2 for file share authentication.

    - **Number of NTLM authentication protocol variant attempts**—Controls the number of login attempts while doing SSO, Select "Low" if you are seeing account lockout issues.

3. Click **Save Changes**.

Related Documentation

- About Basic, NTLM and Kerberos Resources

- About Single Sign-On

Creating UNIX File Bookmarks

You can use two different methods to create UNIX file bookmarks:

- Create bookmarks through existing resource profiles (recommended)—When you select this method, the system automatically populates the bookmark with key parameters (such as the server) using settings from the resource profile. Additionally, while you are creating the associated resource profile, the system guides you through the process of creating any required policies to enable access to the bookmark.

- Create standard bookmarks—When you select this option, you must manually enter all bookmark parameters during configuration. Additionally, you must enable access to the file

browsing at the role level and create resource policies that enable access to the servers defined in the bookmark.

You can create UNIX bookmarks that appear on the welcome page for users mapped to this role. You can insert the user's username in the URL path to provide quick access to the user's network directories.

Related Documentation

- Creating Windows File Bookmarks

- Creating Advanced Bookmarks to UNIX Resources

Creating Advanced Bookmarks to UNIX Resources

Information in this topic is provided for backwards compatibility. We recommend that you configure access to UNIX servers through resource profiles instead, since they provide a simpler, more unified configuration method.

You can create UNIX/NFS bookmarks that appear on the home page. You can insert the user's username in the URL path to provide quick access to the user's network directories.

To create a bookmark to a UNIX/NFS resource:

1. In the admin console, choose **Users > User Roles >** *Role Name* **> Files > UNIX Bookmarks**.

2. Click **New Bookmark** and then enter the server hostname or IP address and the path to the share. If you want to insert the user's username, enter <username> at the appropriate place in the path. If you specify a name and description for the bookmark, this information displays on the home page instead of the server/path.

   Make sure to enter a unique server and path in this field. If you create two bookmarks that contain the same concatenated server and path string, the system deletes one of the bookmarks from the end-user view. You will still be able to see both bookmarks, however, in the administrator console.

3. For Appearance, choose either:

   - **Appear as bookmark on homepage and in file browsing** if you want the bookmark to appear both on a user's welcome page and when browsing network files.

   - **Appear in file browsing only** if you want the bookmark to appear only when browsing network files.

4. For Access, click **Enable auto-allow access to this bookmark** if you want to automatically create a corresponding UNIX/NFS resource policy. Note that this functionality applies only to role bookmarks and not bookmarks created by users. Next, select:

   - **Read-write access** to enable users to save files on the server. Note that users cannot upload files greater than 500 MB to the server.

   - **Include sub-folders** to enable users to view files in directories below the specified bookmark path.

     i NOTE: You may not see the Auto-allow option if you are using a new installation or if an administrator hides the option.

5. Click **Save Changes** or **Save + New** to add another.

Related Documentation

- Creating a File Rewriting Resource Profile

Defining General UNIX File Browsing Options

For NFS file browsing to work properly, you must configure an NIS server on the system before enabling NFS file browsing.

To specify general file browsing options:

1. In the admin console, choose **Users > User Roles >** *Role Name* **> Files > Options**.

2. Under UNIX Network Files, specify which options to enable for users:

   - **User can browse network file shares**—If enabled, users can view and create bookmarks to resources on available UNIX file shares.

   - **User can add bookmarks**—If enabled, users can view and create bookmarks to resources on available UNIX file shares.

   - **Allow automount shares**—If enabled, users access to automount shares specified on a NIS server.

3. Click **Save Changes**.

Related Documentation

- Defining General Windows File Browsing Options

Defining UNIX/NFS File Resource Policies

When you enable the File access feature for a role, you need to create resource policies that specify which Windows and UNIX/NFS resources a user may access, as well as the encoding to use when communicating with Windows and NFS file shares. When a user makes a file request, the system evaluates the resource policies corresponding to the request, such as Windows access resource policies for a request to fetch an MS Word document (.doc file). After matching a user's request to a resource listed in a relevant policy, the system performs the action specified for the resource.

You can create resource policies through the standard interface (as described in this section) or through resource profiles (recommended method).

When writing a File resource policy, you need to supply key information:

- Resources—A resource policy must specify one or more resources to which the policy applies. When writing a File policy, you need to specify File servers or specific shares.

- Roles—A resource policy must specify the roles to which it applies. When a user makes a request, the system determines what policies apply to the role and then evaluates those policies that correspond to the request.

- Actions—Each type of resource policy performs a certain action, which is either to allow or deny a resource or to perform or not perform some function, such as allow a user to write to a directory. You can also write detailed rules that apply more conditions to a user request.

The engine that evaluates resource policies requires that the resources listed in a policy's Resources list follow a canonical format.

Canonical Format: UNIX/NFS File Resources

When writing a resource policy for a UNIX/NFS file resource, you need to understand the following canonical format.

server[/path]

The two components are:

- Server (required)—Possible values:

  - Hostname—The system variable <username> may be used.

  - IP address—The IP address needs to be in the format: a.b.c.d

- Path (optional)—Special characters allowed include:

*	Matches any character
%	Matches any character except back slash (\)
?	Matches exactly one character

If the path is missing, then back slash (\) is assumed, meaning only top-level folders are matched. For example:

```
%.danastreet.net/share/users/<username>
*.\\pulsesecure.net\dana/*
10.11.0.10/web/*
10.11.254.227/public/%.txt
```

Related Documentation

- Defining General File Writing Options

Writing UNIX/NFS Resource Policies

Information in this section is provided for backwards compatibility. We recommend that you configure access to UNIX file servers through resource profiles instead, since they provide a simpler, more unified configuration method.

To write a UNIX/NFS resource policy:

1. In the admin console, choose **Users > Resource Policies > Files > Access > Unix/NFS**.

2. On the UNIX/NFS File Access Policies page, click **New Policy**.

3. Enter a name to label this policy (required) and a description of the policy. (optional)

4. In the Resources section, specify the resources to which this policy applies.

5. In the Roles section, specify:

   - **Policy applies to ALL roles**—To apply this policy to all users.

- **Policy applies to SELECTED roles**—To apply this policy only to users who are mapped to roles in the Selected roles list. Make sure to add roles to this list from the Available roles list.

- **Policy applies to all roles OTHER THAN those selected below**—To apply this policy to all users except for those who map to the roles in the Selected roles list. Make sure to add roles to this list from the Available roles list.

6. In the Action section, specify:

   - **Allow access**—To grant access to the resources specified in the Resources list. Check Read-only to prevent users from saving files on the server.

   - **Deny access**—To deny access to the resources specified in the Resources list.

   - **Use Detailed Rules**—To specify one or more detailed rules for this policy.

7. Click **Save Changes**.

8. On the UNIX/NFS File Access Policies page, order the policies according to how you want to evaluate them. Keep in mind that once the system matches the resource requested by the user to a resource in a policy's (or a detailed rule's) Resource list, it performs the specified action and stops processing policies.

Related Documentation

- Creating a File Rewriting Resource Profile

- Writing a Detailed Rule for Resource Policies

Writing a UNIX/NFS Compression Resource Policy

Information in this section is provided for backwards compatibility. We recommend that you configure access to UNIX file servers through resource profiles instead, since they provide a simpler, more unified configuration method.

Compression policies specify which types of file data to compress when you enable GZIP compression through the Maintenance > System > Options page of the admin console.

The system comes pre-equipped with two file compression policies (*:*/*) which compress all applicable file data. You may enable these policies through the Resource Policies > Files > Compression pages of the admin console.

To write a UNIX/NFS file compression resource policy:

1. In the admin console, choose **Resource Policies > Files > Compression**.

2. Select the **Unix/NFS** tab.

3. Click **New Policy**.

4. Enter a name to label this policy (required) and a description of the policy. (optional)

5. In the Resources section, specify the resources to which this policy applies.

6. In the Roles section, specify:

   - **Allow access**—To grant access to the resources specified in the Resources list. Check Read-only to prevent users from saving files on the server.

- **Deny access**—To deny access to the resources specified in the Resources list.

- **Use Detailed Rules**—To specify one or more detailed rules for this policy.

7. In the Action section, specify:

- **Compress**—Compress the supported content types from the specified resource.

- **Do not compress**—Do not compress the supported content types from the specified resource.

- **Use Detailed Rules**—Select this option to specify one or more detailed rules for this policy.

8. Click **Save Changes**.

Related Documentation

- Creating a File Rewriting Resource Profile

- Specifying Resources for a Resource Policy

- Writing a Detailed Rule for Resource Policies

Defining General UNIX/NFS File Writing Options

You can specify File resource options that apply to your File resource policies. When you enable a File resource policy option, the system compiles a list of hostnames specified in the Resources field of each File resource policy. The system then applies the enabled options to this comprehensive list of hostnames.

To specify options for UNIX/NFS resources:

1. In the admin console, choose **Users > Resource Policies > Files > Options**.

2. Select:

- **IP based matching for Hostname based policy resources**—The the system looks up the IP address corresponding to each hostname specified in a File resource policy. When a user tries to access a server by specifying an IP address rather than the hostname, the system compares the IP to its cached list of IP addresses to determine if a hostname matches an IP. If there is a match, then the system accepts the match as a policy match and applies the action specified for the resource policy.

  This option does not apply to hostnames that include wildcards and parameters.

- **Case sensitive matching for the Path component in File resources**—Select this option to require users to enter a case-sensitive URL to an NFS resource. Use this option when passing username or password data in a URL.

   NOTE: This option does not apply to Windows servers.

- **Encoding**—Select the encoding to use for communicating with the Windows and NFS file shares.

- **NTLM Version**—Select whether to fall back to NTLM version 1 or version 2 authentication if Kerberos authentication of administrator credentials fails.

- **Number of NTLM authentication protocol**—Select **High** to allow a large number of authentication attempt to be made to the backend server. This applies only to NTLM, not basic authentication. If your server locks users out for too many failed attempts, select **Low**.

   NOTE: Many servers do not support the different NTLM protocol variant attempts when you select High. If you find that authentication is failing even though the username and password are correct, set this option to Low.

3. Click **Save Changes**.

Related Documentation

- About Basic, NTLM and Kerberos Resources

CHAPTER 24 Secure Application Manager

Secure Application Manager Overview

The Secure Application Manager option provides secure, application-level remote access to enterprise servers from client applications. You may deploy two versions of the Secure Application Manager:

- Windows version (WSAM)—The Windows version of the Secure Application Manager is a Windows-based solution that enables you to secure traffic to individual client/server applications and application servers.

- Java version (JSAM)—The Java version of the Secure Application Manager provides support for static TCP port client/server applications, including enhanced support for Microsoft MAPI, Lotus Notes, and Citrix NFuse. JSAM also provides NetBIOS support, which enables users to map drives to specified protected resources.

Related Documentation

- Task Summary: Configuring WSAM

- Task Summary: Configuring JSAM

Task Summary: Configuring WSAM

This section provides high-level WSAM configuration steps. These steps do not account for preliminary system configuration steps such as specifying the system's network identity or adding user IDs.

To configure WSAM:

1. Create resource profiles that enable access to client/server applications or destination networks, create supporting autopolicies as necessary, and assign the policies to user roles using settings in the Users > Resource Profiles> SAM pages of the admin console.

   We recommend that you use resource profiles to configure WSAM (as described above). However, if you do not want to use resource profiles, you can configure WSAM using role and resource policy settings in the following pages of the admin console instead:

   1. Enable access to WSAM at the role-level using settings in the Users > User Roles > Role > General > Overview page of the admin console.

   2. Specify which client/server applications and servers WSAM should intermediate using settings in the Users > User Roles > SAM > Applications page of the admin console.

   3. Specify which application servers users can access through WSAM using settings in the Users > Resource Policies > SAM > Access page of the admin console.

4. After enabling access to client/server applications and/or destination networks using WSAM resource profiles or roles and resource policies, you can modify general role and resource options in the following pages of the admin console:

   1. (Optional) Configure role-level options such as whether the system should automatically launch and upgrade WSAM using settings in the Users > User Roles > SAM > Options page of the admin console.

   2. (Optional) Control IP based hostname matching at the resource level using settings in the Users > Resource Policies > SAM > Options page of the admin console.

3. Ensure that an appropriate version of WSAM is available to remote clients using settings in the Maintenance > System > Installers page of the admin console.

4. If you want to enable or disable client-side logging for WSAM, configure the appropriate options through the System > Configuration > Security > Client-side Logs tab of the admin console.

Related Documentation

- WSAM Recommended Operation

- About WSAM Resource Profiles

- Specifying Resource Level WSAM Options

- User Roles Overview

- Specifying Applications and Servers for WSAM to Secure

- Downloading Client Installer Files

- Enabling Client-Side Logging

WSAM Recommended Operation

Pulse Secure recommends the following operation when using WSAM:

- W-SAM supports client-initiated UDP and TCP traffic by process name, by destination hostname, or by destination address range:port range. Except for Passive FTP, W-SAM only supports protocols that do not embed IP addresses in the header or payload. W-SAM also supports unicast client-initiated UDP.

- Users must launch drive maps through W-SAM in one of the following ways:

  - NetUse--At the Command prompt, type **net use * \\server\share /user:username**.

  - Right-click My Computer > Map Network Drive, or in Windows Explorer, go to Tools > Map Network Drive and select **Connect using a different username**.

- When using the W-SAM Access Control List (ACL), administrators should take extra precaution when granting access to hosts. We recommend that administrators use the IP address instead of the hostname. If the hostname is required for security purposes, administrators should try to include additional ACLs with the corresponding IP address or IP addresses for that hostname. Reverse DNS lookups are not supported.

- To run Citrix NFuse through W-SAM, you must define a Caching rule to cache launch.asp files. For example, configure the resource policy to *server name*:**80,443/*.launch.asp**and the Caching Option to **Cache (do not add/modify caching headers)**.

- When using WSAM on Pocket PC, session roaming should be enabled when being used over GPRS because the IP address of the phone may change.

- When using WSAM on Pocket PC, if you have multiple roles defined, select the **Merge settings for all assigned roles** option under Administrators > Admin Realms > *Realm* > Role Mapping.

- When using an external load balancer and accessing J-SAM, W-SAM, Network Connect, or the Online Meeting functionality, persistence must be employed on the load balancer. This persistence should be based on Source IP or Destination Source, depending on the load balancer being used.

Related Documentation

- Task Summary: Configuring WSAM

Launching VPN Tunneling During a WSAM Session

Users can launch VPN Tunneling while signed in to the system via WSAM. If they do, however, the VPN Tunneling installer automatically terminates the WSAM session prior to launching VPN Tunneling. During the process, the system prompts users with a warning message informing them that they are about to terminate their WSAM session in favor of launching VPN Tunneling.

To deal with situation, we recommend that you give users as much access to network resources through VPN Tunneling as through WSAM. If you do, when the users choose to launch VPN Tunneling (simultaneously terminating WSAM), they will still be able to access the same network resources.

Related Documentation

- Launching VPN Tunneling During a Windows Secure Application Manager Session

Debugging WSAM Issues

You can use the Secure Application Manager dialog box on an end-user's system to view the WSAM status and a variety of details about the user's session. For instance, the Secure Application Manager dialog box displays the applications and servers that WSAM is configured to secure, event logs and Winsock data for the user's session, and various system diagnostics and performance data. This information can help you or a Pulse Secure Support representative debug any problems your users may encounter.

To access the Secure Application Manager dialog box, users simply need to double-click the WSAM icon on their Windows task bars:

For more information about viewing information in the Secure Application Manager dialog box, see the end-user help system available from the Help link in the end-user console.

Related Documentation

- Task Summary: Configuring WSAM

About WSAM Resource Profiles

You can create two types of WSAM resource profiles:

- WSAM application resource profiles—These resource profiles configure WSAM to secure traffic to a client/server application. When you create a WSAM application resource profile, the WSAM client intercepts requests from the specified client applications to servers in your internal network.

- WSAM destination network resource profiles—These resource profiles configure WSAM to secure traffic to a server. When you create a WSAM destination network resource profile, the WSAM client intercepts requests from processes running on the client that are connecting to the specified internal hosts.

When creating WSAM resource profiles, note that the resource profiles do not contain bookmarks. To access the applications and servers that WSAM intermediates, users must first launch WSAM and then launch the specified application or server using standard methods (such as the Windows Start menu or a desktop icon).

When you enable JSAM or WSAM through Web rewriting autopolicies in the Users > Resource Profiles > Web Applications/Pages page of the admin console, the system automatically creates JSAM or WSAM autopolicies for you. You can only view these SAM policies through

the appropriate Web resource profile—not through the SAM resource profile pages of the admin console.

Related Documentation

- Resource Profiles
- Task Summary: Configuring WSAM
- Defining a Rewriting Autopolicy
- Enabling WSAM on PDAs

Creating WSAM Client Application Resource Profiles

When you create a WSAM application resource profile, the WSAM client intercepts requests from the specified client applications to servers in your internal network.

To create a WSAM application resource profile:

1. In the admin console, choose **Users > Resource Profiles > SAM > Client Applications**.

2. Click **New Profile**.

3. From the Type list, choose **WSAM**.

4. From the Application list, select one of the following options:

   - **Custom**—When you select this option, you must manually enter your custom application's executable file name (such as telnet.exe). Additionally, you may specify this file's path and MD5 hash of the executable file (although it is not required that you specify the exact path to the executable). If you enter an MD5 hash value, WSAM verifies that the checksum value of the executable matches this value. If the values do not match, WSAM notifies the user that the identity of the application could not be verified and does not forward connections from the application to the system.

   - **Lotus Notes**—When you select this option, WSAM intermediates traffic from the Lotus Notes fat client application.

   - **Microsoft Outlook**—When you select this option, WSAM intermediates traffic from the Microsoft Outlook application.

   - NetBIOS file browsing—When you select this option, WSAM intercepts NetBIOS name lookups in the TDI drivers on port 137.

   - **Citrix**—When you select this option, WSAM intermediates traffic from Citrix applications.

     You can only use WSAM to configure access to a standard application once per user role. For example, you can enable one configuration of Microsoft Outlook and one configuration of Lotus Notes for the "Users" role.

     The system supports several mechanisms for intermediating traffic to the Lotus Notes, Microsoft Outlook, and Citrix applications.

   - **Domain Authentication**—Select this option to allow integrated Windows applications, such as file sharing, Outlook, and so forth to authenticate to the domain controller when the client machine is part of a domain. Before using this option, you must:

     - Specify domain controllers that are reachable through the system in the WSAM Destination list so that LDAP and Kerberos traffic can be proxied and sent to the system.

     - Configure a WSAM Access Control Policy (ACL) to allow access to all domain controllers.

5. Enter a unique name and optionally a description for the resource profile. The system displays this information in the Client Application Sessions section of the end-user home page.

6. In the **Autopolicy: SAM Access Control** section, create a policy that allows or denies users access to the server that hosts the specified application:

   1. If it is not already enabled, select the **Autopolicy: SAM Access Control** check box.

   2. In the Resource field, specify the application server to which this policy applies. You can specify the server as a hostname or an IP/netmask pair. You may also include a port.

      If you select Domain Authentication from the Application list, enter your domain controller server addresses into the Resource field. You can add multiple domain controller servers if more than one is available.

      When enabling auto-policy for any client application for WSAM, avoid entering *.* in the resource list since the access control policies are not restricted to that particular application. This may result in other resources being accessed through client applications for which the access control policies are not defined.

   3. From the Action list, select **Allow** to enable access to the specified server or **Deny** to block access to the specified server.

   4. Click **Add**.

5. Click **Save and Continue**.

6. In the Roles tab, select the roles to which the resource profile applies and click Add.

   The selected roles inherit the autopolicy created by the resource profile. If it is not already enabled, the system also automatically enables the SAM option in the Users > User Roles > Select Role > General > Overview page of the admin console for all of the roles you select.

7. Click **Save Changes**.

Related Documentation

- Creating Resource Profiles Using the Lotus iNotes Template

- Creating Resource Profiles Using the Microsoft OWA Template

- Comparing Access Mechanisms for Configuring Citrix

Creating WSAM Destination Network Resource Profiles

When you create a WSAM destination network resource profile, the WSAM client intercepts requests from processes running on the client to internal hosts.

🛈   NOTE: When destinations (using either IP address or hostnames) are configured on the system, all DNS and NetBIOS names are resolved through the system.

To create a WSAM destination network resource profile:

1. In the admin console, choose **Users > Resource Profiles > SAM > WSAM Destinations**.

2. Click **New Profile**.

3. Enter a unique name and optionally a description for the resource profile.

4. In the WSAM Destinations section, specify which servers you want to secure using WSAM and click **Add**. You can specify the servers as hostname or IP/netmask pairs. You may also include a port.

5. Select the Create an access control policy allowing SAM access to this server check box to enable access to the server specified in the previous step (enabled by default).

6. Click **Save and Continue**.

7. In the Roles tab, select the roles to which the resource profile applies and click Add.

   The selected roles inherit the autopolicy created by the resource profile. If it is not already enabled, the system also automatically enables the SAM option in the Users > User Roles > *Role Name* > General > Overview page of the admin console for all of the roles you select.

Related Documentation

- Using Custom Expressions in Rule Configuration

Specifying Applications and Servers for WSAM to Secure

Information in this section is provided for backwards compatibility. We recommend that you secure traffic using WSAM resource profiles instead, since they provide a simpler, more unified configuration method.

Use the Applications tab to specify applications and servers for which WSAM secures traffic. When WSAM downloads to a client PC, it contains the information you configure on the Applications tab for the role. After a user launches the Secure Application Manager, WSAM intercepts requests from client applications to servers in your internal network and requests from processes running on the client to internal hosts. You define these resources on the Applications tab by configuring two lists:

- WSAM supported applications list—This list contains applications for which you want WSAM to secure client/server traffic between the client and the system.

- WSAM allowed servers list—This list contains hosts for which you want WSAM to secure client/server traffic between the client and the system.

To specify applications for which WSAM secures client/server traffic between the client and the system:

1. In the admin console, choose **Users > User Roles >** *Select Role***> SAM > Application**s.

2. Click **Add Application**.

3. Enter the name of the application and, optionally, a description. This information displays in the Client Application Sessions section of the end-user home page.

4. From the Type list, choose one of the following options:

   - **Standard**—If you select this option, choose one the following applications from the Application Parameters section:

     - **Citrix**—When you select this option, WSAM intermediates traffic from Citrix applications.

     - **Lotus Notes**—When you select this option, WSAM intermediates traffic from the Lotus Notes fat client application.

     - **Microsoft Outlook/Exchange**—When you select this option, WSAM intermediates traffic from the Microsoft Outlook application.

   - **NetBIOS file browsing**—When you select this option, WSAM intercepts NetBIOS name lookups in the TDI drivers on port 137.

     Note that in order to access a share using WSAM with NetBIOS, you need to explicitly specify the server's NetBIOS name (alphanumeric string up to 15 characters) in two

places: on the Add Server page and in a SAM resource policy. (Wildcards are currently not supported.) Alternatively, you can enable the Auto-allow application servers option on the SAM > Options tab, and then the system automatically creates a SAM resource policy that allows access to this server.

- **Custom**—Select this option to specify a custom client/server application. Then:

  1. In the Filename field, specify the name of the file's executable file.

  2. Optionally specify the file's path and MD5 hash of the executable file. If you enter an MD5 hash value, WSAM verifies that the checksum value of the executable matches this value. If the values do not match, WSAM notifies the user that the identity of the application could not be verified and does not forward connections from the application to the system.

3. Click **Save Changes** or **Save + New**.

4. Configure a WSAM resource policy to specify to which enterprise resources (based on IP address/port combination) the system may send the application.

Specifying Servers for WSAM to Secure

To specify servers for which WSAM secures client/server traffic between the client and the system:

1. In the admin console, choose **Users > User Roles >** *Select Role* **> SAM > Applications**.

2. Click **Add Server**.

3. Enter the name of the server and, optionally, a description.

4. Specify the server's hostname (the wild cards '*' or '?' are accepted) or an IP/netmask pair. Specify multiple ports for a host as separate entries.

5. Click **Save Changes** or **Save + New**.

6. Configure a WSAM resource policy to specify to which enterprise resources (based on IP address/port combination) the system may send a server request.

   Alternatively, you can enable the Auto-allow application servers option on the SAM > Options tab, and then the system automatically creates a SAM resource policy that allows access to the specified server. Note that you need to enable this option before specifying the application or server; otherwise, you need to create a SAM resource policy.

Related Documentation

- Creating Resource Profiles Using the Lotus iNotes Template

- Creating Resource Profiles Using the Microsoft OWA Template

- Comparing Access Mechanisms for Configuring Citrix

- Using Custom Expressions in Rule Configuration

Specifying Applications that Need to Bypass WSAM

The WSAM client comes pre-configured with a list of "passthrough" applications bypass WSAM. The WSAM client does not secure traffic for these applications. In addition to bypassing these predefined applications, you may also specify additional applications that should bypass WSAM.

> NOTE: WSAM does not bypass applications on Pocket PCs and other handheld devices.

To specify applications for WSAM to secure:

1. In the admin console, choose **Users > User Roles >** *Select Role* **> SAM > Applications**.

2. Click **Add Bypass Application**. The New Bypass Application page displays.

3. Name the application and provide a description (optional).

4. Provide the file name (required).

5. Enter the absolute path to the application (optional).

6. Select **Save Changes** to add the bypass application to the list or **Save + New** to save the bypass application and create another bypass application.

Default Bypass Applications

The WSAM client is preconfigured to bypass WSAM processing for the following applications:

- apache.exe

- apache*

- licadmin.exe

- vni.exe

- lmgrd.exe

- TNSLSNR.EXE

- ORACLE.EXE

- Agntsrvc.exe

- ONRSD.EXE

- Pagntsrv.exe

- ENCSVC.EXE

- Agntsvc.exe

- EiSQLW.exe

- Sqlservr.exe

- Sqlmangr.exe

- inetinfo.EXE

- xstart.exe

- idsd.exe

- dsTermServ.exe

- dsCitrixProxy.exe

- dsNcService.exe

- dsNetworkConnect.exe

Related Documentation

- Specifying Applications and Servers for WSAM to Secure

Specifying Role-Level WSAM Options

To specify WSAM options at the role level:

1. In the admin console, choose **Users > User Roles >** *Select Role* **> SAM > Options**.

2. If it is not already enabled, select the Windows SAM option at the top of the page.

3. Under Secure Application Manager options, configure the following options:

   - **Auto-launch Secure Application Manager**—If you enable this option, the system automatically launches the Secure Application Manager when a user signs in. If you do not select this option, users must manually start the Secure Application Manager from the Client Applications Sessions section of the end-user home page.

     Although you configure the Secure Application Manager to automatically launch when users sign into the device, users can override this setting through the Preferences > Applications page of the end-user console. If you or the end user disables WSAM from automatically launching, users need to manually start the Secure Application Manager by clicking its link on the home page.

   - **Auto-allow application servers**—If you enable this option, the system automatically creates a SAM resource policy that allows access to the server specified in the WSAM application and server lists.

     You may not see the Auto-allow option if you are using a new installation or if an administrator hides the option.

4. Under Windows SAM Options, configure the following options:

   - **Auto-uninstall Secure Application Manager**—Select this option to automatically uninstall the Secure Application Manager after users sign off.

   - **Prompt for username and password for intranet sites**—Select this option to require users to enter their sign-in credentials before connecting to sites on your internal network. This option changes Internet Explorer's intranet zone setting so that Internet Explorer prompts the user for network sign-in credentials whenever the user wants to access an intranet site.

   - **Auto-upgrade Secure Application Manager**—Select this option to automatically download the Secure Application Manager to a client machine when the version of Secure Application Manager on the system is newer than the version installed on the client. If you select this option, note the following:

     - The user must have Administrator privileges in order for the system to automatically install Secure Application Manager on the client.

     - If a user uninstalls Secure Application Manager and then signs in to the system for which the Auto-upgrade Secure Application Manager option is not enabled, the user no longer has access to Secure Application Manager.

   - **Resolve only hostnames with domain suffixes in the device DNS domains**—If this option is configured, WSAM filters DNS requests (FQDNs) and sends to the system only those DNS requests that have a domain suffix in the list of DNS Domains configured on the Network Overview page. This option is limited to resolution of FQDNs only. No filtering is applied to short names and NetBIOS requests.

   - **Session start script and Session end script**—If you want to run a batch, application, or Win32 service file when the WSAM session starts or ends, enter the name and path for the file. For example, if you want to terminate an application and then restart it, you may use PSKILL.exe (an third-party utility that terminates processes on local or remote systems).

     If you enable the Session start script option or Session end script option, note the following:

- You must either install the specified file on your end-user's computers or specify a path on an accessible network directory.

- To ensure that the system can locate a file on different platforms, you can use Windows variables, such as in a path such as %WINDIR%\system32\log.

- The file must invoke the WSAM launcher using the appropriate command-line options.

5. Click **Save Changes**.

Related Documentation

- Using the WSAM Launcher

Specifying Application Servers that Users Can Access

Information in this section is provided for backwards compatibility. We recommend that you secure traffic using WSAM resource profiles instead, since they provide a simpler, more unified configuration method.

When you enable the Secure Application Manager access feature for a role, you need to create resource policies that specify which application servers a user may access. These policies apply to both the Java version and Windows version of the Secure Application Manager (JSAM and WSAM, respectively). When a user makes a request to an application server, the system evaluates the SAM resource policies. If the system matches a user's request to a resource listed in a SAM policy, the system performs the action specified for the resource.

When writing a SAM resource policy, you need to supply key information:

- Resources—A resource policy must specify one or more resources to which the policy applies. When writing a SAM policy, you need to specify application servers to which a user may connect.

- Roles—A resource policy must specify the roles to which it applies. When a user makes a request, the system determines what policies apply to the role and then evaluates those policies that correspond to the request. SAM resource policies apply to users requests made through either version, JSAM or WSAM.

- Actions—A Secure Application Manager resource policy either allows or denies access to an application server.

You can create resource policies through the standard interface (as described in this section) or through resource profiles (recommended method).

The engine that evaluates resource policies requires that the resources listed in a policy's Resources list follow a canonical format.

To write a Secure Application Manager resource policy:

1. In the admin console, choose **Users > Resource Policies > SAM > Access**.

2. On the Secure Application Manager Policies page, click **New Policy**.

3. Enter a name to label this policy (required) and a description of the policy (optional).

4. In the Resources section, specify the application servers to which this policy applies.

5. In the Roles section, specify:

   - **Policy applies to ALL roles**—Choose this option to apply this policy to all users.

- **Policy applies to SELECTED roles**—Choose this option to apply this policy only to users who are mapped to roles in the Selected roles list. Make sure to add roles to this list from the Available roles list.

- **Policy applies to all roles OTHER THAN those selected below**—Choose this option to apply this policy to all users except for those who map to the roles in the Selected roles list. Make sure to add roles to this list from the Available roles list.

6. In the Action section, specify:

    - **Allow socket access**—Choose this option to grant access to the application servers specified in the Resources list.

    - **Deny socket access**—Choose this option to deny access to the application servers specified in the Resources list.

    - Use Detailed Rules—Choose this option to specify one or more detailed rules for this policy.

7. Click **Save Changes**.

8. On the Secure Application Manager Policies page, order the policies according to how you want to evaluate them. Keep in mind that once the system matches the resource requested by the user to a resource in a policy's (or a detailed rule's) Resource list, it performs the specified action and stops processing policies.

Related Documentation

- About WSAM Resource Profiles

- Specifying Resources for a Resource Policy

Specifying Resource Level WSAM Options

Use the Options tab to specify the SAM resource option to match IP addresses to hostnames specified as resources in your SAM resource policies. When you enable this option, the system looks up IP addresses corresponding to each hostname specified in a SAM resource policy. When a user tries to access a server by specifying an IP address rather than the hostname, the system compares the IP to its cached list of IP addresses to determine if a hostname matches an IP. If there is a match, then the system accepts the match as a policy match and applies the action specified for the resource policy.

When you enable this option, the system compiles a list of hostnames specified in the Resources field of each SAM resource policy. The system then applies the option to this comprehensive list of hostnames.

NOTE: This option does not apply to hostnames that include wildcards and parameters.

To specify the SAM resource option:

1. In the admin console, choose **Users > Resource Policies > SAM > Options**.

2. Select **IP based matching for Hostname based policy resources**. This option looks up the IP address corresponding to each hostname specified in a Secure Application Manager resource policy. When a user tries to access a server by specifying an IP address rather than the hostname, the system compares the IP to its cached list of IP addresses to determine if a hostname matches an IP. If there is a match, then the system accepts the match as a policy match and applies the action specified for the resource policy.

3. Click **Save Changes**.

Related Documentation

-

Using the WSAM Launcher

The WSAM launcher (samlauncher.exe) is a tool that signs a user into a device and then downloads and launches WSAM. The launcher provides a command-line interface that a script or application can call. For example, you can write an application that calls the WSAM executable when needed.

To use the WSAM launcher, you need to:

1. Write a script, batch file, service, or application that calls the WSAM launcher using command line arguments. You need to distribute this file to each client PC that requires it.

2. Download the WSAM launcher from Maintenance > System > Installers page of the admin console and then distribute it to your users.

Use the command-line arguments in Table 90 to invoke the WSAM launcher.

Table90: WSAM Command Line Arguments

Argument	Action
-start	Initiates the WSAM connection.
-stop	Terminates the WSAM connection.
-signout	Terminates the WSAM connection and system user session.
-version	Displays WSAM version information and then exits.
-help	Displays available arguments.
-noupgrade	Cancels automatic upgrade of WSAM software.
-reboot	Automatically reboots if prompted by an upgrade. If reboot flag is not set, WSAM exits and does not reboot during an upgrade. Be sure to set the reboot flag if WSAM is operating automatically on a remote PC.
-u *username*	Specifies the username.
-p *password*	Specifies the password for authentication.
-loginscript file	Specifies the location and name of the script file to run when WSAM launches. This command takes precedence over a script file specified on the Users > User Roles > *Role Name* > SAM > Options page.
-postscript file	Specifies the location and name of the script file to run when WSAM exits. This command takes precedence over a script file specified on the Users > User Roles > *Role Name* > SAM > Options page.
-c *certificateName*	Specifies the certificate submitted by the user for authentication. Note that the user can only use this option if a valid SSL certificate is installed on the device. If the system uses a self-signed certificate, the user must import that certificate into his browser.
-u *URL*	Specifies the sign-in URL for the device.
-r *realm*	Specifies the realm to which the system submits the user's credentials.
-verbose	Prompts users for input through dialog boxes.

ⓘ NOTE: samlauncher does not support secondary authentication.

Table 91 lists the possible codes the WSAM launcher returns when it exits.

Table91: WSAM Launcher Exit Codes

Code	Description
0	Success
1	Invalid argument
2	Could Not Connect.
3	Invalid Credentials
4	Role Not Specified (credentials map to multiple roles)
5	Pre-authentication Error (Host Checker or Cache Cleaner did not load)
6	Installation Failed
7	Reboot Required (if '-reboot' not specified)
8	Unable to perform a required software upgrade
10	Feature not supported
12	Failed to authenticate the client certificate
100	Unable to stop the Secure Application Manager
101	Unable to start the Secure Application Manager due to a software conflict caused by another Layered Service Provider

Running Scripts Manually

Users may manually specify scripts to run when a WSAM session begins or ends using the following command-line arguments.

If you specify scripts to run through the Users > User Roles > *Role Name* > SAM > Options page of the admin console, the configured script does not run if a user manually invokes WSAM using the launcher and specifies a different script.

To manually launch a script after a WSAM session begins:

- At a command prompt, enter -loginscript file followed by a system variable or script file name and location.

To manually launch a script after a WSAM session ends:

- At a command prompt, enter -postscript file followed by a system variable and the script file name and location.

  Place system variables, file paths, and file names in quotes. Precede and append system variables with a percent sign (%)

For example:

```
-loginscript file "%program files:%\Internet Explorer\IEXPLORER.EXE"
```

Running Scripts Automatically

You may automatically run a script when WSAM starts or stops by entering the script path and name in the Session start script field or Session end script field on the Users > User Roles > *Role Name* > SAM > Options page of the admin console. This section includes an example batch file that you can automatically launch.

The following example shows how to use the WSAM launcher to invoke WSAM. This sample batch file generates error messages when WSAM launches:

```
SamLauncher –start –url %1 –user %2 –password %3 –realm %4
if errorlevel 1 goto error_invalid_args
if errorlevel 2 goto error_connect
if errorlevel 3 goto error_credentials
if errorlevel 4 goto error_role
if errorlevel 5 goto error_preauth
if errorlevel 6 goto error_install
if errorlevel 7 goto error_reboot


:error_invalid_args
@echo invalid arguments
goto done

:error_connect
@echo could not connect
goto done

:error_credentials
@echo invalid credentials
goto done

:error_role
@echo invalid role
goto done

:error_preauth
@echo pre auth version checking
goto done

:error_install
@echo install failed
goto done

:error_reboot
@echo reboot required
goto done

:error_success
@echo Secure Application Manager has started
goto done

:done
```

Win32 API example

```
CHAR szCmd = "SamLauncher.exe –stop";
DWORD dwExitCode = 0;
STARTUPINFO si;
PROCESS_INFORMATION pi;
ZeroMemory(&si, sizeof(si));
si.cb = sizeof(si);
ZeroMemory(π, sizeof(pi));
if (!CreateProcess(NULL, szCmd, NULL, NULL, FALSE,
    0, NULL, NULL, &si, π)) {
    printf( "CreateProcess(%s) failed %d", szCmd, GetLastError());
    return -1;
}
WaitForSingleObject(pi.hProcess, 20000);
GetExitCodeProcess(&pi.hProcess,&dwExitCode);
```

```
CloseHandle(pi.hProcess);
CloseHandle(pi.hThread);
printf("SamLauncher return %d\n", dwExitCode);
return 0;
```

If you are using Windows Vista, open the command window as an administrator user. Standard output from the SamLauncher.exe does not display if the command window is opened by a user without administrator privileges.

Related Documentation

- Task Summary: Configuring WSAM

JSAM Overview

The Java version of the Secure Application Manager provides support for static TCP port client/server applications, including enhanced support for Microsoft MAPI, Lotus Notes, and Citrix NFuse. JSAM also provides NetBIOS support, which enables users to map drives to specified protected resources.

JSAM works well in many network configurations but does not support dynamic port TCP-based client/server applications, server-initiated connections, or UDP traffic.

> **ⓘ** NOTE: regedit.exe is required for some JSAM functionality. If regedit.exe is disabled, automatic host mapping and the NetBIOS and Outlook/Exchange applications will not work properly.

For information about the operating systems, Web browsers, and JVMs on which Pulse Secure supports JSAM, see the *Supported Platforms Guide* on the Pulse Secure Global Support Center (PSGSC) Center.

Related Documentation

- Task Summary: Configuring JSAM
- JSAM Linux and Macintosh Support

Task Summary: Configuring JSAM

This topic provides high-level JSAM configuration steps. These steps do not account for preliminary system configuration steps such as specifying the system's network identity or adding user IDs.

To configure JSAM:

1. Create resource profiles that enable access to client/server applications, create supporting autopolicies as necessary, and assign the policies to user roles using settings in the Users > Resource Profiles> SAM pages of the admin console.

   We recommend that you use resource profiles to configure JSAM (as described above). However, if you do not want to use resource profiles, you can configure JSAM using role and resource policy settings in the following pages of the admin console instead:

   1. Enable access to JSAM at the role-level using settings in the Users > User Roles > Select Role > General > Overview page of the admin console.

   2. Specify which client/server applications JSAM should intermediate using settings in the Users > User Roles > SAM > Applications page of the admin console.

3. Specify which application servers users can access through JSAM using settings in the Users > Resource Policies > SAM > Access page of the admin console.

4. After enabling access to client/server applications using JSAM resource profiles or roles and resource policies, you can modify general role and resource options in the following pages of the admin console:

   1. (Optional) Configure role-level options such as whether the system should automatically launch JSAM using settings in the Users > User Roles > SAM > Options page of the admin console.

   2. (Optional) Control IP based hostname matching at the resource level using settings in the Users > Resource Policies > SAM > Access page of the admin console.

3. If you want to enable or disable client-side logging for JSAM, configure the appropriate options through the System > Configuration > Security > Client-side Logs tab of the admin console.

4. If you have multiple internal domains, such as company-a.com and company-b.com, add DNS domains to the system using settings in the System > Network > Overview page of the admin console so that names such as app1.company-a.com and app2.company-b.com resolve correctly.

5. If a remote user's PC is set up to use a Web proxy in Internet Explorer, configure the client machine to bypass the proxy server when the user launches applications that need to connect to the Secure Application Manager.

6. Enable JSAM to associate IP loopback addresses with application servers on specific ports either by enabling JSAM to edit the hosts file on your users' systems or by creating an external DNS to route client application traffic to the JSAM applet.

Related Documentation

- Using JSAM for Client/Server Communications
- JSAM Linux and Macintosh Support
- Creating a JSAM Application Resource Profile
- Specifying Applications for JSAM to Secure
- Specifying Role Level JSAM Options
- Automatically Launching JSAM
- Specifying Resource Level JSAM Options
- User Roles Overview
- Specifying Application Servers that Users Can Access
- Configuring a PC that Connects Through a Proxy Web Server
- Configuring External DNS Servers and User Machines

Using JSAM for Client/Server Communications

JSAM provides secure port forwarding by directing client application traffic to the JSAM applet running on a client machine. To the client application running on the local machine, JSAM appears as the application server. To the application server in your network, the system appears as the client application.

Figure 107 illustrates the interaction between a client application and its server via Connect Secure. (This figure assumes that the user specified a localhost IP address as the server in the client application.)

Figure107: Java Secure Application Manager
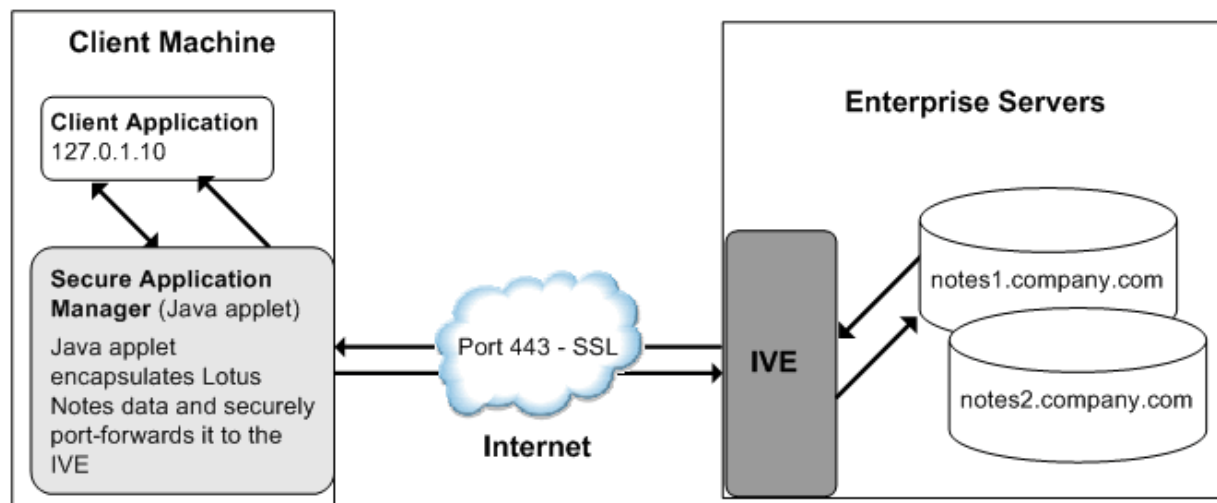
1. The user starts a client application listed in the Client Application Sessions section of the end-user home page. The application resolves the remote server to localhost.

2. The client application connects to JSAM running on the user's machine and starts sending requests.

3. JSAM encapsulates and forwards all client requests to the system over SSL.

4. The system unencapsulates the client data and forwards it to the specified application server.

5. The application server responds with data to the system.

6. The system encapsulates and forwards the response from the application server to JSAM over SSL.

7. JSAM unencapsulates the application server data and forwards it to the client application.

A status indicator on the JSAM window shows the current state of JSAM. If green, JSAM is working correctly. If red, JSAM is unable to send/receive requests to/from the system.

The JSAM window updates the status indicator only when traffic is passed through JSAM. If no traffic is passed through JSAM, the status indicator remains in its current state. For example, if there is a network outage or if the user's session times out, the status indicator remains green even though it cannot send/receive requests to/from the system.

Note the following:

- If a remote user's PC is set up to use a Web proxy in Internet Explorer, you must configure the client machine to bypass the proxy server when the user launches applications that need to connect to the Secure Application Manager.

- JSAM allocates 20-30 MB of RAM when running (the exact amount of memory depends on the Java Virtual Machine (JVM) used) and, if caching is enabled, may leave a .jar file on the client machine. For more information about files left by JSAM on client machines, see the Client-side Changes Guide on the Pulse Secure Global Support Center (PSGSC) Center.

- Users may experience problems waiting for the Secure Application Manager to fully load if they enable pop-up blockers through their Web browsers. This problem occurs because a pop-up window alerting users to accept the Secure Application Manager plug-in may appear in the background (behind the Web browser window) where users cannot see it.

- When launching applications through JSAM, Pulse Secure supports configuration of 1200 unique IP/port combinations on Windows and Mac and 800 unique IP/port combinations on Linux. Note that this limit is based on IP/port combinations, not applications (which may listen on more than one IP address and port). Pulse Secure determined these numbers by testing on Windows machines using default JRE memory settings.

Assigning IP Loopback Addresses to Servers

For JSAM to function, it must listen on loopback addresses for client requests to network application servers. The system assigns these unique IP loopback address to each application server that you specify for a given port. For example, if you specify:

app1.mycompany.com, app2.mycompany.com. app3.mycompany.com,...

for a single port, the system assigns a unique IP loopback address to each application:

127.0.1.10, 127.0.1.11, 127.0.1.12,...

When the system installs JSAM on a user's machine, JSAM listens on the loopback addresses (on the corresponding client port specified for the application server) for client requests to network application servers. You can configure the system to dynamically assign these loopback addresses, or you can configure static loopback addresses yourself through the admin console.

You must enable these associations between IP loopback addresses and applications servers on a specific port in one of two ways:

- Allow the system to edit the hosts file on the client system with IP loopback assignments. The system makes a copy of the current hosts file and then creates a new hosts file with the IP loopback assignments. When the user ends the session, the system deletes the new hosts file and restores the original hosts file.

  If the client system shuts down unexpectedly, the hosts file still points the client to loopback addresses for outside connections. Settings in the hosts file are returned to their original state when the client system reboots.

  Users must have the proper privileges on their machines in order for the system to edit the hosts file.

- Create an external DNS to route client application traffic to the JSAM applet.

Using Static Loopback Addresses

Using an external DNS server with dynamic loopback addresses requires an administrator to update the DNS settings each time the JSAM application configuration changes. On the other hand, configuring an external DNS server using static loopback addresses provides administrators with the highest degree of configuration control. For example, consider the following IP loopback assignments:

app1.mycompany.com - 127.0.1.10
app2.mycompany.com - 127.0.1.11
app3.mycompany.com - 127.0.1.12

If you configure an external DNS server using dynamic loopback address assignments and you delete the first application server, the address assignments change:

app2.mycompany.com - 127.0.1.10
app3.mycompany.com - 127.0.1.11

With static IP loopback addresses in an external DNS, deleting the first application server does not affect the IP loopback assignments for the remaining application servers:

app2.mycompany.com - 127.0.1.11
app3.mycompany.com - 127.0.1.12

You can assign static IP loopback addresses when creating a JSAM custom resource profile through the Users > Resource Profiles > SAM > Client Applications page of the admin console or when enabling JSAM applications through the Users > User Roles > Select Role > SAM > Applications page of the admin console.

If you assign a static IP loopback address while creating a new application, the system checks the address for conflicts against other configured applications in the same role. If another application uses the same address, the system displays an error message prompting you to enter a different IP address.

> NOTE: Static IP loopback addresses apply only to application servers configured by an administrator. The system assigns dynamic IP loopback addresses for user-defined application servers. If the administrator does not assign an IP loopback address to an application server, the system assigns a dynamic address.

IP Loopback Address Considerations When Merging Roles

IP Loopback Address Considerations When Merging Roles

- If two or more roles map to the same application and each mapping contains a different static IP loopback address, all of the static IP loopback addresses remain unchanged.

- If two or more roles map to the same application and only one role uses a static IP loopback address, JSAM uses only the static IP loopback address and binds to only one statically defined socket on the client.

- If two or more roles map to the same application using dynamic IP loopback addresses, only one dynamic IP loopback address is used. The application listener binds to only one dynamically assigned socket on the client.

- If you use the same hostname in multiple roles, either use the same static IP loopback address, or dynamic addresses for all the applications.

- If you use different hostnames associated with the same loopback address and port combination, JSAM cannot distinguish between the two different hosts at the back-end and, hence, cannot accurately direct IP traffic bound for those hosts.

Resolving Hostnames to Localhost

For JSAM to successfully intermediate traffic, a client application on the user's machine needs to resolve the application server to the client localhost. This process enables JSAM to capture and securely port forward the data intended for the application server via Connect Secure. JSAM can perform automatic host-mapping, in which it edits the client's hosts file, to map application servers to localhost. (You can enable automatic host-mapping through the Users > User Roles > Select Role > SAM > Options page of the admin console.)

In order for JSAM to edit a user's hosts file, the user must have the appropriate authority on the client machine:

- Windows users using the FAT file system may belong to any user group. For Exchange MAPI support, however, users must have at least Power User privileges on their machines.

- Windows users using the NTFS file system must have Administrator privileges on their machines.

- Linux (RedHat) users must launch the browser that will launch JSAM as root.

- Macintosh users must supply the Administrator password when prompted by JSAM.

If users do not have the appropriate privileges on their machines, JSAM cannot automatically edit the hosts file, preventing hostname resolution to localhost.

Alternatives for users who do not have the appropriate privileges are:

- You configure your external DNS server to resolve application servers to localhost. If you configure your external DNS server to use a localhost address instead of the application server hostname, remote users need to configure the order in which their machine searches DNS servers to start with the corporate DNS.

- You relax the permissions on the etc directory and the etc\hosts file to enable JSAM to make the necessary modifications.

- Users configure a client application to use the localhost address assigned by the system where they typically specify the application server hostname in the client application.

Related Documentation

- Task Summary: Configuring JSAM

Configuring a PC that Connects Through a Proxy Web Server

If a remote user's PC is set up to use a Web proxy in Internet Explorer, you must configure the client machine to bypass the proxy server and contact the Secure Application Manager instead.

To configure a PC that connects to the system through a Web proxy in Internet Explorer:

1. From the Internet Explorer Tools menu, choose **Internet Options**.

2. On the Connections tab, click the **LAN Settings** button.

3. Under Proxy server, click the **Advanced** button.

4. Under Exceptions, enter the addresses for which you do not want to use a proxy server. Enter all addresses (hostnames and localhost) that the client application uses when connecting through the Secure Application Manager. For example:

   If your application server is app1.company.com, enter the following exceptions:

   app1;app1.company.com;127.0.0.1

   If your Exchange Server is exchange.company.com, enter the following exceptions:

   exchange;exchange.company.com;127.0.0.1

   NOTE: Connect Secure clients parse Internet Explorer's static proxy exception list. We support most exceptions that Internet Explorer supports with the following limitations:

- For IP address exception, we support n.*.*.*, n.n.*.*, n.n.n.*. For example, 10.*.*.*, 10.10.*.*, 10.10.10.*, or 10.10.10.10. We do not support 10* or 10.*.10.* even though Internet Explorer may support them.

- For string expression, we support specific strings such as my.company.net,or a wild card at front of the string, for example, *.my.company.net or *.company.net. We do not support *.company.*, *.company*, *.company.*.com, *.net, *.com and so forth.

Determining the Assigned Loopback Address

Users cannot modify the corporate DNS server for applications they add for port forwarding. If you allow users to specify applications for JSAM to proxy, users need to configure a client application to use the localhost address assigned by the system where they typically enter the server hostname.

The Details pane of the JSAM browser window displays the loopback IP address assigned by the system along with the port specified by the user. To determine what IP address the system assigns to an application specified through the Client Applications page, a user must restart the Secure Application Manager after adding the application. The loopback address assigned to the application appears on the Details pane of the Secure Application Manager browser window.

In the client application, the user needs to enter the system-assigned loopback address as the application server. For example, if a user wants to access a telnet server behind your corporate firewall, the user needs to follow these steps:

1. In the Client Application Sessions section of the end-user home page, click the **Item Properties** icon, then click **Add Application**

2. On the Add Application page, specify:

   - The server's fully qualified domain name or IP address in the Remote Server field, such as terminalserver.pulsesecure.net.

   - The port on which JSAM should listen for client traffic to the server in the Client Port field, such as 3389.

   - The port on which the remote server should listen for traffic from the client application (JSAM) in the Server Port field, such as 3389.

3. Click **Add** to save the information.

4. Close the Secure Application Manager browser window.

5. In the Client Application Sessions section of the end-user home page, click **Start** to restart the Secure Application Manager.

6. In the Secure Application Manager browser window, click Details.

7. On the Details tab, look at which loopback address is assigned to the remote server, such as 127.0.1.18.

8. In the client application, such as Remote Desktop Connection, specify the loopback address in the configuration field for the server. This field appears in different places for different applications. Users may enter this information through a setup wizard or other configuration dialog.

Related Documentation

- Task Summary: Configuring JSAM

Configuring External DNS Servers and User Machines

Client applications must resolve server hostnames to JSAM, which proxies data between a client and a server. On Windows PCs, server hostnames are stored in the hosts file. To intercept data using JSAM, the server names in the hosts file need to resolve to the local machine (localhost) so that the system can intermediate the traffic. The recommended process for mapping application servers to a user's local PC is to enable the automatic host-mapping option, which enables the system to automatically modify the PC hosts file to point application servers to the localhost for secure port forwarding.

For the system to perform automatic host-mapping, however, PC users must have the proper privileges on their machines. If your PC users do not have these privileges, you must ensure that your internal application server names resolve externally to a PC's localhost by adding entries to your external Internet-facing DNS server such as:

127.0.0.1 app1.company-a.com
127.0.0.1 app2.company-b.com
127.0.0.1 exchange1.company-a.com
127.0.0.1 exchange1.company-b.com

If the client application uses an unqualified name for the application server, users need to specify DNS suffixes so that the PC can attach the suffix and contact your external DNS server for name resolution. For example, an MS Outlook client typically has an unqualified name for an MS Exchange server. In order for the qualified name to resolve to 127.0.0.1, users need to specify the appropriate DNS suffixes on their PCs. Adding domain names does not affect other operations on the PC, including use of the client application from within the enterprise.

To configure a user PC with DNS suffixes (Windows 2000):

1. From the Windows Start menu, choose **Settings > Network and Dial-up Connections > Local Area Connection** and then choose **Properties**.

2. Select **Internet Protocol (TCP/IP)** and then click **Properties**.

3. Click **Advanced** and then click the **DNS** tab.

4. Click **Append these DNS suffixes** and then click **Add**.

5. Add your enterprise's internal domains as additional DNS suffixes.

Related Documentation

- JSAM Overview

JSAM Linux and Macintosh Support

Linux users do not have access to ports below 1024 unless they are signed into their machines as root. Macintosh users do not have access to ports below 1024 unless they supply the Administrator password when prompted by JSAM. To support applications that run on privileged ports (ports below 1024), such as a telnet application:

- Users may launch the browser that will launch JSAM as root.

- You or the user may specify a client port number equal to or greater than port 1024 when enabling client applications.

  For example, if you specify 2041 for the client port and 23 for the server port for a telnet application, the command to run the application is:

telnet loopbackIP 2041

where loopbackIP is the loopback IP address assigned to the application server by the system. JSAM listens on port 2041 for traffic from the telnet application and forwards it to the system. The system then forwards the traffic to port 23 on the destination server.

> **NOTE:** Due to the design of the Sun JVM code, Macintosh users cannot relaunch JSAM within the same Safari user session. In order to re-launch JSAM, the user must exit Safari and then launch JSAM again.

Related Documentation

- Task Summary: Configuring JSAM

Standard Application Support: MS Outlook

Remote users can use the Microsoft Outlook client on their PCs to access e-mail, their calendars, and other Outlook features through the system. Versions of MS Outlook currently supported are MS Outlook 2000 and MS Outlook 2002. This ability does not require changes to the Outlook client and does not require a network layer connection, such as VPN.

Refer to the Supported Platforms Document on the Pulse Secure Global Support Center (PSGSC) Center for details on operating system support and dependencies. See *Pulse Connect Secure Client-Side Changes Installation Reference* for details about registry changes made by JSAM.

Also, note that the system does not support Outlook through SVW, since Outlook applications require HKLM registry key changes.

In order for this feature to work for remote users, the network settings of the user's PC must resolve the name of the Exchange Servers embedded in the Outlook client to the local PC (127.0.0.1, the default localhost IP address). We recommend that you configure the system to automatically resolve Exchange server hostnames to the localhost by temporarily updating the hosts file on a client computer through the automatic host-mapping option.

Client/Server Communication Using JSAM

Figure 108 describes the interactions between the Outlook client and an Exchange Server via Connect Secure. This figure assumes that the system is configured to perform automatic host-mapping.

Figure108: Java Secure Application Manager and Enhanced MS Exchange Support
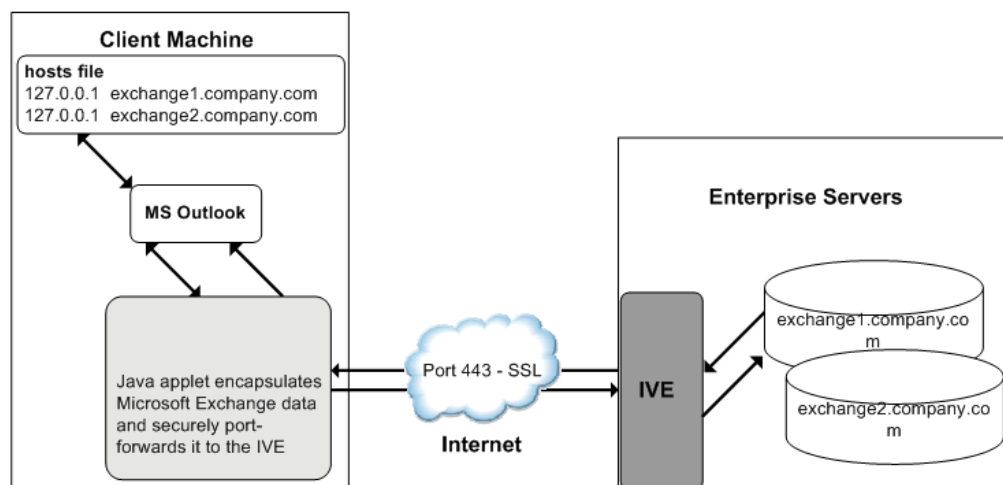
Figure 108 shows the system configured to use automatic host-mapping for the MS Outlook client.

1. The user starts the MS Outlook client. Outlook tries to contact the Exchange Server exchange1.yourcompany.com. The system resolves the Exchange Server hostname to 127.0.0.1 (localhost) through temporary changes to the hosts file.

2. Outlook connects to the Secure Application Manager running on the user's PC and then starts sending requests for e-mail.

3. The Secure Application Manager encapsulates and forwards all the requests from the Outlook client to the system over SSL.

4. The system unencapsulates the client data and looks in the MAPI request to find the target Exchange Server. The request is then forwarded to the target server.

5. Each request in the MAPI protocol encodes the target server for the request. When MAPI requests arrive from the Secure Application Manager, the system looks in each of them and dispatches them to the appropriate target server. This process works transparently even if there are multiple Exchange Servers.

6. The Exchange Server responds to the system with e-mail data.

7. The system encapsulates and forwards the response from the Exchange Server to the Secure Application Manager over SSL.

8. The Secure Application Manager unencapsulates the information sent from the system and forwards the normal MAPI response from the Exchange Server to the Outlook client.

Related Documentation

- Standard Application Support: Lotus Notes

Standard Application Support: Lotus Notes

Remote users can use the Lotus Notes client on their PCs to access e-mail, their calendars, and other features through the system. This ability does not require a network layer connection, such as a VPN.

See the *Supported Platforms Document* for details on operating system support and dependencies.

Client/Server Communication Using JSAM

In order for this feature to work for remote users, they need to configure the Lotus Notes client to use "localhost" as their location setting (that is, their Home Location, Remote Location, or Travel Location setting). The Secure Application Manager then picks up connections requested by the Lotus Notes client. Figure 109 describes the interactions between the Lotus Notes client and a Lotus Notes Server via Connect Secure.

Figure109: Java Secure Application Manager and Enhanced Lotus Notes Support
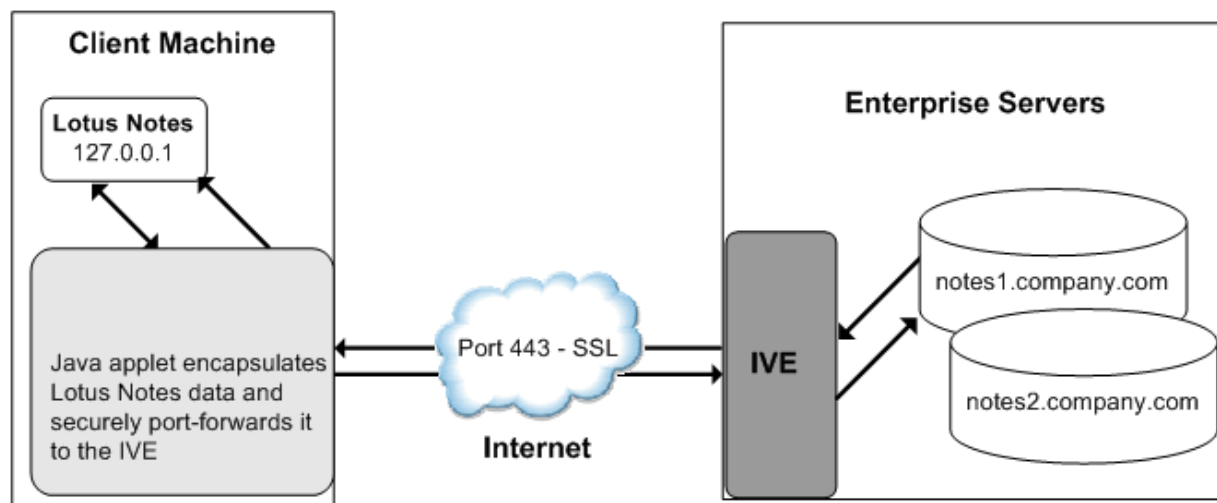
Figure 109 shows the Lotus Notes client location value to be configured to the localhost.

1. The user starts the Lotus Notes client with the location setting. The client uses the HTTP Tunnel proxy setting for its location setting. Note that you must set the HTTP Tunnel proxy setting to use localhost (or 127.0.0.1) as the proxy address and 1352 as the proxy port.

2. The Lotus Notes client connects to the Secure Application Manager and starts sending requests for e-mail.

3. The Secure Application Manager encapsulates and forwards requests from the Lotus Notes client to the system over SSL.

4. The system unencapsulates the client data and looks in the Lotus Notes request to find the target Lotus Notes Server. The request is then forwarded to the target server.

   Each request in the Lotus Notes protocol encodes the target server for the request. When Lotus Notes requests arrive from the application proxy, the system obtains the target server information from the requests and dispatches the requests to the appropriate target server. Thus, this feature works transparently even if there are multiple Lotus Notes Servers accessed by a single user. Note that you must create JSAM ACLs on the system that enable access to these target servers.

5. The Lotus Notes Server responds with e-mail data to the system.

6. The system encapsulates and forwards the response from the Lotus Notes Server to the Secure Application Manager over SSL.

7. The Secure Application Manager unencapsulates the information sent from the system and forwards the normal response from the Lotus Notes Server to the Lotus Notes client.

Related Documentation

- Task Summary: Configuring JSAM

Configuring the Lotus Notes Client

Before a remote user can connect from Lotus Notes to a Lotus Notes Server through the system, the user must edit the Lotus Notes client to set a Location document Proxy field to the PC's localhost port. The Location document edited should be the one used for remote access, such as the Remote Location or Travel Location setting. Setting the Proxy field to the PC's

localhost port enables the system to connect to multiple Lotus Notes Servers, including those set up as pass-through servers.

You should use the following configuration in these cases:

- JSAM is configured to use Lotus Notes as a standard application.

- The Lotus Notes client can connect to multiple Lotus Notes servers.

To configure a Lotus Notes client for use with the system:

1. From the Lotus Notes client, choose **File > Mobile > Locations**.

2. Select the Location used for remote access and then click **Edit Location**.

3. In the Basics tab, click the **Proxy** icon.

4. In the Proxy Server Configuration box, enter **127.0.0.1:1352** in the HTTP Tunnel field.

5. Click **OK**.

Related Documentation

- Standard Application Support: Lotus Notes

Standard Application Support: Citrix Web Interface for MetaFrame (NFuse Classic)

Remote users can use the Citrix Web Interface for MetaFrame server to access a variety of applications via Connect Secure. This process does not require any alterations to the user permissions on the client.

After a user browses to a Citrix Web Interface for MetaFrame server and selects an application, the server sends an ICA file to the client. When the system rewrites the ICA file, it replaces hostnames and IP addresses with pre-provisioned loopback IP addresses. The ICA client then sends application requests to one of the loopback IP addresses. The Secure Application Manager encapsulates the data and sends it to the system. The system unencapsulates the data and sends it to the appropriate MetaFrame server using port 1494 or 2598 (depending on the client)

Note the following:

- The system supports several mechanisms for intermediating traffic between a Citrix server and client, including the Terminal Services, JSAM, WSAM, VPN Tunneling, and hosted Java applets features.

- JSAM does not automatically launch when Embedded Applications are set to "Auto" in the Citrix Web Interface for MetaFrame console. In these cases, we recommend that you configure JSAM to automatically launch after the user signs into the device. Otherwise, end users must manually launch JSAM before using Citrix Web Interface for MetaFrame.

- If a user attempts to use the server discovery feature and then attempts to use application discovery, the application discovery process fails. To resolve this particular situation, shut down and restart Citrix Program neighborhood.

- The system serves as an alternative to deploying the Citrix Secure Gateway (CSG).

- To use the applet-mode of the Java client, make sure to enable Java applet support on the Users > User Roles > *Role Name* > Web > Options page of the admin console.

- If you set the Network Protocol setting in the Citrix Program Neighborhood client to TCP/IP, the system does not support the application through JSAM since the TCP/IP setting produces UDP traffic.

Related Documentation

- Comparing Access Mechanisms for Configuring Citrix
- Enabling Citrix Published Applications on the Citrix Native Client

Enabling Citrix Published Applications on the Citrix Native Client

When enabling Citrix published applications on the Citrix native client through the system, you must complete the following steps:

1. Specify custom application on JSAM to port forward.

2. Configure the Citrix metaframe server for published applications.

3. Configure the Citrix client for published applications.

Note the following:

- These instructions assume that you are not using the Citrix Web Interface for Citrix Presentation Server (formerly known as Nfuse server).

- These instructions do not cover how to configure the standard Citrix application option. (For standard Citrix application instructions, use settings in the Users > Resource Profiles > Web > Web Applications/Pages page of the admin console.) You can enable both the standard Citrix application and the custom Citrix application—these settings do not impact each other.

- The system supports several mechanisms for intermediating traffic between a Citrix server and client, including the Terminal Services, JSAM, WSAM, VPN Tunneling, and hosted Java applets features.

Specifying Custom Applications on JSAM to Port Forward

When configuring JSAM to work with published applications, you must open two port—ports 80 and 1494. Each opened port creates a connection through JSAM to the Citrix Metaframe server.

To specify published applications for JSAM to port forward:

1. Add a custom application through JSAM. When adding the custom application, keep the following settings in mind:

   - Server name—For published applications, you must enter the Metaframe server's fully qualified DNS name, not its IP address.

   - Server port—For published applications, enter 80 and 1494. (Create one entry for port 80 and another for port 1494.) If you have multiple Metaframe servers, you must configure all of them on the same ports.

   - Client port—For published applications, enter 80 and 1494. (Create one entry for port 80 and another for port 1494.)

2. If you have multiple internal domains, such as company-a.com and company-b.com, add DNS domains to the system using settings in the System > Network > Overview page of the admin console so that names such as app1.company-a.com and app2.company-b.com resolve correctly.

3. If a remote user's PC is set up to use a Web proxy in Internet Explorer, configure the client machine to bypass the proxy server when the user launches applications that need to connect to the Secure Application Manager.

4. Enable JSAM to associate IP loopback addresses with application servers on specific ports either by enabling JSAM to edit the hosts file on your users' systems or by creating an external DNS to route client application traffic to the JSAM applet.

5.

Configuring the Citrix Metaframe Server for Published Applications

When enabling Citrix published applications through the system, you must enable the XML service DNS address resolution on the metaframe server. The following instructions describe how to do this on Metaframe XP.

To configure the Citrix metaframe server to work with the system:

1. Open the Citrix Management Console.

2. Right-click on the name of your server farm and click **Properties**.

3. Select the **MetaFrame Settings** tab.

4. Select the **Enable XML Service DNS address resolution** check box.

5. Click **OK**.

Configuring the Citrix Client for Published Applications

When enabling Citrix published applications through the system, you must create an ICA connection on each Citrix client using the instructions that follow.

To configure the Citrix client to work with the system:

1. Open the Citrix Program Neighborhood and choose the **Add ICA Connection** option.

2. In the Add New ICA Connection wizard, select the connection type that your computer uses to communicate.

3. In the next screen:

   1. Enter a description of the new ICA Connection.

   2. Select **TCP/IP + HTTP** as the network protocol.

   3. Select **Published Application**.

   4. Click Server Location, and then:

      1. Deselect the **Use Default** check box.

      2. Click **Add** in the Locate Server or Published Application dialog box.

      3. Confirm that HTTP/HTTPS is selected from the Network Protocol list.

      4. Enter the metaframe server DNS in the Add Server Location Address dialog box.

      5. Enter **80** in the port field.

      6. Click **OK** in the Add Server Location Address dialog box and the Locate Server or Published Application dialog box.

   7. Select an application from the Published Application list.

8. Enter information in the remaining wizard screens as prompted.

Related Documentation

- Configuring a PC that Connects Through a Proxy Web Server

Enabling Citrix Secure Gateways

When enabling Citrix secure gateways (CSGs) through the system, you must:

1. Disable Citrix NFuse as a standard application through the Users > Resource Profiles > Web > Web Applications/Pages page of the admin console.

    > *i* NOTE: You cannot enable the Citrix NFuse standard application and Citrix Secure Gateways (CSGs) custom applications through JSAM on the same device.
    >
    > The system supports several mechanisms for intermediating traffic between a Citrix server and client, including the Terminal Services, JSAM, WSAM, VPN Tunneling, and hosted Java applets features.

2. Specify applications for JSAM to port forward by adding a custom application through JSAM. When adding the custom application, keep the following settings in mind:

    - Server name—For CSGs, you must enter the Citrix secure gateway server's fully qualified DNS name, not its IP address.

    - Server port—For CSGs, enter 443. If you have multiple Citrix secure gateway servers, you must configure all of them on the same port.

    - Client port—For CSGs, enter 443. (Create one entry for port 80 and another for port 443.)

3. If you have multiple internal domains, such as company-a.com and company-b.com, add DNS domains to the system using settings in the System > Network > Overview page of the admin console so that names such as app1.company-a.com and app2.company-b.com resolve correctly.

4. If a remote user's PC is set up to use a Web proxy in Internet Explorer, configure the client machine to bypass the proxy server when the user launches applications that need to connect to the Secure Application Manager.

5. Enable JSAM to associate IP loopback addresses with application servers on specific ports either by enabling JSAM to edit the hosts file on your users' systems or by creating an external DNS to route client application traffic to the JSAM applet.

6. Setup your Citrix Secure Gateway and confirm that it works on your desktop.

7. Add a bookmark to the end-users' home page that points to the list of Citrix secure gateway servers and use the Selective Rewrite feature to turn off rewriting for the URL.

    Or, if you do not want to create a bookmark through the system, simply instruct users to access the URL using their Web browser's address bar instead of the system address bar.

Related Documentation

- Configuring a PC that Connects Through a Proxy Web Server

Creating a JSAM Application Resource Profile

JSAM resource profiles configure JSAM to secure traffic to a client/server application. When you create a JSAM application resource profile, the JSAM client tunnels network traffic generated by the specified client applications to servers in your internal network.

When creating JSAM resource profiles, note that the resource profiles do not contain bookmarks. Therefore, end users will not see a link for the configured application in the end-user interface. To access the applications and servers that JSAM intermediates, users must first launch JSAM and then launch the specified application using standard methods (such as the Windows Start menu or a desktop icon).

Also note that when you enable JSAM or WSAM through rewriting autopolicies for Web resource profiles, the system automatically creates JSAM or WSAM autopolicies for you. You can only view these SAM policies through the appropriate Web resource profile—not through the SAM resource profile pages of the admin console.

To create a JSAM application resource profile:

1. In the admin console, choose **Users > Resource Profiles > SAM > Client Applications**.

2. Click **New Profile**.

3. From the Type list, choose **JSAM**.

4. From the Application list, select one of the following options.

   - **Custom**—Select this option to intermediate traffic to a custom application. Then:

     1. In the Server name field, enter the name or IP address of the remote server. If you are using automatic host mapping, enter the server as it is known to the application. If you enter an IP address, note that end users must connect to JSAM using that IP address in order to connect to the specified server.

     2. In the Server Port field, enter the port on which the remote server listens for client connections. For example, to forward Telnet traffic from a remote machine, specify port 23 for both the client port (on which JSAM listens) and the server port (on which the Telnet server listens).

        To disable the registry change made by JSAM and restore the original copy of the etc/hosts file, users must uninstall the JSAM client using settings in the Preferences > Applications page of the end-user console. To re-enable the change, they need to reboot.

        You can also use the restore system settings script. However, the restore system settings script cannot restore the hosts file successfully if you log in as a different user from the one that originally launched JSAM.

     3. In the Client Loopback IP field, provide a static loopback address. If you do not provide a static IP loopback address, the system assigns an IP loopback address dynamically.

        When configuring an external DNS, do not use IP loopback addresses in the 127.0.2.x range because the system reserves IP loopback addresses in that range for use with Citrix NFuse.

        If you want to modify a static loopback address for a JSAM application server configured on multiple ports, you must delete all applications referring to this application server and re-enter these applications with the new static loopback address.

     4. In the Client Port field, enter the port on which JSAM should listen for client application connections. Typically, the local port value is the same value as the server port; the local port value usually only differs for Linux or Macintosh non-root users who want to add applications for port forwarding that use ports under 1024.

        You may configure more than one application on a single port, such as app1.mycompany.com, app2.mycompany.com, app3.mycompany.com. Either you

assign a static loopback address or the system assigns a dynamic loopback address (127.0.1.10, 127.0.1.11, 127.0.1.12) to each application. JSAM then listens on these multiple loopback addresses on the specified port. For example, when there is traffic on 127.0.1.12 on the specified port, the system forwards the traffic to the app3.mycompany.com destination host.

5. Click **Add**.

6. Select the **Allow JSAM to dynamically select an available port if the specified client port is in use** check box if JSAM is listening for multiple hosts on the same port and you want JSAM to select an available port when the client port you specify is taken. The client application must allow you to specify the port number for the connection in order to use this option.

7. Select the **Create an access control policy allowing SAM access to these servers** check box to enable access to the list of servers specified in the Server column (enabled by default).

- **Lotus Notes**—Select this option to intermediate traffic from the Lotus Notes fat client application. Then, in the Autopolicy: SAM Access Control section, create a policy that allows or denies users access to the Lotus Notes server:

  1. If it is not already enabled, select the **Autopolicy: SAM Access Control** check box.

  2. In the Resource field, specify the application server to which this policy applies. You can specify the server as a fully-qualified hostname or an IP/netmask pair. For example, if the fully-qualified hostname is notes1.yourcompany.com, add notes1.yourcompany.com and notes1 to the Resource field.

  3. From the Action list, select **Allow** to enable access to the specified server or **Deny** to block access to the specified server.

  4. Click **Add**.

     > ℹ️ NOTE: If you select the Lotus Notes option, or you configure the Lotus Notes client to connect to multiple Lotus Notes servers, you should configure the Lotus Notes client appropriately to work with Connect Secure.
     >
     > You can only use JSAM to configure access to one Lotus Notes application per user role.

- **Microsoft Outlook**—Select this option to intermediate traffic from the Microsoft Outlook application. Then:

  1. Enter the hostname for each MS Exchange server in the Servers field. For example, if the fully-qualified hostname is exchange1.yourcompany.com, add exchange1.yourcompany.com to the Servers field.

     You must enter the full name of the servers in this field since the system creates direct one-to-one mappings between the servers you enter here and IP addresses in the etc/hosts file. For more information about registry changes made by JSAM, see the *Client-side Changes Guide* on the Pulse Secure Global Support Center (PSGSC) Center.

     The system does not support Outlook through SVW, since Outlook applications require HKLM registry key changes.

  2. Select the **Create an access control policy allowing SAM access to this server** check box to enable access to the server specified in the previous step (enabled by default).

ⓘ NOTE: You can only use JSAM to configure access to one Microsoft Outlook application per user role.

- **NetBIOS file browsing**—Select this option to tunnel NetBIOS traffic through JSAM. Then:

  1. Enter the fully-qualified hostname for your application servers in the Servers field.

     You must enter the full name of the servers in this field since the system creates direct one-to-one mappings between the servers you enter here and IP addresses in the etc/hosts file. For more information about registry changes made by JSAM, see the *Client-side Changes Guide* on the Pulse Secure Global Support Center (PSGSC) Center.

     If you want to enable drive mapping on a Windows client machine, use the standard NetBIOS file browsing option. When you do, JSAM automatically modifies the registry to disable port 445 on Windows machines, which forces Windows to use port 137, 138, or 139 for drive-mapping. Windows users need to reboot one time to enable the registry change to take effect.

  2. Select the **Create an access control policy allowing SAM access to this server** check box to enable access to the server specified in the previous step (enabled by default).

     ⓘ NOTE: You can only use JSAM to configure NetBIOS file browsing once per user role.

     The system does not support NetBIOS file browsing through SVW, since NetBIOS requires HKLM registry key changes.

3. Enter a unique name and optionally a description for the resource profile. The system displays this information in the Client Application Sessions section of the end-user home page.

4. Click **Save and Continue**.

5. In the Roles tab, select the roles to which the resource profile applies and click Add.

   The selected roles inherit the autopolicy created by the resource profile. If it is not already enabled, the system also automatically enables the SAM option in the Users > User Roles > *Role Name* > General > Overview page of the admin console for all of the roles you select.

6. Click **Save Changes**.

Related Documentation

- Using Custom Expressions in Rule Configuration
- Configuring the Lotus Notes Client

Specifying Applications for JSAM to Secure

Information in this section is provided for backwards compatibility. We recommend that you secure traffic using JSAM resource profiles instead, since they provide a simpler, more unified configuration method.

To specify applications for JSAM to secure:

1. In the admin console, choose **Users > User Roles >** *Select Role* **> SAM > Applications**.

2. Select **Add Application**.

3. Enter the name of the application and, optionally, a description. This information displays in the Client Application Sessions section of the end-user home page.

4. Choose either:

- **Standard application**—Select Citrix NFuse, Lotus Notes, or Microsoft Outlook/Exchange.

  The system does not support the standard JSAM applications Outlook and Netbios file browsing through SVW, since these applications require registry key changes. However, the system does support the Citrix and Lotus Notes JSAM standard applications through SVW.

  If you select the Lotus Notes option, or you configure the Lotus Notes client to connect to multiple Lotus Notes servers, you should configure the Lotus Notes client appropriately to work with Connect Secure.

  The system supports several mechanisms for intermediating traffic to the Lotus Notes, Microsoft Outlook, and Citrix applications.

- **Custom application**

  1. In the Server name field, enter the DNS name of the server or the server IP address. If entering the DNS name, enter name of the remote server as it is known to the application if you are using automatic host mapping.

  2. Enter the server name.

  3. In the Server Port field, enter the port on which the remote server listens for client connections.

     For example, to forward Telnet traffic from a remote machine, specify port 23 for both the client port (on which JSAM listens) and the server port (on which the Telnet server listens).

     To disable the registry change made by JSAM and restore the original copy of the etc/hosts file, users must uninstall the JSAM client using settings in the Preferences > Applications page of the end-user console. To re-enable the change, they need to reboot.

     You can also use the restore system settings script. However, the restore system settings script cannot restore the hosts file successfully if you log in as a different user from the one that originally launched JSAM.

  4. In the Client Loopback IP field, provide a static loopback address. If you do not provide a static IP loopback address, the system assigns an IP loopback address dynamically.

     When configuring an external DNS, do not use IP loopback addresses in the 127.0.2.x range because the system reserves IP loopback addresses in that range for use with Citrix NFuse.

     If you want to modify a static loopback address for a JSAM application server configured on multiple ports, you must delete all applications referring to this application server and re-enter these applications with the new static loopback address.

  5. In the Client Port field, enter the port on which JSAM should listen for client application connections.

     Typically, the local port value is the same value as the server port; the local port value usually only differs for Linux or Macintosh non-root users who want to add applications for port forwarding that use ports under 1024.

You may configure more than one application on a single port, such as app1.mycompany.com, app2.mycompany.com, app3.mycompany.com. Either you assign a static loopback address or the system assigns a dynamic loopback address (127.0.1.10, 127.0.1.11, 127.0.1.12) to each application. JSAM then listens on these multiple loopback addresses on the specified port. For example, when there is traffic on 127.0.1.12 on the specified port, the system forwards the traffic to the app3.mycompany.com destination host.

6. Select the **Allow Secure Application Manager to dynamically select an available port ...** check box if JSAM is listening for multiple hosts on the same port and you want JSAM to select an available port when the client port you specify is taken. The client application must allow you to specify the port number for the connection in order to use this option.

7. Click **Add**.

8. If a remote user's PC is set up to use a Web proxy in Internet Explorer, configure the client machine to bypass the proxy server when the user launches applications that need to connect to the Secure Application Manager.

9. Add DNS domains to the system if you have multiple internal domains, such as company-a.com and company-b.com, so that names such as app1.company-a.com and app2.company-b.com resolve correctly:

   1. In the admin console, choose **System > Network > Overview** .

   2. Under DNS Name Resolution, add a comma-separated list of domains in the to DNS Domains field.

   3. Click **Save Changes**.

Related Documentation

- Creating a JSAM Application Resource Profile

- Using Custom Expressions in Rule Configuration

- Configuring a PC that Connects Through a Proxy Web Server

Specifying Role Level JSAM Options

To specify JSAM options at the role level:

1. In the admin console, choose **Users > User Roles >** *Select Role* **> SAM > Options**.

2. Under Secure Application Manager options, select the options to enable for users:

   - **Auto-launch Secure Application Manager**—Select this option to automatically launches the Secure Application Manager when a user signs in. If you do not select this option, users must manually start the Secure Application Manager from the Client Applications Sessions section of the end-user home page.

     Although you configure the Secure Application Manager to automatically launch when users sign into the system, users can override this setting through the Preferences > Applications page of the end-user console. If disabled from automatically launching, users need to manually start the Secure Application Manager by clicking its link on the home page.

   - **Auto-uninstall Secure Application Manager**—Select this option to automatically uninstall the Secure Application Manager after users sign off.

- **Auto-allow application servers**—Select this option to automatically creates a SAM resource policy that allows access to the server specified in the WSAM application and server lists and the JSAM application list.

  You may not see the Auto-allow option if you are using a new installation or if an administrator hides the option.

3. Under Java SAM Options, select the options to enable for users:

- **User can add applications**—If enabled, users can add applications. For users to add applications, they need to know the application server DNS name and client/server ports.

  When you enable this option, users can set up port forwarding to any host or port in your enterprise. Before providing users with the ability to add applications, please verify that this feature is consistent with your security practices. If a user adds an application, the application remains available to the user even if you later change disable the feature.

- **Automatic host-mapping**—If enabled, the Secure Application Manager edits the Windows PC hosts file and replaces entries of Windows application servers with localhost. These entries are changed back to the original data when a user closes the Secure Application Manager.

  For the Java version of the Secure Application Manager to work, the client application needs to connect to the local PC on which the Secure Application Manager is running as the application server. The recommended process for mapping application servers to a user's local PC is to enable automatic host-mapping, which enables the system to automatically modify the PC's hosts file to point application servers to the PC's localhost for secure port forwarding. Alternatively, you can configure your external DNS server.

- **Skip web-proxy registry check**—If enabled, JSAM does not check a user's registry for a Web proxy. Some users do not have permissions to look at their registries, so if JSAM tries to look at their registries, then users see an error that they do not have permission. This option ensures that users do not see this message.

- **Auto-close JSAM window on sign-out**—If enabled, JSAM automatically closes when a user signs out of the device by clicking Sign Out in the browser window. JSAM continues to run if the user simply closes the browser window.

4. Click **Save Changes**.

Related Documentation

- Task Summary: Configuring JSAM

Automatically Launching JSAM

Use the Launch JSAM tab to write a Web resource policy that specifies a URL for which the system automatically launches JSAM on the client. The system launches JSAM in two scenarios:

- When a user enters the URL in the Address field of the home page.

- When a user clicks a Web bookmark (configured by an administrator) on the home page to the URL.

This feature is useful if you enable applications that require JSAM but don't want to require users to run JSAM unnecessarily. This feature requires, however, that users access the URL through the home page. If users enter the URL in a browser Address field, the system does not serve the request.

The system provides tight integration with Citrix. If you specify Citrix as a standard JSAM application, the system automatically launches JSAM when a user selects an ICA file even if the URL is not configured as a resource policy.

To write a Launch JSAM resource policy:

1. In the admin console, choose **Users > Resource Policies > Web**.

2. If your administrator view is not already configured to show Launch JSAM policies, make the following modifications:

    1. Click the **Customize** button in the upper right corner of the page.

    2. Select the **Launch JSAM** check box.

    3. Click **OK**.

4. Select the **Launch JSAM** tab.

5. On the JSAM Autolaunch Policies page, click **New Policy**.

6. Enter a name to label this policy (required) and a description of the policy (optional).

7. In the Resources section, specify the URLs to which this policy applies.

    > NOTE: The resource policies configured for the JSAM auto launch policy must be a specific URL and not include wildcards. The URL should specify the entry point of the web application for which JSAM tunneling is needed.

8. In the Roles section, specify:

    - **Policy applies to ALL roles**—Choose this option to apply this policy to all users.

    - **Policy applies to SELECTED roles**—Choose this option to apply this policy only to users who are mapped to roles in the Selected roles list. Make sure to add roles to this list from the Available roles list.

    - **Policy applies to all roles OTHER THAN those selected below**—Choose this option to apply this policy to all users except for those who map to the roles in the Selected roles list. Make sure to add roles to this list from the Available roles list.

9. In the Action section, specify:

    - **Launch JSAM for this URL**—The system downloads the Java Secure Application Manager to the client and then serves the requested URL.

      JSAM launches automatically for the specified URL only if a user enters the URL or selects a bookmark to the URL on the home page (Browsing > Bookmarks). The bookmark does not launch the application that is configured through JSAM, but launches JSAM itself.

    - **Don't Launch JSAM for this URL**—The system does not download the Java Secure Application Manager to the client for the requested URL. This option is useful if you want to temporarily disable JSAM auto-launching for the specified URLs.

    - **Use Detailed Rules**—To specify one or more detailed rules for this policy.

10. Click **Save Changes**.

Related Documentation

- Task Summary: Configuring JSAM

Specifying Application Servers that Users Can Access

Information in this topic is provided for backwards compatibility. We recommend that you secure traffic using JSAM resource profiles instead, since they provide a simpler, more unified configuration method.

When you enable the Secure Application Manager access feature for a role, you need to create resource policies that specify which application servers a user may access. These policies apply to both the Java version and Windows version of the Secure Application Manager (JSAM and WSAM, respectively). When a user makes a request to an application server, the system evaluates the SAM resource policies. If the system matches a user's request to a resource listed in a SAM policy, the system performs the action specified for the resource.

When writing a SAM resource policy, you need to supply key information:

- Resources—A resource policy must specify one or more resources to which the policy applies. When writing a SAM policy, you need to specify application servers to which a user may connect.

- Roles—A resource policy must specify the roles to which it applies. When a user makes a request, the system determines what policies apply to the role and then evaluates those policies that correspond to the request. SAM resource policies apply to users requests made through either version, JSAM or WSAM.

- Actions—A Secure Application Manager resource policy either allows or denies access to an application server.

You can create resource policies through the standard interface (as described in this section) or through resource profiles (recommended method).

The engine that evaluates resource policies requires that the resources listed in a policy's Resources list follow a canonical format.

To write a Secure Application Manager resource policy:

1. In the admin console, choose **Users > Resource Policies > SAM > Access**.

2. On the Secure Application Manager Policies page, click **New Policy**.

3. Enter a name to label this policy (required) and a description of the policy (optional).

4. In the Resources section, specify the application servers to which this policy applies.

5. In the Roles section, specify:

    - **Policy applies to ALL roles**—Choose this option to apply this policy to all users.

    - **Policy applies to SELECTED roles**—Choose this option to apply this policy only to users who are mapped to roles in the Selected roles list. Make sure to add roles to this list from the Available roles list.

    - **Policy applies to all roles OTHER THAN those selected below**—Choose this option to apply this policy to all users except for those who map to the roles in the Selected roles list. Make sure to add roles to this list from the Available roles list.

6. In the Action section, specify:

    - **Allow socket access**—Choose this option to grant access to the application servers specified in the Resources list.

- **Deny socket access**—Choose this option to deny access to the application servers specified in the Resources list.

- **Use Detailed Rules**—Choose this option to specify one or more detailed rules for this policy.

7. Click **Save Changes**.

8. On the Secure Application Manager Policies page, order the policies according to how you want to evaluate them. Keep in mind that once the system matches the resource requested by the user to a resource in a policy's (or a detailed rule's) Resource list, it performs the specified action and stops processing policies.

Related Documentation

- Task Summary: Configuring JSAM

Specifying Resource Level JSAM Options

Use the Options tab to specify the SAM resource option to match IP addresses to hostnames specified as resources in your SAM resource policies. When you enable this option, the system looks up IP addresses corresponding to each hostname specified in a SAM resource policy. When a user tries to access a server by specifying an IP address rather than the hostname, the system compares the IP to its cached list of IP addresses to determine if a hostname matches an IP. If there is a match, then the system accepts the match as a policy match and applies the action specified for the resource policy.

When you enable this option, the system compiles a list of hostnames specified in the Resources field of each SAM resource policy. The system then applies the option to this comprehensive list of hostnames.

> 🛈 NOTE: This option does not apply to hostnames that include wildcards and parameters.

To specify the SAM resource option:

1. In the admin console, choose **Users > Resource Policies > SAM > Options**.

2. Select IP based matching for Hostname based policy resources. When you select this option, the the system looks up the IP address corresponding to each hostname specified in a Secure Application Manager resource policy. When a user tries to access a server by specifying an IP address rather than the hostname, the system compares the IP to its cached list of IP addresses to determine if a hostname matches an IP. If there is a match, then the system accepts the match as a policy match and applies the action specified for the resource policy.

3. Click **Save Changes**.

Related Documentation

- Task Summary: Configuring JSAM

CHAPTER 25 Telnet/SSH

About Telnet/SSH

The Telnet/SSH option enables users to connect to internal server hosts in the clear using Telnet protocols or to communicate over an encrypted Secure Shell (SSH) session through a Web-based terminal session emulation. This feature supports the following applications and protocols:

- Network Protocols—Supported network protocols include Telnet and SSH.

- Terminal Settings—Supported terminal settings include VT100, VT320, and derivatives and screen buffers.

- Security—Supported security mechanisms include Web/client security using SSL and host security (such as SSH if desired).

You can create secure terminal session bookmarks that appear on the welcome page for users mapped to a specific role. A terminal session bookmark defines Terminal Session information for Telnet or SSH sessions that users may launch. These sessions give users access to a variety of networked devices, including UNIX servers, networking devices, and other legacy applications, that utilize terminal sessions. The system supports SSH versions V1 and V2 and uses the following SSH versions: OpenSSH 5.2, OpenSSH_2.9.9p1, SSH protocols 1.5/2.0, and OpenSSL 0x0090607f.

When communicating over an encrypted Secure Shell (SSH) session, note that the Telnet/SSH feature does not support using the ^J character combination. (Some applications use this character combination to justify text). If you want to use this character combination, we recommend that you find a java applet that supports it and upload that applet through the system using the hosted Java applets feature.

Related Documentation

Task Summary: Configuring the Telnet/SSH Feature

To configure the Telnet/SSH feature:

1. Create resource profiles that enable access to Telnet and SSH servers, include bookmarks that link to those servers, and assign the bookmarks to user roles using settings in the Users > Resource Profiles > Telnet/SSH page of the admin console.

We recommend that you use resource profiles to configure Telnet/SSH (as described above). However, if you do not want to use resource profiles, you can configure Telnet/SSH using role and resource policy settings in the following pages of the admin console instead:

- Create resource policies that enable access to Telnet and SSH servers using settings in the Users > Resource Policies > Telnet/SSH > Sessions page of the admin console.

- Determine which user roles may access the Telnet and SSH servers that you want to intermediate, and then enable Telnet/SSH access for those roles through the Users > User Roles > Select Role > General > Overview page of the admin console.

- Create bookmarks to your Telnet and SSH servers using settings in the Users > User Roles > Select Role > Telnet/SSH > Access page of the admin console.

2. After configuring Telnet/SSH using resource profiles or roles and resource policies, you can modify general role and resource options in the following pages of the admin console:

- (Optional) Enable users to create their own connections to Telnet and SSH sessions using settings in the Users > User Roles > Select Role > Telnet/SSH > Options page of the admin console.

- (Optional) Enable the system to match IP addresses to hostnames and disable the auto-allow bookmarks option using settings in the Users > Resource Policies> Telnet/SSH > Options page of the admin console.

Related Documentation

- About Telnet/SSH

- Creating a Telnet/SSH Resource Profile

- Associating Bookmarks with Telnet/SSH Resource Profiles

- Configuring General Telnet/SSH Options

- Writing a Telnet/SSH Resource Policy

Creating a Telnet/SSH Resource Profile

A Telnet/SSH resource profile is a resource profile that enables users to connect to internal server hosts in the clear using Telnet protocols or to communicate over an encrypted Secure Shell (SSH) session through a Web-based terminal session emulation.

To create a Telnet/SSH resource profile:

1. In the admin console, choose **Users > Resource Profiles > Telnet/SSH**.

2. Click **New Profile**.

3. From the Type list, specify the session type (Telnet or SSH) for this resource profile.

4. Enter a unique name and optionally a description for the resource profile. (This name becomes the default bookmark's name.)

5. In the Host field, enter the name or IP address of the server to which this resource profile should connect.

6. Select the **Create an access control policy allowing Telnet/SSH access to this server** check box to enable access to the server specified in the previous step (enabled by default).

7. In the Port field, enter the port on which the system should connect to the server. (By default, the system populates this field with port number 23 if you select Telnet and port number 22 if you select SSH.)

8. If you want to pass the user's credentials to the server, enter a static username, the <username> variable, or another appropriate session variable in the Username field. (Required for SSH sessions.)

9. Click **Save and Continue**.

10. In the Roles tab, select the roles to which the resource profile applies and click **Add**.

    The selected roles inherit the autopolicy and bookmarks created by the resource profile. If it is not already enabled, the system also automatically enables the Telnet/SSH option in the Users > User Roles > Select Role > General > Overview page of the admin console for all of the roles you select.

11. Click **Save Changes**.

12. (Optional) In the Bookmarks tab, modify the default bookmark created by the system and/or create new ones. (By default, the system creates a bookmark to the server defined in the Host field and displays it to all users assigned to the role specified in the Roles tab.)

Related Documentation

- Task Summary: Configuring the Telnet/SSH Feature
- Using Custom Expressions in Rule Configuration

Associating Bookmarks with Telnet/SSH Resource Profiles

When you create a Telnet/SSH resource profile, the system automatically creates a bookmark that links to the host that you specified in the resource profile. The system enables you to modify this bookmark as well as create additional bookmarks to the same host.

You can use two different methods to create Telnet/SSH session bookmarks:

- Create bookmarks through existing resource profiles (recommended)—When you select this method, the system automatically populates the bookmark with key parameters (such as the host, port, username, and session type) using settings from the resource profile. Additionally, while you are creating the associated resource profile, the system guides you through the process of creating any required policies to enable access to the bookmark.

- Create standard bookmarks—When you select this option, you must manually enter all bookmark parameters during configuration. Additionally, you must enable access to the Telnet/SSH feature and create resource policies that enable access to the servers defined in the bookmark.

Creating Bookmarks Through Existing Resource Profiles

When configuring bookmarks, note that:

- To change the host, port, or username for a Telnet/SSH bookmark created through a resource profile, you must edit the values through the resource profile's Resource tab (not its Bookmark tab).

- You can only assign bookmarks to roles that you have already associated with the resource profile—not all of the roles defined on the system. To change the list of roles associated with the resource profile, use settings in its Roles tab.

- Bookmarks simply control which links are displayed to users—not which resources the users can access. For example, if you enable access to a Telnet server through a resource profile but do not create a corresponding bookmark to that server, the user can still access the server by entering it into the Address field of the home page.

- Make sure to enter a unique set of parameters when defining a Telnet/SSH bookmark. If you create two bookmarks that contain the same set of parameters, the system deletes one of

the bookmarks from the end-user view. You will still be able to see both bookmarks, however, in the administrator console.

To associate bookmarks with Telnet/SSH resource profiles:

1. If you want to create a resource profile bookmark through the standard resource profiles page:

   - Choose **Users > Resource Profiles > Telnet/SSH> Select Resource Profile > Bookmarks**.

   - Click the appropriate link in the Bookmark column if you want to modify an existing bookmark. Or, click **New Bookmark** to create an additional bookmark.

   Alternatively, if you want to create a resource profile bookmark through the user roles page:

   1. Choose **Users > User Roles > Select Role > Telnet/SSH > Sessions**.

   2. Click **Add Session**.

   3. From the Type list, choose **Telnet/SSH Resource Profile**. (The system does not display this option if have not already created a Telnet/SSH resource profile.)

   4. Select an existing resource profile. (The system automatically populates the Host and Port fields using settings from the selected resource profile.)

   5. Click **OK**. (If you have not already associated the selected role with the resource profile, the system automatically makes the association for you. The system also enables any access control policies for the role that are required by the resource profile.)

   6. If this role is not already associated with the selected resource profile, the system displays an informational message. If you see this message, click **Save Changes** to add this role to the resource profile's list of roles and to update the profile's autopolicies as required. Then, repeat the previous step to create the bookmark.

      > *i*
      >
      > NOTE: When you create a resource profile bookmark through the user roles page (instead of the standard resource profiles page), the system only associates the generated bookmark with the selected role. The system does not assign the bookmark to all of the roles associated with the selected resource profile.

   7. Optionally change the name and description of the bookmark. (By default, the system names the bookmark the same as the resource profile name.)

   8. If you want to change the font size in the server display window, choose one of the following options in the Font Size section:

      - **Fixed size of _ pixels**—Enter a size from 8 to 36 pixels. (The default font size is 12.)

      - **Resize to fit window**—Dynamically changes the font size as you resize the window. This option requires Internet Explorer. (Enabled by default.)

   9. If you want to change the size of the server display window, select an option from the Screen Size drop-down list. The default window size is 80 characters by 24 rows.

   10. If you want to change the number of rows that the server window retains to display during scrolling, change the value in the Screen Buffer field. The default buffer size is 100 rows.

   11. If you are configuring the bookmark through the resource profile pages, under Roles, specify the roles to which you want to display the bookmark:

       - **ALL selected roles**—Select this option to display the bookmark to all of the roles associated with the resource profile.

- **Subset of selected roles**—Select this option to display the bookmark to a subset of the roles associated with the resource profile. Then select roles from the ALL Selected Roles list and click Add to move them to the Subset of selected roles list.

12.   Click **Save Changes**.

Creating Standard Bookmarks

Information in this topic is provided for backwards compatibility. We recommend that you configure access to Telnet and SSH servers through resource profiles instead, since they provide a simpler, more unified configuration method.

To create a bookmark for secure terminal sessions:

1. In the admin console, choose **Users > User Roles > Select Role > Telnet/SSH > Sessions**.

2. Click **Add Session**. The New Telnet/SSH Session page loads.

3. From the Type list, choose **Standard**. (The system only displays the Type list if you have already created a Telnet/SSH resource profile.)

4. Enter a bookmark name and description for the new Telnet/SSH session (optional). If you specify a bookmark name and description, this information appears on the Terminal Sessions page.

5. Enter the name or IP address of the remote host for this session in the Host field.

6. Select the Session Type, either **Telnet** or **SSH Secure Shell**, and specify the port if different from the pre-populated port assignment.

7. Provide a username or use the <username> or other appropriate session variable.

8. Specify the Font Size by selecting one of the following:

   - **Fixed size of _ pixels**—enter a size from 8 to 36 pixels.

   - **Resize to fit window**—dynamically changes the font size as you resize the window. This option requires Internet Explorer.

9. Select the Screen Size using the drop-down list.

10.   Specify the Screen Buffer size.

11.   Click **Save Changes** or **Save + New**.

In addition to creating bookmarks for secure terminal sessions, you must create a resource policy allowing Telnet/SSH sessions for the role, or enable Auto-allow role Telnet/SSH sessions on the Telnet/SSH > Options tab to automatically allow access to the resources defined in the session bookmark.

Make sure to enter a unique set of parameters when defining a Telnet/SSH bookmark. If you create two bookmarks that contain the same set of parameters, the system deletes one of the bookmarks from the end-user view. You will still be able to see both bookmarks, however, in the administrator console.

Related Documentation

- Task Summary: Configuring the Telnet/SSH Feature

Configuring General Telnet/SSH Options

You can enable users to create their own Telnet/SSH bookmarks, browse to a terminal session, and to configure the system to create Telnet/SSH resource policies that allow access to the servers specified in the session bookmarks.

When you allow users to browse to a terminal session, note that they can use two different methods:

- Use the homepage—Users can enter the server and port that they want to access into the Address field of the home page. Valid formats for the URL include:

  - Telnet://host:port

  - SSH://host:port

  For example: Telnet://terminalserver.yourcompany.com:3389

- Use the Web browser's address bar—Users can enter the server and port that they want to access (as well as session parameters such as font and window size) into the address bars of their Web browsers using standard Web protocol. For example:

  https://iveserver/dana/term/newlaunchterm.cgi?
  protocol=telnet&host=termsrv.yourcompany.com&port=23&
  username=jdoe&fontsize=12&buffer=800&size=80x25

To specify general Telnet/SSH options:

1. In the admin console, choose **Users > User Roles > Select Role > Telnet/SSH > Options**.

2. Enable **User can add sessions** to allow users to define their own session bookmarks and to allows users to browse to a terminal session using telnet:// and ssh:// syntax as well as the /dana/term/newlaunchterm.cgi syntax. When you enable this option, the Add Terminal Session button appears on the Terminal Sessions page the next time a user refreshes the welcome page.

3. Enable **Auto-allow role Telnet/SSH sessions** to enable the system to automatically allow access to the resources defined in the session bookmark (rather than having to create resource policies). Note that this only applies to role bookmarks, not user bookmarks.

   You may not see the Auto-allow option if you are using a new installation or if an administrator hides the option.

4. Click **Save Changes**.

Related Documentation

- Task Summary: Configuring the Telnet/SSH Feature

Writing a Telnet/SSH Resource Policy

When you enable the Telnet/SSH access feature for a role, you need to create resource policies that specify which remote servers a user may access. If the system matches a user's request to a resource listed in a Telnet/SSH policy, it performs the action specified for the resource.

You can create resource policies through the standard interface (as described in this topic) or through resource profiles (recommended method).

When writing a Telnet/SSH resource policy, you need to supply key information:

- Resources—A resource policy must specify one or more resources to which the policy applies. When writing a Telnet/SSH policy, you need to specify remote servers to which a user may connect.

- Roles—A resource policy must specify the roles to which it applies. When a user makes a request, the system determines what policies apply to the role and then evaluates those policies that correspond to the request.

- Actions—A Telnet/SSH resource policy either allows or denies access to a server.

The engine that evaluates resource policies requires that the resources listed in a policy's Resources list follow a canonical format.

Writing Telnet/SSH Resource Policies

Information in this section is provided for backwards compatibility. We recommend that you configure access to Telnet and SSH servers through resource profiles instead, since they provide a simpler, more unified configuration method.

To write a Telnet/SSH resource policy:

1. In the admin console, choose **Users > Resource Policies > Telnet/SSH > Access**.

2. On the Telnet/SSH Policies page, click **New Policy**.

3. On the New Policy page, enter a name to label this policy and optionally a description.

4. In the Resources section, specify the servers to which this policy applies.

5. In the Roles section, specify:

   - **Policy applies to ALL roles**—Use this field to apply this policy to all users.

   - **Policy applies to SELECTED roles**—Use this field to apply this policy only to users who are mapped to roles in the Selected roles list. Make sure to add roles to this list from the Available roles list.

   - **Policy applies to all roles OTHER THAN those selected below**—Use this field to apply this policy to all users except for those who map to the roles in the Selected roles list. Make sure to add roles to this list from the Available roles list.

6. In the Action section, specify:

   - **Allow access**—Use this field to grant access to the servers specified in the Resources list.

   - **Deny access**—Use this field to deny access to the servers specified in the Resources list.

   - **Use Detailed Rules**—Use this field to specify one or more detailed rules for this policy.

7. Click **Save Changes**.

8. On the Telnet/SSH Policies page, order the policies according to how you want to evaluate them. Keep in mind that once the system matches the resource requested by the user to a resource in a policy's (or a detailed rule's) Resource list, it performs the specified action and stops processing policies.

Matching IP Addresses to Hostnames

You can configure Telnet/SSH to match IP addresses to hostnames specified as resources in your Telnet/SSH resource policies. When you enable this option, the system looks up IP address corresponding to each hostname specified in a Telnet/SSH resource policy. When a user tries to access a server by specifying an IP address rather than the hostname, the system compares the IP to its cached list of IP addresses to determine if a hostname matches an IP. If there is a match, then the system accepts the match as a policy match and applies the action specified for the resource policy.

When you enable this option, the system compiles a list of hostnames specified in the Resources field of each Telnet/SSH resource policy. The system then applies the option to this comprehensive list of hostnames.

This option does not apply to hostnames that include wildcards and parameters.

To specify the Telnet/SSH resource option:

1. In the admin console, choose **Users > Resource Policies > Telnet/SSH > Options**.

2. Select **IP based matching for Hostname based policy resources**.

   The system looks up the IP address corresponding to each hostname specified in a Telnet/SSH resource policy. When a user tries to access a server by specifying an IP address rather than the hostname, the system compares the IP to its cached list of IP addresses to determine if a hostname matches an IP. If there is a match, then the system accepts the match as a policy match and applies the action specified for the resource policy.

3. Click **Save Changes**.

Related Documentation

- Task Summary: Configuring the Telnet/SSH Feature

CHAPTER 26 Terminal Services

About Terminal Services

Use the Terminal Services feature to enable a terminal emulation session on a Windows terminal server, Citrix NFuse server, or Citrix Metaframe server. You can also use this feature

to deliver the terminal services through the system, eliminating the need to use another Web server to host the clients.

The system supports several mechanisms for intermediating traffic between a Citrix server and client, including the Terminal Services, JSAM, WSAM, VPN Tunneling, and hosted Java applets features.

Terminal Services User Experience

From an end-user perspective, accessing secured terminal services resources through the system is simple. When you enable the Terminal Services feature for a user role, the end user simply needs to do the following tasks:

1. Specify the resource that the user wants to access—The user can specify the resource he wants to access by clicking a link or entering the resource in the system browse bar. Or, if you enable auto-launch for a bookmark, the system automatically launches the resource for the user when he signs into the device.

2. Enter credentials for the resource—When the user accesses a resource, the system prompts him to enter his username and password (if required by the resource). Or if you enable SSO, the system automatically sends this information to the resource without prompting the user. Once the resource verifies the credentials, the system launches the resource.

Users can access terminal services resources using the following methods:

- Session bookmarks—A session bookmark defines various information, including the server to which the user can connect, the terminal session's window parameters, and the username and password that the system sends to the Windows terminal server or Metaframe server. You can create any number of session bookmarks for a role, enabling the user to access multiple servers using different session bookmarks for each. (Users can simultaneously open multiple sessions to the same terminal server or to different servers.)

- URLs from other web sites—In most cases, users access session bookmarks directly from the end-user console. If you do not want to require users to sign into the end-user console to find and access terminal services links, you can create URLs on other web sites that point to session bookmarks that you have already created. Or, you can create URLs that include all of the parameters that you want to pass to the Terminal Services program, such as the host, ports, and terminal window parameters.

  NOTE: If you create links on external servers to terminal services bookmarks on the system and you are using multiple customized sign-in URLs, some restrictions occur.

- Connect Secure browse bar—In addition to enabling users to link to terminal services links through bookmarks and URLs, you can also enable them to access these resources through the system browse bar on Windows systems. Users can access Citrix Metaframe or Nfuse servers by entering ica://*hostname* in the browse box. Or, users can access Microsoft terminal services or remote desktop sessions by entering rdp://*hostname* in the browse box.

- Server address—By entering a terminal server IP address or hostname, users can launch a remote desktop connection to any accessible server.

Related Documentation

- Task Summary: Configuring the Terminal Services Feature

Task Summary: Configuring the Terminal Services Feature

To configure the Terminal Services feature:

1. Create resource profiles that enable access to Windows terminal servers or Citrix servers, include session bookmarks that link to those servers, and assign the session bookmarks to user roles using settings in the Users > Resource Profiles > Terminal Services page of the admin console.

   We recommend that you use resource profiles to configure terminal services (as described here). However, if you do not want to use resource profiles, you can configure the Terminal Services feature using role and resource policy settings in the following pages of the admin console instead:

   - Create resource policies that enable access to Windows terminal servers and Citrix servers using settings in the Users > Resource Policies > Terminal Services > Access page of the admin console.

   - Determine which user roles may access the Windows terminal servers and Citrix servers that you want to intermediate, and then enable Terminal Services access for those roles through the Users > User Roles > Select_Role > General > Overview page of the admin console.

   - Create session bookmarks to your Windows terminal servers and Citrix servers using settings in the Users > User Roles > Select_Role > Terminal Services > Sessions page of the admin console.

2. (Optional.) Modify general role and resource options after configuring terminal services using resource profiles or roles and resource policies. Use the following pages of the admin console:

   - (Optional.) Enable users to define their own terminal services sessions, specify the local devices to which users can connect, and set display and performance options using settings in the Users > User Roles > Select_Role > Terminal Services > Options page of the admin console. If you choose to enable users to define their own terminal services sessions, you must also create corresponding resource policies or resource profiles that enable access the specified resources, as explained in earlier in this topic.

   - (Optional.) Create links to a terminal services session that users can access from an external web site.

   - (Optional.) Enable the system to match IP addresses to hostnames using settings in the Users > Resource Policies> Terminal Services > Options page of the admin console.

3. (Citrix only) Specify where the system should obtain the Citrix client to upload to the users' systems through settings in the Users > User Roles > Select_Role > Terminal Services > Options page of the admin console.

   Additionally, if you specify that the system should obtain a Citrix client from an external web site, you must:

   - Create a Web access resource policy that enables access to the web site where the Citrix client resides through settings in the Users > Resource Policies > Web > Access > Web ACL page of the admin console.

   - Create a Web caching resource policy through settings in the Users > Resource Policies > Web > Caching page of the admin console so the user's browser can deliver the Citrix client. (Note that you must select the Unchanged (do not add/modify caching headers) option.)

Related Documentation

- About Terminal Services

- Terminal Services Execution

- About Terminal Services Resource Profiles

- Specifying General Terminal Services Options

- Configuring Terminal Services Resource Policies

- Specifying the Terminal Services Resource Option

Terminal Services Execution

When a user tries to access a terminal services resource, the system completes the following steps to initiate and intermediate the terminal services session:

1. The system checks for a Java client.

   To enable a terminal services session, the user either needs an RDP client on his system (to access a Windows terminal server) or an ICA client (to access a Citrix Metaframe server or server farm). These clients come in both Windows and Java versions and enable the user to run an application on the server while only transmitting keyboard, mouse, and display information over the network.

   The system enables you to upload a Java version of the RDP or ICA client through a terminal services resource profile (but not role). If you have uploaded a client to the system and specified that the system always use it to run your users' terminal sessions, the system launches the specified Java client.

2. (Citrix only.) If necessary, the system checks for a Windows client.

   If you have not uploaded a Java client, the system checks for a Windows version of the ICA client. If it cannot find a Windows ICA client, it installs the version you specified in the Users > User Roles > Role >Terminal Services > Options page of the admin console.

3. The system checks for the terminal services proxy.

   To intermediate a Windows or Citrix session, the user either needs a Juniper Windows Terminal Services proxy on his system or a Juniper Networks Citrix Terminal Services proxy. The system checks for the appropriate proxy on the user's computer, and if it cannot find it, installs a new one. Depending on the user's rights, the system either uses an ActiveX component or Java component to install the proxy.

4. The proxy tries to invoke the Windows client.

   Once the system has confirmed that a proxy is installed on the user's computer, the proxy attempts to invoke the Windows RDP or ICA client. If successful, the client initiates the user's terminal services session and the proxy intermediates the session traffic.

5. The proxy tries to invoke the Java client.

   If a Windows client is not present on the user's machine (for instance, because it was deleted or because the user does not have the proper privileges to install it), but you have uploaded one to the system through the terminal services resource profile, the system uses the uploaded Java applet to launch the session.

   As part of the installation, the system asks the user if he wants to always use the Java client or only for this session. The system then stores the user's preference as a persistent cookie. Once the Java client is installed, the client initiates the user's terminal services session and the proxy intermediates the session traffic.

For information about the specific files installed by the system when you enable the Terminal Services feature, as well as the rights required to install and run the associated clients, see the *Client-side Changes Guide* on the Pulse Secure Global Support Center (PSGSC) Center.

Related Documentation

- Task Summary: Configuring the Terminal Services Feature

Configuring Citrix to Support ICA Load Balancing

The Service Terminal Services feature supports connecting to Citrix server farms in which published applications are preconfigured (as described later in this topic). The system does not support load balancing configurations in which Nfuse servers dynamically retrieve a list of Citrix published applications within a server farm.

Citrix Load Balancing Overview

The system supports the following Citrix load balancing scenario:

1. The Citrix administrator makes a published application available to multiple Citrix servers in a farm by generating a custom ICA file. The generated ICA file contains a parameter called HTTPBrowserAddress that points to the IP address and port number of the master browser (that is, the server that performs the load balancing).

2. When the ICA client attempts to launch a published application on the user's computer, it uses the HTTPBrowserAddress parameter to connect to the master browser.

3. The master browser pings servers in the farm to determine their respective loads and returns the IP address of the least busy server to the ICA client.

4. The ICA client uses the IP address returned by the master browser to connect to the appropriate terminal server.

Configuring Citrix Load Balancing

For the system to work properly with a Citrix farm, you must configure the Citrix farm and Connect Secure as described in the following steps. Note that these instructions are based on using a Citrix Metaframe Presentation Server 3.0.

1. On the Citrix server, enable a server (or multiple servers) in your farm as a master browser:

    1. Right-click a server in the Metaframe Farm and select **Properties**.

    2. Select **Metaframe Settings**.

    3. Enter the TCP/IP port for the Citrix XML service.

4. On the Citrix server, publish the applications that are hosted on MetaFrame XP servers in the farm:

    1. Right-click the Applications link and select **Publish applications**.

    2. Specify which desktop or application to publish.

    3. Follow the prompts in the wizard.

    4. Specify the list of servers that host the application you are publishing and click **Finish**.

        The specified published application appears in the server farm.

5. On the Citrix server, generate a corresponding Citrix ICA file for the published application:

    1. Select the application you published in Step 2 and select **Create ICA file**.

    2. Follow the prompts in the wizard.

3. On the TCP/IP + HTTP Server page, enter the name of the HTTP browser server and the port number. (The port should match the Citrix XML Service port that you set up in Step 1).

4. Save the ICA file.

5. On Connect Secure, upload the ICA file using settings in either of the following admin console pages:

- Users > User Roles > *Role* > Terminal Services > Sessions

- Users > Resource Profiles > *Profile*

6. On Connect Secure, create a resource policy for the HTTP browser server and port entered in Step 3.

7. On Connect Secure, test the configuration by launching the bookmark as an end user.

> NOTE: One of the Citrix servers in the farm performs the load balancing, not Connect Secure. If the ICA client is already installed on the user's desktop then administrator rights are not required.
>
> For more information about the rights required to use the Terminal Services feature, see *Pulse Connect Secure Client-Side Changes Installation Reference*.
>
> If the XML response from the master browser contains the hostname, it will not work through Connect Secure. To ensure that the response is in dot-port format (an IP address), clear the Enable XML service DNS address resolution check box during the browser server configuration. This option controls whether the destination Citrix server is represented as a hostname or as an IP address.

Related Documentation

- Comparing Access Mechanisms for Configuring Citrix

About Terminal Services Resource Profiles

Terminal Services resource profile configuration instructions vary depending on whether you want to configure access to a Windows terminal server (which requires a RDP client) or Citrix terminal server (which requires an ICA client). Furthermore, if you choose to configure access to a Citrix server using a custom ICA file, you include many of your configuration settings in the ICA file itself and therefore do not need to configure them through the system. If you configure access to a Citrix server using the default ICA file on the system, however, you must configure additional settings.

You may want to create multiple bookmarks for the same terminal services resource in order to provide easy access to multiple applications. For instance, the server defined in your resource profile may provide access to multiple applications (such as Siebel and Outlook). To easily provide access to each of these applications, you can create resource profile bookmarks to each. Or, you may want to use multiple bookmarks to configure single sign-on to one application, but not another.

When configuring session bookmarks, note that:

- To change the host or ports for a terminal services session bookmark created through a resource profile, you must edit the values through the resource profile's Resource tab (not its Bookmark tab).

- You can only assign session bookmarks to roles that you have already associated with the resource profile—not all of the roles defined on the system. To change the list of roles associated with the resource profile, use settings in its Roles tab.

- Session bookmarks simply control which links to display to users—not which resources the users can access. For example, if you enable access to a terminal server through a resource profile but do not create a corresponding session bookmark to that server, the user can still access the server by entering it into the Address box of the home page.

- Make sure to enter a unique set of parameters when defining a terminal services bookmark. If you create two bookmarks that contain the same set of parameters, the system deletes one of the bookmarks from the end-user view. You can still see both bookmarks, however, in the administrator console.

Related Documentation

- Task Summary: Configuring the Terminal Services Feature

- Defining Resource Profile Bookmarks

Configuring a Windows Terminal Services Resource Profile

This topic describes how to configure a terminal services resource profile that enables access to a Windows terminal server using an RDP client.

Users can use RDP7 features through the Juniper Networks terminal services if an RDP7 client is present. However, the true multi-monitor and bidirectional audio features of RDP7 are not supported with this release.

To create a Windows terminal services resource profile:

1. In the admin console, select **Users > Resource Profiles > Terminal Services**.

2. Click **New Profile**. Or select an existing profile from the list.

3. Select **Windows Terminal Services** from the Type list.

4. Enter a unique name and optionally a description for the resource profile. (This name becomes the default session bookmark's name.)

5. Specify the server and port to which this resource profile should connect in the Host field. When entering the server, you may enter a hostname or IP address.

6. Enter the port on which the terminal server listens in the Server port box. (By default, the system populates this box with port number 3389.)

7. Select the Cre**ate an access control policy allowing Terminal Service access to this server** check box to enable access to the server specified in the Server Port box (enabled by default).

8. If you want to enable intermediation using a Java client, select **Enable Java support** and then specify which Java client the system should use.

9. Click **Save and Continue**.

10. Select the roles to which the resource profile applies in the Roles tab and click Add.

    The selected roles inherit the autopolicy and session bookmarks created by the resource profile. If it is not already enabled, the system also automatically enables the Terminal Services option in the Users > User Roles > Select_Role > General > Overview page of the admin console for all of the roles you select.

11. Click **Save Changes**.

12.   (Optional.) Modify the default session bookmark created by the system in the Bookmarks tab and/or create new ones. By default, the system creates a session bookmark to the server defined in the Host box and displays it to all users assigned to the role specified in the Roles tab.)

Related Documentation

- Using Custom Expressions in Rule Configuration

Defining a Hosted Java Applet Autopolicy

Hosted Java applet autopolicies enable you to store terminal services Java clients directly on the system without employing a separate Web server to host them. You can then associate these Java applets with the resource profile and specify that the system always use them to intermediate traffic, or that the system fall back to the applets when other terminal services clients are not available on the user's system.

Although you can use a Java applet to intermediate traffic to an SSO-enabled resource, we do not recommend it because the applet may require the user's password to be presented as plain text.

A default Premier Java RDP applet is shipped with each device and cannot be deleted. The HOB applet is available through the New Terminal Services Resource Profile window and the Users > User Roles > Users > Terminal Services > Options window. To use the Pulse Secure-supplied HOB applet, you must contact Pulse Secure Customer Care to purchase a license including the number of concurrent users you want to support.

The HOB applet is similar to any other Java applet accessed through the system or uploaded to the system. You must install a code-signing certificate to avoid seeing a warning similar to "This applet was signed by "Pulse Connect Secure" but Java cannot verify the authenticity of the signature's certificate. Do you trust this certificate?" Install a valid Applet signing certificate (JavaSoft) in the **Configuration > Certificates > Code-signing Certificates** window.

> NOTE: The HOB applet is for RDP connections and appears only for Windows Terminal Services. It is not applicable for Citrix Terminal Services profiles. The supported HOB version is 3.3.0754.

You can purchase HOB applets directly from HOB; however, Pulse Secure will support them only to the extent of uploading them. If you have any problems configuring or running the applet, you must contact HOB support.

To create a hosted Java applet autopolicy:

1. Create a terminal services resource profile.

2. Select **Enable Java support** within the resource profile.

3. Select the Java applet that you want to associate with the resource profile from the Applet to use list. Or, if the applet that you want to use is not currently available in the list, click Edit Applet. Then:

   1. Click **New Applet** to add an applet to this list. Or, select an existing applet and click **Replace** (to replace an existing applet with a new applet) or **Delete** (to remove an applet from the system).

      > NOTE: If you replace an existing archive, make sure that the new applet archive contains all of the necessary files for the applet to successfully launch

You can also add any additional HTML or JavaScript that you choose to this Web page definition. The system rewrites all of the code that you enter in this box.

Make sure to enter unique HTML in this box. If you create two bookmarks with the same HTML code, the system deletes one of the bookmarks in the end-user view. You can still see both bookmarks, however, in the administrator console.

For dynamic drive mapping to work with HOB Applet 3.3.0692, you must enable both the **AUTOLDM** and **TWAUTOMAPDRIVE** parameters. See the *Premier Java Applet Configuration Options* document located on the Pulse Secure support site for more details on these two parameters.

7. Select **Use this Java applet as a fallback mechanism** to use this applet only when the Windows client fails to launch. Or **select Always use this Java applet** to use this applet regardless of whether or not the Windows client launches.

8. Click **Save Changes**.

Related Documentation

- About Hosted Java Applet Templates
- Task Summary: Configuring the Terminal Services Feature
- Required Attributes for Uploaded Java Applets
- Required Parameters for Uploaded Java Applets

Defining a Bookmark for a Windows Terminal Services Profile

When you create a terminal services resource profile, the system automatically creates a bookmark that links to the terminal server that you specified in the resource profile. The system allows you to modify this bookmark as well as create additional bookmarks to the same terminal server.

To configure resource profile bookmarks for Windows terminal services:

1. In the admin console, select **Users > Resource Profiles > Terminal Services>***Resource Profile Name* **> Bookmarks**.

2. Click the appropriate link in the Bookmark column if you want to modify an existing session bookmark. Or, click **New Bookmark** to create an additional session bookmark.

   Although it is generally easiest to create a resource profile session bookmark through the resource profile configuration page, you can choose to create one through the user roles page as well.

3. (Optional.) Change the name and description of the session bookmark. (By default, the system populates and names the session bookmark using the resource profile name.)

4. Specify how the terminal emulation window should appear to the user during a terminal session by configuring options in the Settings area of the bookmark configuration page.

5. Pass user credentials from the system to the terminal server so that users can sign onto the terminal server without having to manually enter their credentials. You can do this by configuring options in the Session area of the bookmark configuration page.

6. Allow users to access specific applications on the terminal server by configuring options in the Start Application area of the bookmark configuration page. In addition, you can use settings in this area to define auto-launch and session reliability options.

7. Allow users to access local resources such as printers and drives through the terminal session by configuring options in the Connect Devices area of the bookmark configuration page.

8. Specify how the terminal emulation window should appear to the user during a terminal session by configuring options in the Desktop Settings area.

9. Specify the roles to which you want to display the session bookmarks if you are configuring the session bookmark through the resource profile pages, under Roles:

   - **ALL selected roles**—Displays the session bookmark to all of the roles associated with the resource profile.

   - **Subset of selected roles**—Displays the session bookmark to a subset of the roles associated with the resource profile. Then select roles from the ALL Selected Roles list and click **Add** to move them to the Subset of selected roles list.

10. Click **Save Changes**.

Related Documentation

- Task Summary: Configuring the Terminal Services Feature

- Defining SSO Options for the Windows Terminal Services Session

- Defining Application Settings for the Windows Terminal Services Session

- Defining Desktop Settings for the Windows Terminal Services Session

Creating a Windows Terminal Services Bookmark Through the User Roles Page

It is generally easiest to create a terminal services bookmark through the resource profile configuration pages. However, you can choose to create a resource profile session bookmark through the user roles page using the following instructions:

1. In the admin console, select **Users > User Roles > Select Role > Terminal Services> Sessions**.

2. Click **Add Session**.

3. Select **Terminal Services Resource Profile** from the Type list. (This option displays only after you have already created a terminal services resource profile.)

4. Select an existing resource profile that connects to a Windows terminal server on the system. (The system automatically populates the Host and Server Port boxes using settings from the selected resource profile.)

5. Click **OK**. (If you have not already associated the selected role with the resource profile, the system automatically makes the association for you. The system also enables any access control policies for the role that are required by the resource profile.)

6. If this role is not already associated with the selected resource profile, the system displays an informational message. If you see this message, click **Save Changes** to add this role to the resource profile's list of roles and to update the profile's autopolicies as required. Then, repeat the previous steps to create the session bookmark.

   When you create a resource profile session bookmark through the user roles page (instead of the standard resource profiles page), the system only associates the generated session bookmark with the selected role. The system does not assign the session bookmark to all of the roles associated with the selected resource profile.

7. (Optional.) Change the name and description of the session bookmark. By default, the resource profile name is used as the bookmark name.

8. Configure the bookmark's settings.

Related Documentation

- Task Summary: Configuring the Terminal Services Feature

Defining Display Options for the Windows Terminal Services Session

When configuring a terminal services bookmark, you can specify how the terminal emulation window should appear to users during their terminal sessions.

To define display options and auto-launch options:

1. Create a terminal services bookmark or edit an existing bookmark.

2. Scroll to the Settings area of the bookmark configuration page.

3. Select an option from the Screen Size drop-down list if you want to change the size of the terminal services window on the user's workstation. The default window size is full screen.

> *i* NOTE: If you select the Full Screen option and are connecting to a Windows terminal server, the system modifies the user's hosts file to display the correct hostname in the terminal services window. If the user does not have the proper rights to modify the hosts file, the system displays the loopback address instead.
>
> Also note that to restore the hosts file to its original state after running the terminal services window, the user must properly close his application. Otherwise, other applications that use the hosts file (such as JSAM and Host Checker) might not run properly. The user can also restore his hosts file to its original state by rebooting his system or by renaming the backup hosts file (**hosts_ive.bak**).

4. Select **8-bit**, **15-bit**, **16-bit**, **24-bit**, or **32-bit** color from the Color Depth list if you want to the change the color-depth of the terminal session data. The default color depth is 8-bit.

5. Click **Save Changes**.

Related Documentation

- Defining a Bookmark for a Windows Terminal Services Profile

Defining SSO Options for the Windows Terminal Services Session

When configuring a terminal services bookmark, you can configure the system to pass user credentials from the system to the terminal server so that the user does not have to manually enter his username and password. The system passes the specified credentials when a user clicks the session bookmark. If the credentials fail, the server prompts the user to manually enter his username and password.

To define single sign-on options:

1. Create a terminal services bookmark or edit an existing bookmark.

2. Scroll to the Authentication area of the bookmark configuration page.

3. Specify the username to pass to the terminal server. You can enter a static username or a variable. Enter the <username> variable to pass the username stored in the system's primary authentication server. Or use the following syntax to submit the username for the secondary authentication server: <username@SecondaryServerName> or <username[2]>.

4. Select Password if you want to specify a static password or select Variable Password if you want use the password stored in the system's primary or secondary authentication server. To use the password from the primary authentication server, enter the <password> variable. Or use the following syntax to submit the password for the secondary authentication server: <Password@SecondaryServerName> or <Password[2]>.

5. Click **Save Changes**.

Related Documentation

- Defining a Bookmark for a Windows Terminal Services Profile

- Task Summary: Configuring the Terminal Services Feature

Defining Application Settings for the Windows Terminal Services Session

When configuring a terminal services bookmark, you can specify that users can only access specific applications on the terminal server.

To define applications that users can access:

1. Create a terminal services bookmark or edit an existing bookmark.

2. Scroll to the Start Application area of the bookmark configuration page.

3. Select the **Launch seamless window** check box to have the Windows application server manage the display of the application. This allows an application's windows to behave in the same way as an application running on a Windows application server, regardless of the user's desktop environment.

> **NOTE:** If SSO is not configured, seamless window is supported only on Remote Desktop Protocol (RDP) 6.0 and later.
>
> The Launch seamless window check box is applicable only for servers running Windows 2008 and later.

Enter the server alias name (applicable only for servers running Windows 2008 and later) in the Alias name box.

4. Specify where the application's executable file resides on the terminal server in the Path to application box (visible only when you clear Launch seamless window). For example, you might enter the following directory for the Microsoft Word application:

    C:\Program Files\Microsoft Office\Office10\WinWord.exe

5. Specify where the terminal server should place working files for the application in the Working directory box. For example, you might specify that Microsoft Word should save files to the following directory by default:

    C:\Documents and Settings\username\My Documents

> **NOTE:** You can use session variables such as <username> and <password> n the Path to application and Working directory boxes. For example, when specifying an application path, you might want to include the <username> variable to personalize the location. For example: C:\Documents and Settings\<username>\My Documents.

6. Select the **Auto-launch** check box if you want to automatically launch this Terminal Service session bookmark when users sign into device. When you select this option, the system launches the terminal services application in a separate window after the user signs in.

7. Click **Save Changes**.

Related Documentation

- Defining a Bookmark for a Windows Terminal Services Profile

- Using Custom Expressions in Rule Configuration

Defining Device Connections for the Windows Terminal Services Session

When configuring a terminal services bookmark, you can specify local resources that users can access through the terminal session.

> **NOTE:** The system does not support providing users access to local resources when intermediating traffic using Java applets. Therefore, if you select the Enable Java Applets option when creating a Windows Terminal Services resource profile, note that the Connect Devices options described below might not work.
>
> When you enable local resources through the terminal server, each user can only access his own local resources. For instance, user 1 cannot see user 2's local directories.

To define local resources that users can access:

1. Create a terminal services bookmark or edit an existing bookmark.

2. Scroll to the Connect Devices area of the bookmark configuration page.

3. Select **Connect local drives** to connect the user's local drive to the terminal server, enabling the user to copy information from the terminal server to his local client directories.

4. Select **Connect local printers** to connect the user's local printers to the terminal server, enabling the user to print information from the terminal server to his local printer.

5. Select **Connect COM Ports** to connect the user's COM ports to the terminal server, allowing communication between the terminal server and the devices on his serial ports.

6. Select **Allow Clipboard Sharing** to allow the contents of the clipboard to be shared between the user's host computer and the terminal server. Because of limitations in RDP client earlier than version 6.0, clearing the Allow Clipboard Sharing option will automatically disable the Connect local drives, Connect local printers, and Connect COM Ports options.

7. Select **Connect smart cards** to allow users to use smart cards to authenticate their remote desktop sessions.

   > **NOTE:** Smart cards are supported by Microsoft Remote Desktop Protocol versions 5.1 and later.

8. Select **Sound Options** to enable sound during the remote session. Select Bring to this computer to redirect audio to the local computer. Select **Leave at remote computer** to play the audio only at the server.

   > **NOTE:** Sound options are supported by Microsoft Remote Desktop Protocol versions 5.1 and later.

9. Click **Save Changes**.

Related Documentation

- Defining a Bookmark for a Windows Terminal Services Profile

Defining Desktop Settings for the Windows Terminal Services Session

When configuring a terminal services bookmark, you can specify how the terminal emulation window should appear to the user during a terminal session.

> *ⓘ* NOTE: The options in this topic only apply to Windows Terminal Services bookmarks.

To define display settings for the users' sessions:

1. Create a terminal services bookmark or edit an existing bookmark

2. Scroll to the Display Settings area of the bookmark configuration page.

3. Select **Desktop background** to display a wallpaper background to users. If you do not select this option, the background is blank.

4. Select **Show contents of window while dragging** to show the contents of the Windows Explorer window while users move the windows on their desktops.

5. Select **Menu and window animation** to animate the movement of windows, menus, and lists.

6. Select **Themes** to allow users to set Windows themes in their terminal server windows.

7. Select **Bitmap Caching** to improve performance by minimizing the amount of display information that is passed over a connection.

8. Select **Font Smoothing (RDP 6.0 onwards)** to make text smoother and easier to read. This option only works on Windows Vista computers running RDP clients that are version 6.0 or later.

9. Select **Desktop Composition (RDP 6.0 onwards)** to allow desktop composition. With desktop composition, individual windows no longer draw directly to the screen. Instead, their drawing is redirected to video memory, which is then rendered into a desktop image and presented on the display.

10. Click **Save Changes**.

Related Documentation

- Defining a Bookmark for a Windows Terminal Services Profile

Creating a Citrix Terminal Services Resource Profile Using Default ICA Settings

This topic describes how to configure access to a Citrix Metaframe server using a default ICA configuration file.

To create a Citrix Terminal Services resource profile that uses default ICA settings:

1. In the admin console, select **Users > Resource Profiles > Terminal Services**.

2. Click **New Profile**. Or select an existing profile from the list.

3. Select **Citrix using default ICA** from the Type list.

4. (Existing resource profiles only) If you want to customize the default ICA file that comes with the system, click the **Open** link, customize the file, and upload it.

5. Enter a unique name and optionally a description for the resource profile. (This name becomes the default session bookmark's name.)

6. Specify the server and port to which this resource profile should connect in the Host box. When entering the server, you may enter a hostname or IP address.

7. Enter the port on which the terminal server listens in the Server Port field. (By default, the system populates this field with port number 1494 for Citrix.)

8. Select the **Create an access control policy allowing Terminal Service access to this server** check box to enable access to the server specified in the Server Port box (enabled by default).

9. Enable intermediation using a Java client by selecting **Enable Java support** and then specifying which Java client the system should use.

10. Click **Save and Continue**.

11. Select the roles to which the resource profile applies in the Roles tab and click Add.

    The selected roles inherit the autopolicy and session bookmarks created by the resource profile. If it is not already enabled, the system also automatically enables the Terminal Services option in the Users > User Roles > Select_Role > General > Overview page of the admin console for all of the roles you select.

12. Click **Save Changes**.

13. (Optional.) Modify the default session bookmark created by the system in the Bookmarks tab and/or create new ones. (By default, the system creates a session bookmark to the server defined in the Host box and displays it to all users assigned to the role specified in the Roles tab.)

Related Documentation

- Using Custom Expressions in Rule Configuration

- Creating a Citrix Resource Profile That Uses a Custom ICA File

- Defining a Hosted Java Applet Autopolicy

Defining a Bookmark for a Citrix Profile Using Default ICA Settings

When you create a Terminal Services resource profile, the system automatically creates a bookmark that links to the terminal server that you specified in the resource profile. The system enables you to modify this bookmark as well as create additional bookmarks to the same terminal server.

To configure resource profile bookmarks for Citrix Terminal Services using default ICA settings:

1. In the admin console, select **Users > Resource Profiles > Terminal Services> Select Resource Profile > Bookmarks**.

2. Click the appropriate link in the Bookmark column if you want to modify an existing session bookmark. Or, click **New Bookmark** to create an additional session bookmark.

   NOTE: Although it is generally easiest to create a resource profile session bookmark through the resource profile configuration page, you can choose to create one through the user roles page as well.

3. (Optional.) Change the name and description of the session bookmark. (By default, the system populates and names the session bookmark using the resource profile name.)

4. Specify how the terminal emulation window should appear to the user during a terminal session use configuration options in the Settings area of the bookmark configuration page.

5. Pass user credentials from the system to the terminal server so users can sign onto the terminal server without having to manually enter their credentials. You can do this by using the configuration options in the Session area of the bookmark configuration page.

6. Allow users to access specific applications on the terminal server by using configuration options in the Start Application area of the bookmark configuration page. In addition, you can use settings in this section to define auto-launch and session reliability options.

7. Allow users to access local resources such as printers and drives through the terminal session by using the configuration options in the Connect Devices section of the bookmark configuration page.

8. Specify the roles to which you want to display the session bookmark if you are configuring the session bookmark through the resource profile pages, under Roles:

   - **ALL selected roles**—Displays the session bookmark to all of the roles associated with the resource profile.

   - **Subset of selected roles**—Displays the session bookmark to a subset of the roles associated with the resource profile. Then select roles from the ALL selected roles list and click Add to move them to the Subset of selected roles list.

9. Click **Save Changes**.

Creating a Citrix Terminal Services Bookmark Through the User Roles Page

It is generally easiest to create a terminal services bookmark through the resource profile configuration pages, as explained in the previous topic. However, you can choose to create a resource profile session bookmark through the user roles page using the following instructions:

1. In the admin console, select **Users > User Roles > Select_Role > Terminal Services> Sessions**.

2. Click **Add Session**.

3. Choose **Terminal Services Resource Profile** from the Type list. (The system does not display this option if you have not already created a terminal services resource profile.)

4. Select an existing resource profile that connects to a Citrix server using the default ICA file. (The system automatically populates the Host and Server Port fields using settings from the selected resource profile.)

5. Click **OK**. (If you have not already associated the selected role with the resource profile, the system automatically makes the association for you. The system also enables any access control policies for the role that are required by the resource profile.)

6. If this role is not already associated with the selected resource profile, the system displays an informational message. If you see this message, click **Save Changes** to add this role to the resource profile's list of roles and to update the profile's autopolicies as required. Then, repeat the previous steps to create the session bookmark.

   > *i* NOTE: When you create a resource profile session bookmark through the user roles page (instead of the standard resource profiles page), the system only associates the generated session bookmark with the selected role. The system does not assign the session bookmark to all of the roles associated with the selected resource profile.

7. (Optional.) Change the name and description of the session bookmark. By default, the resource profile name is used as the session bookmark name.

8. Configure the bookmark's settings.

Related Documentation

- About Terminal Services Resource Profiles

- Defining Display Options for the Citrix Terminal Services Session

- Defining SSO Options for the Citrix Terminal Services Session

- Defining Application, Auto-Launch, and Session Reliability Settings for the Citrix Terminal Services Session

- Defining Device Connections for the Citrix Terminal Services Session

Defining Display Options for the Citrix Terminal Services Session

When configuring a terminal services bookmark, you can specify how the terminal emulation window should appear to users during their terminal sessions.

To define display, auto-launch, and session reliability options:

1. Create a terminal services bookmark or edit an existing bookmark.

2. Scroll to the Settings area of the bookmark configuration page.

3. Select an option from the Screen Size drop-down list if you want to change the size of the terminal services window on the user's workstation. The default window size is full screen.

4. Select **8-bit**, **15-bit**, **16-bit**, **24-bit**, or **32-bit** color from the Color Depth list if you want to the change the color-depth of the terminal session data. The default color depth is 8-bit.

   NOTE: When configuring a Citrix session bookmark, note that the setting you choose here and the user's local desktop setting both affect the client's color-depth display. If these settings do not match, the user sees the lower of the two color-depths during his session. For example, if you select 16-bit color during configuration, but the user's local desktop is set to 8-bit, the user sees 8-bit color depth during his session.

5. Click **Save Changes**.

Related Documentation

- Defining a Bookmark for a Citrix Profile Using Default ICA Settings

Defining SSO Options for the Citrix Terminal Services Session

When configuring a terminal services bookmark, you can configure the system to pass user credentials from the system to the terminal server so that the user does not have to manually enter his username and password. The system passes the specified credentials when a user clicks the session bookmark. If the credentials fail, the server prompts the user to manually enter his username and password.

To define single sign-on options:

1. Create a terminal services bookmark or edit an existing bookmark.

2. Scroll to the Authentication area of the bookmark configuration page.

3. Specify the username to pass to the terminal server in the Username field. You can enter a static username or a variable. Enter the <username> variable to pass the username stored in the system's primary authentication server. Or use the following syntax to submit the username for the secondary authentication server: <username@SecondaryServerName> or <username[2]>.

4. Select **Password** if you want to specify a static password or select Variable Password if you want to use the password stored in the system's primary or secondary authentication server. To use the password from the primary authentication server, enter the <password> variable.

Or use the following syntax to submit the password for the secondary authentication server: <Password@SecondaryServerName> or <Password[2]>.

5. (Default ICA file and listed applications only.) Select **Use domain credentials** to pass the user's cached domain credentials to the Citrix Metaframe server (also called pass-through authentication). When you select this option, the system uses the Citrix Program Neighborhood client to intermediate the Citrix terminal session.

> **NOTE:** If you want to download the Program Neighborhood client, select **Users > User Roles > Select_Role > Terminal Services > Options** in the admin console and enter the URL in the Download from URL box. See the Citrix web site for the location of the latest Program Neighborhood client cab file.

When you select the **Use domain credentials** option, you must also enable SSO through the user's settings file (appsrv.ini). If the user has already successfully signed into the Metaframe server using cached domain credentials, this setting should already be enabled. Otherwise, you or the user must:

- **Set EnableSSOnThruICAFile=On** in appsrv.ini. You can locate appsrv.ini in the %HOMEPATH%\Application Data\ICAClient directory.

- **Set UseLocalUserAndPassword=On** in the ICA file.

If you have not enabled SSO through the INI file, the user is prompted to manually enter his credentials when he tries to access the Metaframe server through the system.

6. Click **Save Changes**.

Related Documentation

- Defining a Bookmark for a Citrix Profile Using a Custom ICA File

Defining Application, Auto-Launch, and Session Reliability Settings for the Citrix Terminal Services Session

When configuring a terminal services bookmark, you can specify that users can only access specific applications on the terminal server.

To define applications that users can access:

1. Create a terminal services bookmark or edit an existing bookmark.

2. Scroll to the Start Application area of the bookmark configuration page.

3. Select the **Launch seamless window** check box to have the Windows application server manage the display of the application. This allows an application's windows to behave in the same way as an application running on a Windows application server, regardless of the user's desktop environment.

> **NOTE:** If SSO is not configured, seamless window is supported only on Remote Desktop Protocol (RDP) 6.0 and later.

Enter the server alias name in the Alias Name field (applicable only for servers running Windows 2008 and later).

4. Specify where the application's executable file resides on the terminal server in the Path to application box (visible only when you clear Launch seamless window). For example, you might enter the following directory for the Microsoft Word application:

C:\Program Files\Microsoft Office\Office10\WinWord.exe

5. Specify where the terminal server should place working files for the application in the Working directory field. For example, you might specify that Microsoft Word should save files to the following directory by default:

C:\Documents and Settings\<username>\My Documents

> **(i)** NOTE: You can use system session variables such as <username> and <password> in the Path to application and Working directory boxes. For example, when specifying an application path, you might want to include the <username> variable to personalize the location. For example, C:\Documents and Settings\<username>\My Documents.

6. Select the **Auto-launch** check box if you want to automatically launch this terminal service session bookmark when users sign into the device. When you select this option, the system launches the terminal services application in a separate window when the user signs in.

7. Select **Session Reliability and Auto-client reconnect** to keep ICA sessions active and on the user's screen when network connectivity is interrupted. Users continue to see the application they are using until the network connectivity resumes or the session reliability time-out has expired (the time-out value is defined by the Citrix product). Enter the port to use in the Port to be enabled box.

8. Click **Save Changes**.

Related Documentation

- Using Custom Expressions in Rule Configuration

Defining Device Connections for the Citrix Terminal Services Session

When configuring a terminal services bookmark, you can specify local resources that users can access through the terminal session.

For the Connect Devices settings to take effect, they must also be enabled on the Metaframe server. For example, if you enable Connect Drives on the system, but disable it on the Metaframe server, then the Metaframe server will block access to local drives. Note that if you clear the Configure access to local resources check box, the settings on the Metaframe server take effect.

To define local resources that users can access:

1. Create a terminal services bookmark or edit an existing bookmark.

2. Scroll to the Connect Devices area of the bookmark configuration page.

3. Select **Connect local drives** to connect the user's local drive to the terminal server, enabling the user to copy information from the terminal server to his local client directories.

4. Select **Connect local printers** to connect the user's local printers to the terminal server, enabling the user to print information from the terminal server to his local printer.

5. Select **Connect COM Ports** to connect the user's COM ports to the terminal server, allowing communication between the terminal server and the devices on his serial ports.

When you enable local resources through the terminal server, each user can only access his own local resources. For instance, user 1 cannot see user 2's local directories.

6. Click **Save Changes**.

Related Documentation

- Defining a Bookmark for a Citrix Profile Using Default ICA Settings

- Defining a Bookmark for a Windows Terminal Services Profile

Creating a Citrix Resource Profile That Uses a Custom ICA File

Use this type of resource profile to enable a terminal session to a Citrix Metaframe server using settings that you specify in a customized ICA file. Use custom ICA files to enable terminal sessions to Citrix Metaframe servers or NFuse servers governing Citrix server farms (in other words, to load balance). You may also use custom ICA files to link to single servers, if necessary. When you select this option, the system uses the session parameters defined in the specified custom ICA file.

To enable the connection between the system and the Citrix server farm, you must use the TCP/IP+HTTP protocol for browsing and specify the Citrix Metaframe or NFuse server port and IP address. The system does not support UDP port-forwarding.

To create a Citrix resource profile that uses a custom ICA file:

1. In the admin console, select **Users > Resource Profiles > Terminal Services**.

2. Click **New Profile**. Or select an existing profile from the list.

3. Select **Citrix using custom ICA file** from the Type list.

4. Specify the ICA file that contains the session parameters that you want use in the Custom ICA File box. Note that you may download and customize the following ICA files from the system:

   - ICA file that comes with the system—To customize this file, click the **Open** link, save the file to your local machine, customize the file as required, and upload it back to the system using the Browse option. If you customize this file, you must replace the following parameters in the default.ica file: <CITRIX_CLIENT_NAME>, <APPDATA> and <TARGET_SERVER>.

   - ICA file that you have already associated with the resource profile—To customize this file, click the **Current ICA File** link, save the file to your local machine, and customize the file as required. Once you make changes, you must upload the revised version using the Browse option.

     Before uploading the ICA file, you should test it to make sure it initiates the Citrix session correctly. To test, create an ICA file and access it directly. If the file displays the Citrix session correctly then it should work through the system.

     If SSO is configured in the custom ICA bookmark, seamless mode is ignored and the application is launched in non-seamless mode.

     When using the Java rewriting technology to tunnel Citrix JICA applets through the system, you must set the proxyType parameter in the ICA file to None (even if a client-side proxy is configured in the browser).

5. Enter a unique name and optionally a description for the resource profile. (This name becomes the default session bookmark's name.)

6. Enable access to the servers specified in the custom ICA file:

   1. Select the **Autopolicy: Terminal Services Access Control** check box.

   2. Specify the Metaframe servers to which you want to enable access in the Resource field.

   3. Choose **Allow** to enable access to the specified resource or **Deny** to block access to the specified resource from the Action list.

   4. Click **Add**.

5. Enable intermediation using a Java client by selecting Enable Java support.

   If you select the **Enable Java support** option and have a custom ICA file that you uploaded to the system, your HTML file is auto-populated with references to your custom ICA file. No additional HTML code needs to be added.

6. Click **Save and Continue**.

7. Select the roles to which the resource profile applies in the Roles box and click **Add**.

   The selected roles inherit the autopolicy and session bookmarks created by the resource profile. If it is not already enabled, the system also automatically enables the Terminal Services option in the Users > User Roles > Select_Role > General > Overview page of the admin console for all of the roles you select.

8. Click **Save Changes**.

9. (Optional) Modify the default session bookmark created by the system in the Bookmarks tab and/or create new ones. (By default, the system creates a session bookmark to the server defined in your custom ICA file and displays it to all users assigned to the role specified in the Roles tab.)

Related Documentation

- Defining a Hosted Java Applet Autopolicy

- Defining a Bookmark for a Citrix Profile Using a Custom ICA File

Defining a Bookmark for a Citrix Profile Using a Custom ICA File

When you create a terminal services resource profile, the system automatically creates a bookmark that links to the terminal server that you specified in the resource profile. You can modify this bookmark as well as create additional bookmarks to the same terminal server.

To configure resource profile bookmarks for Citrix Terminal Services using custom ICA settings:

1. In the admin console, select **Users > Resource Profiles > Terminal Services> Select_Resource_Profile > Bookmarks**.

   Click the appropriate link in the Bookmark column if you want to modify an existing session bookmark. Or, click **New Bookmark** to create an additional session bookmark.

   Although it is generally easiest to create a resource profile session bookmark through the resource profile configuration page, you can choose to create one through the user roles page as well.

2. (Optional.) Change the name and description of the session bookmark. (By default, the system populates and names the session bookmark using the resource profile name.)

3. Pass user credentials from the system to the terminal server so that users can sign onto the terminal server without having to manually enter their credentials. You can do this by configuring options in the Session area of the bookmark configuration page.

4. Automatically launch this terminal service session bookmark when a user signs in to the device by selecting the **Auto-launch** check box. When you select this option, the system launches the terminal services application in a separate window when the user signs in.

5. Under Roles, specify the roles to which you want to display the session bookmark:

   - **ALL selected roles**—Displays the session bookmark to all of the roles associated with the resource profile.

- **Subset of selected roles**—Displays the session bookmark to a subset of the roles associated with the resource profile. Then select roles from the ALL selected roles list and click **Add** to move them to the Subset of selected roles list.

6. Click **Save Changes**.

Related Documentation

- About Terminal Services Resource Profiles

- Defining SSO Options for the Citrix Terminal Services Session

Creating a Citrix Profile That Lists Published Applications

Citrix created published applications to satisfy the need for security. It is dangerous to allow any executable to be run on the server. With published applications, only applications that are allowed to be run are published.

These published applications are displayed on the system index page as terminal services bookmarks.

> **NOTE:** The Citrix Desktop Toolbar Viewer is enabled only for XenDesktop. It is not enabled for XenApp. If you require the Citrix Desktop Toolbar Viewer, use the XenDesktop configuration on Connect Secure. Do not configure a desktop as part of the Citrix Listed Applications feature.

To create a Citrix profile that lists published applications:

1. In the admin console, select **Users > Resource Profiles > Terminal Services**.

2. Click **New Profile**.

3. Select **Citrix Listed Applications** from the Type list.

4. Enter a unique name and optionally a description for the resource profile. This name becomes the default session bookmark's name.

5. Enter the IP address and port of the Citrix MetaFrame server where the XML service is running.

   You do not need to enter the port number if you are using the default value. The default port is 80 (if SSL is selected, the default port is 443).

   You can enter more than one server. If the connection fails on one server, the next server in the list is used.

6. Click the **Use SSL for connecting to Citrix XML Service** check box to send the password through SSL instead of cleartext.

   > **NOTE:** Although cleartext is supported, we recommend you always use SSL to avoid any security issues.

7. Enter the username, password, and domain name for connecting to the Citrix Metaframe server where the XML service is running.

   You can enter variable credentials such as <USERNAME> and <PASSWORD>. If you use variable credentials, the Subset of selected Applications option is disabled in the Bookmarks window.

   When the user accesses the application list, their credentials are submitted to the Citrix XML service, substituting the session context variables <USERNAME> and <PASSWORD>. Only the user's specific applications (as determined by the Citrix administrator) are returned.

8. Enable access to the servers specified in the custom ICA file:

   1. Select the **Autopolicy: Terminal Services Access Control** check box.

   2. Specify the Metaframe servers to which you want to enable access in the Resource field.

   3. Choose **Allow** to enable access to the specified resource or **Deny** to block access to the specified resource from the Action list.

   4. Click **Add**.

5. Enable intermediation using a Java client by selecting **Enable Java support** and then specifying which Java client to use.

6. Click **Save and Continue**.

7. Select the roles to which the resource profile applies in the Roles tab and click Add.

   The selected roles inherit the autopolicy and session bookmarks created by the resource profile. If it is not already enabled, the system also automatically enables the Terminal Services option in the Users > User Roles > Select_Role > General > Overview page of the admin console for all of the roles you select.

8. Click **Save Changes**.

9. (Optional.) Modify the default session bookmark created by the system in the Bookmarks box and/or create new ones.

Related Documentation

- Defining a Bookmark for a Citrix Profile Listing Applications
- Using Custom Expressions in Rule Configuration

Defining a Bookmark for a Citrix Profile Listing Applications

When you create a terminal services resource profile, the system automatically creates a bookmark that links to the terminal server that you specified in the resource profile. You can modify this bookmark as well as create additional bookmarks to the same terminal server.

To configure resource profile bookmarks for Citrix terminal services list applications:

1. In the admin console, select **Users > Resource Profiles > Terminal Services> Resource_Profile > Bookmarks**.

   Click the appropriate link in the Bookmark column if you want to modify an existing session bookmark. Or, click **New Bookmark** to create an additional session bookmark.

   Although it is generally easiest to create a resource profile session bookmark through the resource profile configuration page, you can choose to create one through the user roles page as well.

2. (Optional.) Change the name and description of the session bookmark. By default, the resource profile name is used as the session bookmark name.

3. Under Applications, select the applications you want available to the end user.

   - **ALL Applications**—Allow all executables on the server to be available to the end user.

   - **Subset of selected applications**—Select the executables from the Available list and click **Add** to allow only those applications to be run. The Available list is automatically populated from the Metaframe server.

     This option is disabled when you enter variable credentials, such as <USERNAME> and <PASSWORD> while defining the resource profile.

4. Under Settings, specify how the terminal emulation window should appear to users during their terminal sessions.

> ℹ️ NOTE: You cannot change the IP address or XML Service running port for connecting to the XML Service or the Java client to use for intermediation.

- Select an option from the Screen Size drop-down list if you want to change the size of the terminal services window on the user's workstation.

- (Optional.) Select **8-bit**, **15-bit**, **16-bit**, **24-bit**, or **32-bit** color from the Color Depth list if you want to the change the color-depth of the terminal session data.

5. Under Session, you can configure the system to pass user credentials from the system to the terminal server so that the user does not have to manually enter his username and password.

- Specify the username to pass to the terminal server in the Username box. You can enter a static username or a variable.

- Select **Password** if you want to specify a static password or select **Variable** Password if you want to use the password stored in the system's primary or secondary authentication server.

- Select **Use domain credentials** to pass the user's cached domain credentials to the Citrix Metaframe server (also called pass-through authentication). When you select this optionthe system uses the Citrix Program Neighborhood client to intermediate the Citrix terminal session.

  > ℹ️ NOTE: If you want to download the Citrix Program Neighborhood client, select **Users > User Roles > _Role Name_ > Terminal Services > Options** of the admin console and enter the following URL in the Download from URL box: http://download2.citrix.com/FILES/en/products/client/ica/client9230/wficat.cab

  When you select the Use domain credentials option, you must also enable SSO through the user's settings file (appsrv.ini).

6. Under Connect Devices, specify which user devices to connect to the terminal server.

- **Connect local drives**—Connect the user's local drive to the terminal server, enabling the user to copy information from the terminal server to his local client directories.

- **Select Connect local printers**—Connect the user's local printers to the terminal server, enabling the user to print information from the terminal server to his local printer.

- **Select Connect COM Ports**—Connect the user's COM ports to the terminal server, allowing communication between the terminal server and the devices on his serial ports.

7. Under Roles, specify the roles to which you want to display the session bookmark:

8. Click **Save Changes**.

Related Documentation

- Creating Resource Profiles Using Citrix Web Applications

- Defining SSO Options for the Citrix Terminal Services Session

Creating Session Bookmarks to Your Terminal Server

When you enable the Terminal Services option through the admin console, you can create session bookmarks to your terminal server. A terminal services session bookmark defines information about the terminal server to which users can connect and (optionally) applications

that they can use on the terminal server. The session bookmarks that you define appear on the Terminal Services panel in the end-user console for users who map to the appropriate role.

You can use two different methods to create terminal services session bookmarks:

- Create session bookmarks through existing resource profiles (recommended)—When you select this method, the system automatically populates the session bookmark with key parameters (such as the session type) using settings from the resource profile. Additionally, while you are creating the associated resource profile, the system guides you through the process of creating any required policies to enable access to the session bookmark.

- Create standard session bookmarks—With this option, you must manually enter all session bookmark parameters during configuration. Additionally, you must enable access to the Terminal Services feature and create resource policies that enable access to the servers defined in the session bookmark.

> **NOTE:** If you enable the Terminal Services option but do not give users the ability to create their own session bookmarks, make sure that you configure session bookmarks for them. Otherwise, users cannot use this feature.

You can also enable users to create their own session bookmarks on the homepage and browse to the terminal servers using the system browse bar. Or, you can create links from external sites to a terminal services bookmarks.

Related Documentation

- Defining a Bookmark for a Windows Terminal Services Profile
- Defining a Bookmark for a Citrix Profile Using a Custom ICA File
- Creating Advanced Terminal Services Session Bookmarks
- Creating Links from an External Site to a Terminal Services Session Bookmark
- Specifying General Terminal Services Options

Creating Advanced Terminal Services Session Bookmarks

The information in this topic is provided for backwards compatibility. We recommend that you configure access to Windows terminal servers and Citrix servers through resource profiles instead, because they provide a simpler, more unified configuration method. Resource profile also contain features (such as the ability to use Java RDP clients to support Macintosh and Linux users) which are not available through roles.

Make sure to enter a unique set of parameters when defining a terminal services bookmark. If you create two bookmarks that contain the same set of parameters, the system deletes one of the bookmarks from the end user view. You can still see both bookmarks in the administrator console.

To create a session bookmark for terminal sessions:

1. In the admin console, select **Users > User Roles > Role > Terminal Services > Sessions**.

2. Click **Add Session**.

3. Select **Standard** in the Type drop-down list.

4. Specify the type of user session you want to create from the Session Type list:

   - **Windows Terminal Services**—Enables a terminal session to a Windows terminal server.

- **Citrix using default ICA**—Enables a terminal services session to a Citrix Metaframe server. When you select this option, the system uses the default Citrix session parameters.

  (Existing sessions only.) You can also use the Open link to open the system's default ICA file, which you can then save to your local machine and customize as required. If you customize this file, you must replace the following parameters in the default.ica file: <CITRIX_CLIENT_NAME>, <APPDATA>, and <TARGET_SERVER>.

- **Citrix using custom ICA file**—Enables a terminal services session to a Citrix Metaframe or NFuse server governing a Citrix server farm. When you select this option, the system uses the session parameters defined in the specified custom ICA file, thus removing the Session Reliability, Start Application, and Connect Devices configuration items from the current page.

  > **NOTE:** Because the system does not support UDP port-forwarding, you must use the TCP/IP+HTTP protocol for browsing and specify the Citrix Metaframe or NFuse server port and IP address to enable the connection between Connect Secure and the Citrix server farm.

5. Enter a name and (optionally) a description for the session bookmark.

6. In the Host field, specify the hostname or IP address of the Windows terminal server or Metaframe terminal server.

7. In the Client Port and Server Port fields, enter the ports on which the user client communicates and terminal server listens.

   If you specify a client port and the Juniper terminal services client is unable to bind to this port, then the terminal services client will fail. However, if you leave the Client Port field blank, the Juniper terminal services client dynamically selects an available port.

8. (Windows Terminal Services and Citrix using default ICA only) If you want to specify the screen size and color depth options for the terminal emulation window, use configuration options in the Settings section.

9. If you want to pass user credentials from the system to the terminal server, enabling users to sign onto the terminal server without having to manually enter their credentials, use configuration options in the Session section.

10. If you only want to allow users to access specific applications on the terminal server, use configuration options in the Start Application section of the bookmark configuration page. In addition, you can use settings in this section to define auto-launch and session reliability options.

11. (Windows Terminal Services and Citrix using default ICA only) If you want to allow users to access local resources such as printers and drives through the terminal session, use configuration options in the Connect Devices section of the bookmark configuration page.

12. (Windows Terminal Services only) If you want to specify how the terminal emulation window should appear to the user during a terminal session, use configuration options in the Desktop Settings section.

13. Click **Save Changes** or **Save + New**.

Defining Screen Size and Color Depth Options for the Terminal Services Session

When configuring a terminal services bookmark, you can specify how the terminal emulation window should appear to users during their terminal sessions.

The options in this section only apply to Windows Terminal Services bookmarks, Citrix using default ICA bookmarks and Citrix listed applications bookmarks.

To define display, auto-launch, and session reliability options:

1. Create a terminal services session bookmark or edit an existing session bookmark.

2. Scroll to the Settings section of the bookmark configuration page.

3. Select an option from the Screen Size drop-down list if you want to change the size of the terminal services window on the user's workstation. The default window size is full screen.

   If you select the Full Screen option and are connecting to a Windows terminal server, the system modifies the user's hosts file in order to display the correct hostname in the terminal services window. If the user does not have the proper rights to modify the hosts file, the system displays the loopback address instead.

   Also note that in order to restore the hosts file to its original state after running the terminal services window, the user must properly close his application. Otherwise, other applications that use the hosts file (such as JSAM and Host Checker) might not run properly. The user can also restore his hosts file to its original state by rebooting his system or by renaming the backup hosts file (hosts_ive.bak).

4. Select a value from the Color Depth list if you want to the change the color-depth of the terminal session data. The default color depth is 8-bit.

   When configuring a Citrix session bookmark, note that the setting you choose here and the user's local desktop setting both affect the client's color-depth display. If these settings do not match, the user sees the lower of the two color-depths during his session. For example, if you choose 16-bit color during system configuration, but the user's local desktop is set to 8-bit, the user sees 8-bit color depth during his session.

5. Click **Save Changes** or **Save + New**.

Defining SSO Options for the Terminal Services Session

When configuring a terminal services bookmark, you can pass user credentials from the system to the terminal server so that the user does not have to manually enter his username and password. The system passes the specified credentials when a user clicks the session bookmark. If the credentials fail, the server prompts the user to manually enter his username and password.

To define single sign-on options:

1. Create a terminal services session bookmark or edit an existing session bookmark.

2. Scroll to the Authentication section of the bookmark configuration page.

3. In the Username field, specify the username to pass to the terminal server. You can enter a static username or a variable. Enter the <username> variable to pass the username stored in the system's primary authentication server. Or use the following syntax to submit the username for the secondary authentication server: <username@SecondaryServerName> or <username[2]>.

4. Select **Password** if you want to specify a static password or select **Variable Password** if you want use the password stored in the system's primary or secondary authentication server. To use the password from the primary authentication server, enter the <password> variable. Or use the following syntax to submit the password for the secondary authentication server: <Password@SecondaryServerName> or <Password[2]>.

5. (Citrix using default ICA or listed applications) Select **Use domain credentials** to pass the user's cached domain credentials to the Citrix Metaframe server (also called pass-through authentication). When you select this option, the system uses the Citrix Program Neighborhood client to intermediate the Citrix terminal session.

> **NOTE:** If you want to download the Program Neighborhood client, go to the Users > User Roles > Select Role > Terminal Services > Options page of the admin console and enter the following URL in the Download from URL field:
> http://download2.citrix.com/FILES/en/products/client/ica/client9230/wficat.cab

When you select the Use domain credentials option, you must also enable SSO through the user's settings file (appsrv.ini). If the user has already successfully signed into the Metaframe server using cached domain credentials, this setting should already be enabled. Otherwise, you or the user must:

- Set EnableSSOnThruICAFile=On in appsrv.ini. You can locate appsrv.ini in the %HOMEPATH%\Application Data\ICAClient directory.

  Set UseLocalUserAndPassword=On in the ICA file.

If you have not enabled SSO through the INI file, the user is prompted to manually enter his credentials when he tries to access the Metaframe server through the system.

6. Click **Save Changes** or **Save + New**.

Defining Application Settings for the Terminal Services Session

When configuring a terminal services bookmark, you can specify that users can only access specific applications on the terminal server. Additionally, you can define auto-launch and session reliability options for the session.

To define applications that users can access:

1. Create a terminal services session bookmark or edit an existing session bookmark.

2. Scroll to the Start Application section of the bookmark configuration page.

   If you specify Citrix using custom ICA file in the Session Type configuration section, the Start Application configuration item is not available.

3. (Windows Terminal Services and Citrix using default ICA only) In the Path to application field, specify where the application's executable file resides on the terminal server. For example, you might enter the following directory for the Microsoft Word application:

   C:\Program Files\Microsoft Office\Office10\WinWord.exe

4. (Windows Terminal Services and Citrix using default ICA only) In the Working directory field, specify where the terminal server should place working files for the application. For example, you might specify that Microsoft Word should save files to the following directory by default:

   C:\Documents and Settings\<username>\My Documents

   You can use session variables such as <username> and <password>in the Path to application and Working directory fields. For example, when specifying an application path, you might want to include the <username> variable to personalize the location. For example: C:\Documents and Settings\<username>\My Documents.

5. (Citrix using default ICA only) Select **Session Reliability and Auto-client reconnect** to keep ICA sessions active and on the user's screen when network connectivity is interrupted. Users continue to see the application they are using until the network connectivity resumes

or the session reliability time-out has expired (the time-out value is defined by the Citrix product). Enter the port to use in Port to be enabled.

6. Select the **Auto-**launch check box if you want to automatically launch this Terminal Service session bookmark when users sign into the device. When you select this option, the system launches the terminal services application in a separate window when the user signs in.

7. Click **Save Changes** or **Save + New**.

Defining Device Connections for the Terminal Services Session

When configuring a terminal services bookmark, you can specify local resources that users can access through the terminal session.

The options in this section only apply to Windows Terminal Services bookmarks and Citrix using default ICA bookmarks.

The Connect Devices options that you specify at the role-level control whether end users can enable or disable access to local resources when they configure their own bookmarks. These role-level options do not control whether users can access local resources through a bookmark created by the system administrator.

To define local resources that users can access:

1. Create a terminal services session bookmark or edit an existing session bookmark.

2. Scroll to the Connect Devices section of the bookmark configuration page.

   If you specify Citrix using custom ICA file in the Session Type configuration section, the Connect Devices configuration item is not available.

3. Select **Connect local drives** to connect the user's local drive to the terminal server, enabling the user to copy information from the terminal server to his local client directories.

4. Select **Connect local printers** to connect the user's local printers to the terminal server, enabling the user to print information from the terminal server to his local printer.

5. Select **Connect COM Ports** to connect the user's COM ports to the terminal server, allowing communication between the terminal server and the devices on his serial ports.

6. (Windows Terminal Services only) Select **Allow Clipboard Sharing** if you want to allow the contents of the clipboard to be shared between the user's host computer and the terminal server. Due to the limitations in the pre-6.0 versions of the RDP client, disabling the Allow Clipboard Sharing option will automatically disable the Connect local drives, Connect local printers, and Connect COM Ports options.

   When you enable local resources through the terminal server, each user can only access his own local resources. For instance, user 1 cannot see user 2's local directories.

7. (Windows Terminal Services only) Select **Connect smart cards** to allow users to use smart cards to authenticate their remote desktop sessions.

8. (Windows Terminal Services only) Select **Sound Options** to enable sound during the remote session. Choose Bring to this computer to redirect audio to the local computer. Choose Leave at remote computer to play the audio only at the server.

   ⓘ   NOTE: Smart cards and sound options are supported by Microsoft Remote Desktop Protocol versions 5.1 and later.

9. Click **Save Changes** or **Save + New**.

Defining Desktop Settings for the Terminal Services Session

When configuring a terminal services bookmark, you can specify how the terminal emulation window should appear to the user during a terminal session.

The options in this section only apply to Windows Terminal Services bookmarks.

To define display settings for the users' sessions:

1. Create a terminal services session bookmark or edit an existing session bookmark.

2. Scroll to the Display Settings section of the bookmark configuration page.

3. Select **Desktop background** if you want to display a wallpaper background to users. If you do not select this option, the background is blank.

4. Select **Show contents of window while dragging** if you want to show the contents of the Windows Explorer window while users move the windows on their desktops.

5. Select **Menu and window animation** if you want to animate the movement of windows, menus, and lists.

6. Select **Themes** if you want to allow users to set Windows themes in their terminal server windows.

7. Select **Bitmap Caching** if you want to improve performance by minimizing the amount of display information that is passed over a connection.

8. Select **Font Smoothing (RDP 6.0 onwards)** to make text smoother and easier to read. This option only works on Windows Vista computers running RDP clients that are version 6.0 or later.

9. Select **Desktop Composition (RDP 6.0 onwards)** to allow desktop composition. With desktop composition, individual windows no longer draw directly to the screen. Instead, their drawing is redirected to video memory, which is then rendered into a desktop image and presented on the display.

10. Click **Save Changes** or **Save + New**.

Related Documentation

- About Terminal Services Resource Profiles

- Using Custom Expressions in Rule Configuration

- Specifying the Terminal Services Resource Option

- Defining SSO Options for the Citrix Terminal Services Session

- Defining Application Settings for the Windows Terminal Services Session

- Defining Device Connections for the Windows Terminal Services Session

- Defining Desktop Settings for the Windows Terminal Services Session

- Multiple Sign-In Credentials Execution

Creating Links from an External Site to a Terminal Services Session Bookmark

When creating a link to a terminal services session bookmark from another site, you can construct either of the following types of URLs:

- URL that includes all necessary parameters—Create a URL that includes all of the parameters that you want to pass to the terminal services program, such as the host, ports, and terminal window parameters. When constructing the URL, use the following syntax:

https://<SASeriesAppliance>/dana/term/winlaunchterm.cgi?<param1>=<value1>&<param2>=<value2>

When constructing your URL, you can use the case-insensitive parameter names described in Table 92. If you want to include more than one parameter in the session bookmark, string them together using ampersand characters (&). For example:

https://YourSA.com/dana/term/winlaunchterm.cgi?host=yourtermserver.yourdomain.com&type=Windows&clientPort=1094&serverPort=3389&user=john&password=abc123&screenSize=fullscreen

- URL to terminal services bookmark—Create a URL that simply points to a session bookmark that you have already created on the system. When constructing the URL, use the following syntax:

https://<SASeriesAppliance>/dana/term/winlaunchterm.cgi?bmname=<bookmarkName>

Within the URL, only define the bmName parameter.

When using the system to host Terminal Services session bookmarks, you must:

- Enable the **User can add sessions** option in the Users > User Roles > Select Role > Terminal Services > Options page. If you do not select this option, users cannot link to the Terminal Services session bookmarks from external sites.

- Create a policy that prevents the system from rewriting the link and the page that contains the link using settings in the Users > Resource Policies > Web > Rewriting > Selective Rewriting page of the admin console.

Additionally, we recommend that you use https protocol instead of http. Otherwise, when users launch the session bookmark, they see an insecure site warning.

> 🛈 NOTE: If you create links on external servers to Terminal Services bookmarks on the system and you are using multiple customized sign-in URLs, some restrictions occur.

Table92: Case-Insensitive Terminal Services Session Bookmark Parameter Names

Parameter Name	Description and Possible Values	Example

host	Required. Hostname or IP address of the Windows terminal server or Metaframe terminal server.	https://YourSystem.com/dana/term/winlaunchterm.cgi?host=YourTermServer
type	Type of terminal server. Possible values include Windows or Citrix.	https://YourSystem.com/dana/term/winlaunchterm.cgi?host=YourTermServer&type=Windows
clientPort	Port on which the user client communicates.	https://YourSystem.com/dana/term/winlaunchterm.cgi?host=YourTermServer&clientPort=1094
serverPort	Port on which the terminal server listens. Default values are 3389 for Windows and 1494 for Citrix.	https://YourSystem.com/dana/term/winlaunchterm.cgi?host=YourTermServer&serverPort=3389
user	Username to pass to the terminal server. You can enter a static username, such as JDoe, or a variable username such as <user> or <username>	https://YourSystem.com/dana/term/winlaunchterm.cgi?host=YourTermServer&user=jDoe
password	Password to pass the terminal server.	https://YourSystem.com/dana/term/winlaunchterm.cgi?host=YourTermServer&user=jDoe&password=<password>
bmname	Specifies the session bookmark name	https://YourSystem.com/dana/term/winlaunchterm.cgi?bmname=<bookmarkname>

screenSize	Terminal services window's size. Possible values: • fullScreen • 800x600 • 1024x768 • 1280x1024	https://YourSystem.com/dana/term/winlaunchterm.cgi?host=YourTermServer&screenSize=full Screen
colorDepth	Terminal services window's color depth, in bits. Possible values: • 8 • 15 • 16 • 24 • 32	https://YourSystem.com/dana/term/winlaunchterm.cgi?host=YourTermServer&colorDepth=32
startApp	Specifies the path of an application executable to start.	https://YourSystem.com/dana/term/winlaunchterm.cgi?host=YourTermServer&startApp=C:\Program Files\Microsoft Office\Office10\WinWord.exe
startDir	Specifies where the terminal server should place working files for the application.	https://YourSystem.com/dana/term/winlaunchterm.cgi?host=YourTermServer&startapp=C:\temp
connectDrives	Specifies whether the user can connect his local drive to the terminal server. Possible values: • Yes • No	https://YourSystem.com/dana/term/winlaunchterm.cgi?host=YourTermServer&connectDrives=Yes

connectPrinters	Specifies whether the user can connect his local printer to the terminal server. Possible values: • Yes • No	https://YourSystem.com/dana/term/winlaunchterm.cgi?host=YourTermServer&connectPrinters=Yes
connectComPorts	Specifies whether the user can connect his COM ports to the terminal server. Possible values: • Yes • No	https://YourSystem.com/dana/term/winlaunchterm.cgi?host=YourTermServer&connectComPorts=Yes
allowclipboard	Specifies whether the user can share the contents of the clipboard between the user's host computer and the terminal server. Possible values: • Yes • No	https://YourSystem.com/dana/term/winlaunchterm.cgi?host=YourTermServer&allowclipboard=Yes
desktopbackground	Specifies whether to display your current wallpaper setting. Possible values: • Yes • No	https://YourSystem.com/dana/term/winlaunchterm.cgi?host=YourTermServer&desktopbackground=Yes

showDragConte nts	Specifies whether to show the contents of the Windows Explorer window while moving the window around your desktop. Possible values: • Yes • No	https://YourSystem.com/dana/term/winlaunchterm.cgi?host=YourTermServer&showDragCont ents=Yes
showMenuAnima tion	Specifies whether to animate the movement of windows, menus, and lists. Possible values: • Yes • No	https://YourSystem.com/dana/term/winlaunchterm.cgi?host=YourTermServer&showMenuAni mation=Yes
themes	Specifies whether to allow users to set Windows themes in their terminal server windows. Possible values: • Yes • No	https://YourSystem.com/dana/term/winlaunchterm.cgi?host=YourTermServer&themes=Yes
bitmapcaching	Specifies whether to improve performanc e by minimizing the amount of display information that must be passed over a connection. Possible values: • Yes • No	https://YourSystem.com/dana/term/winlaunchterm.cgi?host=YourTermServer&bitmapcaching =Yes

fontsmoothing	Specifies whether to make text smoother and easier to read. This option only works on Windows Vista computers running RDP clients that are version 6.0 or later. Possible values: • Yes • No	https://YourSystem.com/dana/term/winlaunchterm.cgi?host=YourTermServer&fontsmoothing=Yes
desktopcomposition	Specifies whether to enable desktop composition. Possible values: • Yes • No	https://YourSystem.com/dana/term/winlaunchterm.cgi?host=YourTermServer&desktopcomposition=Yes
soundoptions	Specifies whether to enable sound. Possible values: • 0—disable sound • 1—bring sound to this computer • 2—leave sound at remote computer	https://YourSystem.com/dana/term/winlaunchterm.cgi?host=YourTermServer&soundOptions=1

Related Documentation

- Creating Advanced Terminal Services Session Bookmarks
- Defining a Sign-In Policy

Specifying General Terminal Services Options

Users can create their own terminal services session bookmarks and can configure the system to create terminal services resource policies that enable access to the servers specified in the terminal services session bookmarks.

To specify general Terminal Services options:

1. In the admin console, choose **Users > User Roles > Role > Terminal Services > Options**.

2. If you are enabling Citrix sessions, under Citrix client delivery method, specify where the system should obtain the ICA client to download to users' systems:

   - **Download from the Citrix website**—The system installs the latest version of the ICA client from the Citrix web site. You can edit the URL to point to a new location if the one listed is no longer valid.

   - **Download from the IVE**—Use the Browse button to browse to the ICA client on your local network. You can upload a CAB, MSI or EXE file. Once you upload the client, the system uses it as the default and downloads it to your users' systems when necessary. You must also specify the exact version number of the ICA client.

     If you upload an MSI or EXE file, an open/save dialog box appears to download and install the client. If Java fallback is configured, you are given the option to bypass this download and use Java instead.

   - **Download from a URL**—The system installs the ICA client of your choice from the specified web site. You must also specify the exact version number of the ICA client. If Java fallback is configured, you are given the option to bypass this download and use Java instead.

     > **NOTE:** We recommend that you test the Citrix client that you expect the system to download with the custom ICA file that you have uploaded to the system. Perform this testing without the system to determine if the Citrix client supports the features in the custom ICA file. If the features do not work without the system, they will not work through the system either.

     If you choose to download an ICA client from the Citrix web site or a URL, the system secures the download transaction by processing the URL through the Content Intermediation Engine. Therefore, you must choose a site that is accessible by the system and enabled for users within this role.

     To determine if the ICA web client is already installed on a machine, check for the following entry in your Windows registry: HKEY_CLASSES_ROOT\CLSID\{238F6F83-B8B4-11CF-8771-00A024541EE3}

     You can determine the version number of an ICA client by extracting the cab file (for example, wfica.cab), looking for an inf file in the archive (for example, wficat.inf), and then locating the information about each ocx in the inf file. For example, wficat.inf (in wficat.cab) might contain the following information:

     ```
     [wfica.ocx]
     file-win32-x86=thiscab
     clsid={238F6F83-B8B4-11CF-8771-00A024541EE3}
     FileVersion=8,00,24737,0
     [wfica32.exe]
     file-win32-x86=thiscab
     FileVersion=8,00,24737,0
     ```

     In this case, "8,00,23737,0" is the file version. (Note that the version includes commas instead of periods.)

3. Enable the **User can add sessions** option to enable users to define their own terminal session bookmarks and to enable users to access terminal servers through the system browse bar on the home page. When you enable this option, the **Add Terminal Services**

Session button appears on the Terminal Services page the next time a user refreshes the user console.

4. Select the **Deny single sign-on for sessions added by user** option if you do not want the user Add Terminal Service Session page to include the Authentication section used for single sign-on. This setting is disabled by default. When enabled, it disallows SSO for all user-added terminal services sessions, even if the user had previously configured SSO authentication credentials when that was permitted. This option adds a security measure to protect against exploitation of a security breach. If SSO is allowed, an attacker who gains access to a user's home page could gain access to the terminal services added by the user.

5. Enable the **Auto-allow role Terminal Services** sessions option to enable the system to automatically enable access to the resources defined in the terminal session bookmark (rather than having to create resource policies). Note that this only applies to role bookmarks, not user bookmarks.

    You may not see the Auto-allow option if you are using a new installation or if an administrator hides the option.

6. If you want to allow users to enable access to local devices through the bookmarks they create, select from the following options in the Allow users to enable local resources defined below section:

    - **Users can connect drives**—Enables the user to create bookmarks that connect the his local drives to the terminal server, enabling the user to copy information from the terminal server to his local client directories.

    - **User can connect printers**—Enables the user to create bookmarks that connect his local printers to the terminal server, enabling the user to print information from the terminal server to his local printer.

    - **User can connect COM ports**—Enables the user to create bookmarks that connects his COM ports to the terminal server, allowing communication between the terminal server and the devices on his serial ports.

    - **Allow Clipboard Sharing**—Enables the user to create bookmarks that shares the contents of the clipboard between the user's host computer and the terminal server. Due to the limitations in the pre-6.0 versions of the RDP client, disabling the Allow Clipboard Sharing option will automatically disable the Connect local drives, Connect local printers, and Connect COM Ports options.

        When you enable local resources through the terminal server, each user can only access his own local resources. For instance, user 1 cannot see user 2's local directories.

        The Connect Devices options that you specify at the role-level override any Connect Devices options that you set at the bookmark level.

    - **User can connect smart cards**—Allows users to use smart card readers connected to their system for authenticating their remote desktop session.

    - **User can connect sound devices**—Allows users to redirect audio from the remote desktop session to their local system.

        NOTE: Smart cards redirecting audio are supported by Microsoft Remote Desktop Protocol versions 5.1 and later.

7. In the Allow users to modify Display settings below section:

    - Select **Desktop background** to display your current wallpaper setting. If you do not select this option, your background is blank.

- Select **Show contents of window while dragging** to show the contents of the Windows Explorer window while moving the window around your desktop.

- Select **Menu and window animation** to animate the movement of windows, menus, and lists.

- Select **Themes** to allow Windows themes to be set in the terminal server window.

- Select **Bitmap Caching** to improve performance by minimizing the amount of display information that must be passed over a connection.

- Select **Font Smoothing (RDP 6.0 onwards)** to make text smoother and easier to read. This option only works on Windows Vista computers running RDP clients that are version 6.0 or later.

8. Click **Save Changes**.

Related Documentation

- Task Summary: Configuring the Terminal Services Feature

Configuring Terminal Services Resource Policies

When you enable the Terminal Services feature for a role, you need to create resource policies that specify which remote servers a user can access. You can create resource policies through the standard interface (as described in this section) or through resource profiles (recommended method).

The information in this section is provided for backwards compatibility. We recommend that you configure access to Windows terminal servers and Citrix servers through resource profiles instead, since they provide a simpler, more unified configuration method.

When writing a Terminal Services resource policy, you need to supply key information:

- Resources—A resource policy must specify one or more resources to which the policy applies. When writing a Terminal Services policy, you need to specify the terminal server to which users can connect.

- Roles—A resource policy must specify the roles to which it applies. When a user makes a request, the system determines what policies apply to the role and then evaluates those policies that correspond to the request.

- Actions—A Terminal Services resource policy either allows or denies access to a terminal server.

The system's engine that evaluates resource policies requires that the resources listed in a policy's Resources list follow a canonical format.

To write a Terminal Services resource policy:

1. In the admin console, choose **Users > Resource Policies > Terminal Services > Access**.

2. On the Terminal Services Policies page, click **New Policy**.

3. On the New Policy page, enter a name to label this policy and optionally description.

4. In the Resources section, specify the servers to which this policy applies.

5. In the Roles section, specify which roles to which this policy applies.

6. In the Action section, specify:

- **Allow access**—To grant access to the servers specified in the Resources list.

- **Deny access**—To deny access to the servers specified in the Resources list.

- **Use Detailed Rules**—To specify one or more detailed rules for this policy.

7. Click **Save Changes**.

8. On the Terminal Services Policies page, order the policies according to how you want the system to evaluate them. Keep in mind that once the system matches the resource requested by the user to a resource in a policy's (or a detailed rule's) Resource list, it performs the specified action and stops processing policies.

Related Documentation

- About Terminal Services Resource Profiles

- Specifying Resources for a Resource Policy

- Writing a Detailed Rule for Resource Policies

Specifying the Terminal Services Resource Option

Use the Options tab to match IP addresses to hostnames specified as resources in your terminal services resource policies. When you enable this option, the system looks up IP address corresponding to each hostname specified in a Terminal Services resource policy. When a user tries to access a server by specifying an IP address rather than the hostname, the system compares the IP to its cached list of IP addresses to determine if a hostname matches an IP. If there is a match, then the system accepts the match as a policy match and applies the action specified for the resource policy.

When you enable this option, the system compiles a list of hostnames specified in the Resources field of each Terminal Services resource policy. The system then applies the option to this comprehensive list of hostnames.

This option does not apply to hostnames that include wildcards and parameters.

To specify the Terminal Services resource option:

1. In the admin console, choose **Users > Resource Policies > Terminal Services > Options**.

2. Select **IP based matching for Hostname based policy resources**.

3. Click **Save Changes**.

Related Documentation

- Configuring Terminal Services Resource Policies

Using the Remote Desktop Launcher

End users can connect to a terminal server by:

- Entering rdp://hostname in the system browser bar

- Creating a terminal services bookmark

- Using the remote desktop launcher (RDPLauncher)

RDPLauncher uses the Terminal Services section in the end-user home page and allows the end user to enter a terminal service IP address or hostname. The default server port is 3389.

RDPLauncher provides only the screen. User experience options are not available through RDPLauncher. For example, the following options in the New Terminal Services Sessions window do not apply to terminal services launched through RDPLauncher:

- Client port

- Authentication settings

- Start application settings

- Connect Devices settings

- Display Settings

- Remote Audio

To allow end users to use RDPLauncher,

1. Select the **Terminal Services** option in Users > User Roles > *Role Name* > General > Overview.

2. Select **Enable Remote Desktop Launcher** in Users > User Roles > *Role Name* > Terminal Services > Options.

3. (optional) If your end users are on non-Windows systems, such as a Macintosh or Linux system, select **Enable Java for Remote Desktop Launcher** and select the applet to use.

> NOTE: If you select **Hob-Pulse Secure RDP Applet** from the Applet to Use menu, you must select the **Configure HTML for the default applet** check box in order to edit the HTML. Otherwise, the default HTML is used.

Screen size and color depth parameters for the RDPLauncher terminal services session are defined through Preferences > General on the end-users home page.

Related Documentation

- Creating Advanced Terminal Services Session Bookmarks

- Defining a Hosted Java Applet Autopolicy

CHAPTER 27 Pulse Collaboration

Pulse Collaboration Overview

Pulse Collaboration (formerly Secure Meeting) allows users to securely schedule and hold online meetings between both Connect Secure users and non-Connect Secure users. In meetings, users can share their desktops and applications with one another over a secure connection, allowing everyone in the meeting to instantaneously share electronic data on-screen. Meeting attendees can also securely collaborate online by remote-controlling one another's desktops and through text chatting using a separate application window that does not interfere with the presentation.

The number of meetings and users doubles in a cluster configuration compared to a single unit. For example, if you have x meeting/y users in a single unit, then you have $2x$ meeting/$2y$ users in a two-plus cluster unit.

> **NOTE:** During installation, if the Juniper Installer Service is not present Pulse Collaboration prompts for the administrator credentials. If you do not know the administrator credentials, Pulse Collaboration will install but the remote controlling of higher privilege processes feature will not be enabled. If you enter the administrator credentials correctly, this feature is enabled.

All Pulse Collaboration online meetings must be scheduled by a system user. The meeting creator specifies meeting details such as the meeting name, description, start time, start date, recurrence pattern, duration, password, and a list of invitees.

Meeting creators can use either of the following applications to schedule meetings:

- System end-user console—When the meeting creator uses the system end-user console to schedule a meeting, Pulse Collaboration displays it in the Meetings page of meeting-enabled system invitees. If you choose to enable a Simple Mail Transfer Protocol (SMTP) e-mail server, Pulse Collaboration also sends a notification e-mail to each invitee with a known e-mail address.

- Microsoft Outlook—When the meeting creator uses Microsoft Outlook to schedule a meeting, Outlook displays it in the Calendar page of other Outlook-enabled invitees and sends a notification e-mail to each invitee through the Outlook e-mail server. Pulse Collaboration also displays the meeting in the Meetings page of the end-user console for the meeting creator (but does not send e-mail notifications through the SMTP server).

Meeting creators can bypass these scheduling mechanisms if they choose to create instant meetings or support meetings instead of standard meetings.

You can create meetings with static URLs for a particular type of meeting (for example, weekly status meetings). You do not need to schedule these types of meetings. The host starts the meeting and the invitees enter the URL to attend the meeting.

Related Documentation

- Enabling and Configuring Pulse Collaboration

- Troubleshooting Pulse Collaboration

- Task Summary: Configuring Pulse Collaboration

Task Summary: Configuring Pulse Collaboration

To configure Pulse Collaboration:

1. Specify a network identity through the System > Network > Overview page of the admin console. Pulse Collaboration uses this hostname when constructing meeting URLs for e-mail notifications.

2. Configure role-level settings using settings in the following pages of the admin console:

   - Use settings in the Users > User Roles > *Role Name* > General page to enable Pulse Collaboration at the role level.

   - Use settings in the Users > User Roles > *Role Name* > Meetings > Options page to configure role-level meeting restrictions.

3. Specify which authentication servers meeting creators can access and search using settings in the following pages of the admin console:

   - Use settings in the Users > User Roles > Select Role > Meetings > Auth Servers page to specify which authentication servers meeting creators can access and search.

   - If you want to allow meeting creators to invite users from an LDAP server, use settings in Authentication > Auth. Servers > Select LDAP Server> Meetings page to enable the server.

4. If you want to change the default sign-in page or URL that meeting attendees use to sign into meetings, use settings in the following pages of the admin console to configure meeting sign-in policies:

- Use settings in the Authentication > Signing In > Sign-in Pages page to customize the pages that meeting attendees see when they sign into a meeting.

- Use settings in the Authentication > Signing In > Sign-in Policies > *Meeting Policy* page to define the URL that meeting invitees must use in order to access a meeting. You can also use this page to associate a meeting page with the URL.

- Use settings in the Authentication > Signing In > Sign-in Policies > *User Policy* page to associate your meeting sign-in policy with a user sign-in policy. The system applies the specified meeting URL to any meeting created by a user who signs into the associated user URL.

5. Configure system-level meeting settings, include session timeouts, SMTP server information, time zone settings, and color-depth settings using options in the System > Configuration > Pulse Collaboration page of the admin console.

6. If you want to enable client-side logging, use settings in the following pages of the admin console:

   - Use settings in the System > Log/Monitoring > Client Logs > Settings page of the admin console to enable client-side logging. You must enable this option in order to generate logs for end users and for meeting attendees.

   - Use settings in the System > Log/Monitoring > Uploaded Logs page of the admin console to view the logs.

     NOTE: Pulse Collaboration installs client files in different directories depending on your operating system and privileges. For more information, see *Pulse Connect Secure Client-Side Changes Installation Reference*.

Related Documentation

- Configuring Network Services

- Specifying Authentication Servers That Meeting Creators Can Access

- Configuring LDAP Search Attributes for Meeting Creators

- Configuring System-Level Meeting Settings

- Enabling Client-Side Logging

Scheduling Meetings Through the End-User Console

If you enable meeting creation abilities at the role level, system users can create meetings through the Meetings page of the end user console. When they do, they must specify all of the standard meeting details such as the meeting name, description, start time, start date, recurrence pattern, duration, password, and a list of invitees. Additionally, they must categorize all invitees into one of two categories:

- System invitees—A system invitee (also called an in-network invitee) is a user who signs into the same device or cluster as the meeting creator. When inviting a system user to a meeting, the meeting creator must specify the user's username and authentication server.

- Non-system invitees—A non-system invitee (also called an out-of-network invitee) is a non-system user or a system user who signs into a different device or cluster than the meeting creator. When inviting a non-system user to a meeting, the meeting creator must specify the user's e-mail address.

  NOTE: If a system invitee uses the meeting URL instead of the Meetings page in the end user console to join a meeting, Pulse Collaboration classifies the user as a non-system invitee.

Related Documentation

- Specifying Authentication Servers That Meeting Creators Can Access

- Joining Meetings

Scheduling Meetings Through Microsoft Outlook

If you enable meeting creation abilities at the role level, system users can create meetings through the Microsoft Outlook calendar using the Pulse Collaboration Outlook plug-in.

You must use the same Outlook profile to remove the Pulse Collaboration plug-in for Outlook as the one used to install the plug-in. Switching profiles between the installation and removal of the Plug-In is not supported.

NOTE: The Pulse Collaboration Outlook plug-in is only supported on Windows machines with Outlook 2007 or later.

To use this plug-in, the user must:

1. Install the plug-in from the Meetings page in the end user console.

2. Open the Pulse Collaboration scheduling form in Outlook by clicking the Collaboration button located on the Outlook ribbon. See Figure 110.

Figure110: Pulse Collaboration Button in the Outlook Ribbon
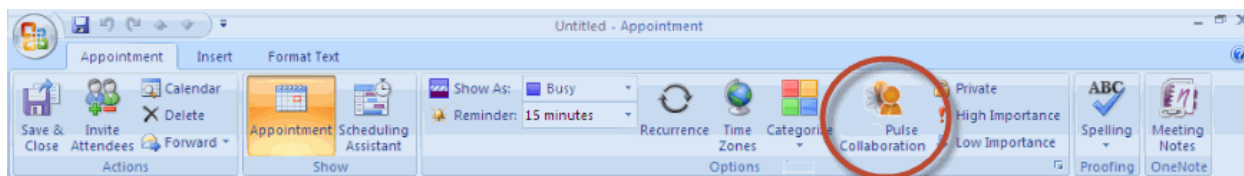
3. Use the Options window to enter details about the server on which the meeting should be scheduled as well as the user's sign-in credentials, realm, and a meeting password. See Figure 111. To open the Options window:

- (Outlook 2010) Select **File > Help > Options > Add-Ins** and then click **Add-in Options**.

- (Outlook 2007) Select **Tools > Options** and then click the **Pulse Collaboration** tab.

Figure111: Using the Pulse Collaboration Tab in the Options Window

Due to limitations with Microsoft Outlook, not all meeting details cross-populate between Microsoft Outlook and Connect Secure. For example, if the user schedules a meeting through Connect Secure, Microsoft Outlook does not display the meeting in its calendar. For a complete list of restrictions, see the Pulse Collaboration for Outlook document available from the end user help system as well as the Pulse Collaboration for Outlook plug-in installer.

4. Use the Outlook Scheduling and Appointment tabs to schedule the meeting and add invitees using standard Outlook functionality. Note that Pulse Collaboration supports creating standard or recurring meetings through Outlook.

NOTE: Depending on your Outlook version, the Appointment tab may have an Online Meeting button. This button is not related to the Meeting Server or the Pulse Collaboration Outlook Plug-in and cannot be used by a third-party plug-in. See Figure 112.

Figure112: Online Meeting Button Is Not Used
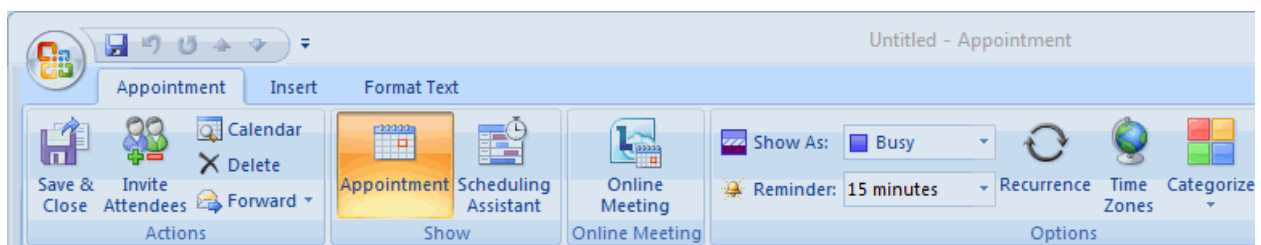
5. Save the calendar entry to send the information to the Pulse Collaboration server. Note that when saving a meeting, the user might see a certificate warning because the plug-in is communicating with a secure server.

6. Outlook sends invitation e-mails to the invitees using the text and meeting URL link constructed by the Pulse Collaboration Outlook plug-in. Outlook also adds the meeting to the Outlook calendars of meeting invitees. This calendar item includes all of the standard information recorded by Outlook as well as an additional Pulse Collaboration tab containing the information specified by the meeting creator in the Pulse Collaboration tab. Note that the system does not send an additional e-mail using the SMTP server.

7. To delete a meeting, right-click the meeting in the Outlook window and select Delete. See Figure 113.

Figure113: Deleting the Meeting
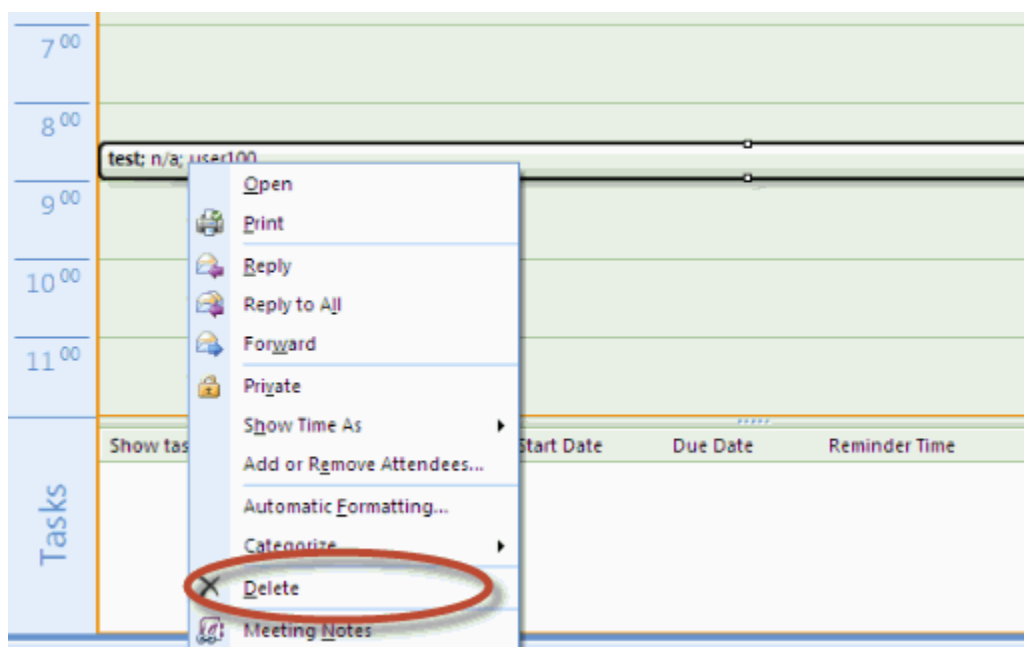
The Pulse Collaboration Outlook plug-in authentication does not work if the realm enables Host Checker policies or requires users to select a role.

Related Documentation

- Sending Notification E-mails

Sending Notification E-mails

You can configure Pulse Collaboration or Outlook to send notification e-mails to invitees when the meeting creator saves a new or modified meeting. The e-mail contains meeting details, a link that the invitee can use to join the meeting, and another link that the invitee can use to check whether his system is compatible with Pulse Collaboration.

ⓘ NOTE: You cannot send e-mail notifications for meetings.

If your users are scheduling meetings through the end user console, you must enable an SMTP server in the Users > Resource Policies > Meetings page of the admin console in order to send e-mail notifications to invitees.

If your users are scheduling a meeting through Microsoft Outlook, the Pulse Collaboration Outlook plug-in uses the e-mail addresses that are stored on the Outlook e-mail server.

If the person creating a Pulse Collaboration is using e-mail invitations and accesses the device using a URL that is not the fully-qualified domain name (for example, **https://sa** instead of **https://sa.company.com**), the e-mail invitation may display **https://sa** in the invitation information and not the true hostname. As a result, e-mail recipients may not be able to access the link from the e-mail. We recommend you configure the system's network identity. If configured, Pulse Collaboration invitations use that hostname.

Related Documentation

- Monitoring Pulse Collaboration
- About MyMeeting Meetings
- Specifying Authentication Servers That Meeting Creators Can Access

Pulse Collaboration Bridge Profile

- Pulse Collaboration Bridge Profile Overview
- Creating a Pulse Collaboration Bridge Profile
- Viewing, Editing and Deleting Your Pulse Collaboration Bridge Profile
- Entering the Pulse Collaboration Bridge Profile Conference and PIN IDs

Pulse Collaboration Bridge Profile Overview

You can embed the phone bridge number, or a list of phone numbers, and the conference access PIN into your e-mail so that mobile phone users can just click the link to join a Pulse Collaboration meeting. The Pulse Collaboration meeting bridge profile lets you enter "structured" data for audio bridges. This data can then be passed from the meeting server to a meeting client and used for dialing the audio bridge.

An end user opens the meeting invite e-mail through their mobile phone and selects a local phone number from the list. The mobile phone dials the selected phone number. After the call connects, the access code is automatically entered and the user connects to the Pulse Collaboration meeting.

Depending on how the bridge profile is configured, mobile phone users can use their location services feature or IP address to display the local phone number at the top of the list. For example, if the administrator creates a profile with "US", "China" and "India" numbers and includes the correct country codes, mobile phone users joining from the US will see the US number displayed at the top of the list.

> **NOTE:** The meeting bridge profile applies only to mobile phones. Joining audio conference bridges from a desktop client is not supported.

Creating a Pulse Collaboration Bridge Profile

The system ships with a sample Pulse Collaboration bridge profile template. You can download and edit the sample bridge profile template, shown below, to create your own custom version. For example, you may want to create a meeting bridge profile for each audio bridge provider your company uses.

```
<?xml version="1.0" encoding="UTF-8"?>
<!--
dialing-number: You can configure 1-500 dialing numbers in one profile.
    display-name: Please use UTF-8 encoding.
    country-code: Country code of the phone number.
    phone-number: Please enter digits only, without ' ', '(', '-' or ')'.
        You can enter 3-15 digits in each phone number.
```

moderator:
    Please enter the extra dialing steps for the moderator (host).
    In this example, after the moderator dials one of the phone numbers, she
    needs to enter a "Meeting Code" and a "PIN" to start the bridge.
    You only need to specify the names of the extra steps here, the moderator
    will configure their own pin codes.

participant:
    Please enter the extra dialing steps for other participants.
    In this example, after a participant dials one of the phone numbers, she
    needs to enter a "Meeting Code" to connect to the bridge.
    You only need to specify the names of the extra steps here, the moderator
    will configure their own pin codes.

```xml
-->
<jnpr:configuration xmlns:jnpr="http://xml.juniper.net/">
  <dialing-numbers>
   <dialing-number>
      <display-name>US Local</display-name>
      <country-code>1</country-code>
      <phone-number>8001112222</phone-number>
   </dialing-number>
   <dialing-number>
      <display-name>China Southern</display-name>
      <country-code>86</country-code>
      <phone-number>108001234567</phone-number>
   </dialing-number>
   <dialing-number>
      <display-name>Germany</display-name>
      <country-code>49</country-code>
      <phone-number>01801003825</phone-number>
   </dialing-number>
   <dialing-number>
      <display-name>India, Mumbai</display-name>
      <country-code>91</country-code>
      <phone-number>2261501727</phone-number>
   </dialing-number>
  </dialing-numbers>
  <moderator>
     <display-name>Meeting Code</display-name>
     <display-name>PIN Code</display-name>
  </moderator>
  <participant>
     <display-name>Meeting Code</display-name>
  </participant>
</jnpr:configuration>
```

The moderator and participant sections contain <display-name> parameters. These strings are
the text included in the collaboration e-mails and in the end-user Preference page that identify
the meeting conference code and the host's PIN number. For example, if you change the
participant <display-name> from "Meeting Code" to "Participant Passcode":

```xml
<participant>
    <display-name>Participant Passcode</display-name>
 </participant>
```

The collaboration e-mail invite will look similar to this:

```
Teleconference Info
--------------------------------------
Participant Passcode: 1234567
```

Local Access Numbers:
US Local: 12 123456789012345
China Southern: 86 108001234567
Germany: 49 01801003825
India, Mumbai: 91 0008001007678

Any bridge profile configured by the system administrator is available to all user roles.

To download and edit the sample bridge profile template:

1. Select **System > Configuration > Pulse Collaboration > Teleconference Bridge Profiles**.

2. Click **Download Template**.

3. Select **Save** to save the template to your hard drive, or click **Open** to open the template in your default editor.

4. Edit the template according to the instructions located at the top of the template and save it to your local hard drive. See the above example for the template format.

To upload your modified bridge profile:

1. Select **System > Configuration > Pulse Collaboration > Teleconference Bridge Profiles**.

2. Click **Upload New Profile**.

3. Enter a unique name to identify your new bridge profile template.

4. Click **Browse** to located your bridge profile and click **Open**.

5. Click **Create New Profile** and then click **Close Window**.

   The system parses the profile for syntax errors. If no errors are found, the new bridge profile appears in the Available Bridge Profiles section.

Viewing, Editing and Deleting Your Pulse Collaboration Bridge Profile

Once your Pulse Collaboration bridge profile is created, you can perform the following operations:

- Viewing the Contents of the Pulse Collaboration Bridge Profile

- Editing a Pulse Collaboration Bridge Profile

- Deleting a Pulse Collaboration Bridge Profile

Viewing the Contents of the Pulse Collaboration Bridge Profile

You can view a bridge profile's content at any time. The content includes the:

- Name of each country listed in the profile

- Country code

- Phone number

- Moderator information

- Participant information

To view the contents of a Pulse Collaboration bridge profile:

1. Select **System > Configuration > Pulse Collaboration > Teleconference Bridge Profiles**.

2. Click the **View** button next to the profile name that you want to display.

A new window appears displaying the contents of the bridge profile.

3. When you're done, click the **Close** icon located in the lower right corner of the window.

Editing a Pulse Collaboration Bridge Profile

You can edit a Pulse Collaboration bridge profile at any time. However, you must download the bridge profile to your local system before you can edit it. Updated information will be present the next time a meeting invite e-mail is sent.

To download and edit a Pulse Collaboration bridge profile:

1. Select **System > Configuration > Pulse Collaboration > Teleconference Bridge Profiles**.

2. Click the **Edit/Modify** button next to the profile name you just updated.

   1. Click **Download Profile**.

   2. Click **Open** to open the profile in your default text editor or click **Save** to save the profile to your local drive.

   3. Edit your bridge profile.

4. Click **Next**.

5. Click **Browse** to locate the bridge profile you just edited and then click **Open**.

6. Click **Update Profile**.

   The system parses the profile for syntax errors. If no errors are found, an upload successful message appears and the bridge profile is saved.

7. Click **Close Window**.

Deleting a Pulse Collaboration Bridge Profile

When you no longer need a Pulse Collaboration bridge profile you can remove it. Once you delete a bridge profile, you cannot retrieve it later.

To delete a Pulse Collaboration bridge profile:

1. Select **System > Configuration > Pulse Collaboration > Teleconference Bridge Profiles**.

2. Click the **Delete** button next to the name of the bridge profile you want to remove.

3. Click **Delete Profile** to delete the bridge profile or click **Cancel** to close the dialog box without deleting the bridge profile.

Entering the Pulse Collaboration Bridge Profile Conference and PIN IDs

End users can select the default bridge profile and assign their conference ID and host PIN ID through the Preference window of their home page.

When creating a meeting through the web interface, the meeting host can choose whether to include the bridge profile details (phone number, conference code, and PIN number) in the collaboration e-mail.

If you are using the Pulse Collaboration Outlook plug-in, you are given the option to include the bridge profile information in your e-mail.

i NOTE: If the option to include the bridge profile information is disabled, verify that you selected a default bridge profile.

End users must log in to their home page. To set up the conference ID and PIN numbers:

1. Click **Preferences** and then click **Pulse Collaboration**.

2. Under Meeting Bridge Profile, do the following:

   1. Select the bridge profile to use from the Default Meeting Bridge Profile pull-down menu.

   2. Enter the moderator meeting code and PIN number.

   3. Enter the participant PIN number.

      Typically, there is a slight delay after entering a PIN number and the meeting operator confirming your PIN number. To reduce this delay, append a pound sign (#) to the end of the PIN. For example, **1234#**.

4. Click **Save Changes**.

Joining Meetings

Invitees are allowed to join up to 15 minutes before the meeting is scheduled to start. Pulse Collaboration holds its online meetings on the system, allowing both system users and non-system users to attend meetings. (However, non-system meeting attendees cannot access anything on the system except the meeting to which they were invited.)

To join a meeting, Pulse Collaboration invitees must navigate to the meeting site on the Pulse Collaboration server using one of the following methods:

- Use the link provided in the Meetings page (system invitees only).

- Use the link provided in the notification e-mail.

- Enter the meeting URL in a Web browser.

  ℹ️　　NOTE: MyMeeting support only entering the meeting URL in a Web browser.

To obtain the URL for a meeting, the meeting creator can look on the Join Meeting page. Or, if you choose to use the default meeting URL, any meeting invitee can determine the appropriate URL by entering the applicable values into the following URL:

`https://<YourSA>/meeting/<MeetingID>`

Where:

- *<YourSA>* is the name and domain of the device hosting the meeting, such as *YourServer.yourcompany.com*. Pulse Collaboration pulls this name from the Hostname field in the System > Network > Overview tab, if defined. Otherwise, Pulse Collaboration pulls the system name from the meeting creator's browser.

- meeting is a literal string. (This string is always the same.) Note that meeting must start with a lower-case "m."

- *<MeetingID>* is the unique 8-digit identification number that Pulse Collaboration generates for the meeting. If the user does not include the meeting ID in the URL, Pulse Collaboration prompts him for it when he signs into the meeting. For example:

`https://connect.acmegizmo.com/meeting/86329712`

  ℹ️　　NOTE: You can choose to customize the meeting URL using the customized sign-in pages feature. If you do, users cannot access a meeting using the URL described here.

Once they have navigated to the meeting site, authenticated system users can directly join the meeting—they do not need to enter a username or password to access the meeting site on the system since they are already authenticated.

Non-system users must enter a name and password in the meeting sign-in page, however, since they are not yet authenticated. Pulse Collaboration authenticates the non-system users based on the meeting IDs and passwords that they enter in the sign-in page. (Note that the system does not use the invitees' names for authentication—it only uses the names for display purposes during the meeting.)

When an invitee chooses to join a meeting, Pulse Collaboration downloads and launches either a Windows client or a Java applet on to the invitee's system. This client-side component contains a meeting viewer, presentation tools, and a text messaging application. Once Pulse Collaboration launches the Windows client or Java applet on the user's desktop, the user becomes a meeting attendee and can begin participating in the meeting.

> **NOTE:** When configuring Pulse Collaboration, note that:
>
> - Pulse Collaboration does not work with PAC files on Macintosh or Linux systems.
> - Pulse Collaboration allows Windows users to join meetings through an NTLM proxy with or without authentication, provided that their browsers properly support proxies. Pulse Collaboration does not support NTLM proxies on Macintosh or Linux clients.

-

A few things to note about joining a meeting from an Apple iOS device using the Pulse Client:

- If the Pulse client is not installed on the Apple iOS device, nothing happens when a user clicks the mobile meeting invitation link.
- The Pulse client on an Apple iOS device supports only meetings from a system software 7.2 or later meeting host. If you try to access a meeting from a meeting host that has pre-7.2 software, the client will receive an "Invalid username, password or URL" error.
- If a user joins a meeting from an Apple iOS device, the user is removed from the attendee list if the iOS devices goes to sleep or is locked. The user is added back to the attendee list after the device is awakened or unlocked.

Sign-in URL Cookie

The last meeting sign-in URL ID is saved in a cookie. If you access another sign-in URL, this cookie is overwritten. If you access the collaboration server without a sign-in URL information, the web server uses the saved cookie and redirects you to the last sign-in URL you used.

For example, suppose User A schedules a meeting and makes User B the meeting host. User B receives and e-mail invite with the meeting host URL similar to this: https://vpn.mycompany.com/dana/meeting/weekly.cgi?t=65837483. When User B clicks this URL to start the meeting, User B is directed to the last sign-in URL they accessed from that system. If they previously accessed https://vpn.mycompany.com/vpn, clicking the URL in the e-mail invite redirects User B to https://vpn.mycompany.com/vpn. If User B did not previously access any sign-in URL, clicking the URL in the e-mail invite redirects User B to https://vpn.mycompany.com as expected.

A workaround is for User B to log in to the collaboration server and to start the meeting manually. This guarantees that the correct sign-in URL is used.

Related Documentation

- About MyMeeting Meetings
- About Sign-In Policies

Attending Meetings

By default, as soon as an attendee joins a meeting, he can see the names of other users who are attending the meeting and can start sending text messages to them using the Pulse Collaboration Chat window. However, you can choose to disable these capabilities in order to make meetings more secure or productive.

For instance, if your company's CFO chooses to hold a meeting with your company's analyst community, you can choose to hide attendee names in order to keep the identities of the analysts confidential. Additionally, you can choose to disable text chatting so that the meeting attendees cannot disrupt the CFO's presentation.

You can disable text chatting and enable hidden names for individual user roles. Or, you can specify that meeting creators within the role can decide themselves whether or not Pulse Collaboration hides attendee names. If you do, meeting creators can make this choice in the following situations:

- When scheduling or modifying a meeting from the Meetings page of the standard system interface. (The meeting creator cannot choose to hide attendee names from the Microsoft Outlook scheduling interface.)

- When joining a standard meeting or instant meeting. (Note, however, that the meeting creator can only choose to hide attendee names if he is the first person to join the meeting. If another attendee joins the meeting before the creator, Pulse Collaboration automatically displays the names of the meeting attendees and does not allow the meeting creator to change the display setting.)

If you or the meeting creator chooses to hide attendee names, Pulse Collaboration users can only see their own names and the names of the meeting host and presenter.

The Pulse Collaboration Chat functionality only supports users using the same language encoding (based on the Web browser settings) in a single meeting. Using a different encoding results in garbled text. Meeting invitations are sent based on the language setting in the creator's Web browser when meetings are created or saved.

Related Documentation

- Conducting Meetings

- Presenting Meetings

Conducting Meetings

The meeting host is a system user who is responsible for starting the meeting. Pulse Collaboration grants the host the following responsibilities and capabilities in order to help him effectively run his meeting:

- Starting the meeting presentation—Before the host joins, the other attendees can only chat. They cannot view or make a presentation because the host is also the default meeting presenter. The meeting presenter starts the meeting presentation by sharing his desktop or applications with other attendees.

- Passing host and presenter rights—The meeting host can choose to pass some or all of his responsibilities to another meeting attendee. For instance, after joining the meeting, the host can specify that another attendee should start the meeting presentation by passing that attendee presenter rights. The host can pass his host rights to any other system user and pass his presenter rights to any other system user or non-system user.

- Monitoring the meeting— The meeting host is responsible for expelling meeting attendees if necessary. The meeting host can also see the names of all meeting attendees so that he can determine who is attending.

- Ending the meeting—The meeting host is responsible for extending the meeting if it runs over the scheduled duration and closing the meeting when it is done.

Related Documentation

- Presenting Meetings

- Attending Meetings

Presenting Meetings

Once the presenter begins sharing, a meeting viewer automatically opens on all of the meeting attendees' desktops and displays the presenter's shared applications. Pulse Collaboration grants the presenter the following capabilities in order to help him effectively present to other users:

- Sharing multiple applications—The presenter can share a single application, multiple applications, or his entire desktop with other meeting attendees. (Note that Macintosh users cannot share individual applications. The can only share their desktops.)

- Passing controller rights—The meeting presenter can designate a controller. A *meeting controller* uses his own mouse and keyboard to remote control the presenter's shared desktop or applications. The presenter can pass remote control rights to any other attendee. When the presenter wants to regain control of his remote-controlled applications, he simply needs to click and Pulse Collaboration returns control to the presenter.

Like the meeting host, the meeting presenter can also see the names of all meeting attendees. Pulse Collaboration allows him to view all attendee names so that he knows to whom he is passing controller rights.

> NOTE: Pulse Collaboration cannot display the content of meeting presenter's desktop if it is locked.

Viewers on Linux and Macintosh clients may take a while to load the presentation if the presenter's desktop screen area is larger than 1856 x 1392.

Related Documentation

- Conducting Meetings

- Attending Meetings

About Instant Meetings and Support Meetings

Instant meetings and support meetings are meeting that users can quickly create without going through Connect Secure or Microsoft Outlook scheduling pages. Instead, a Connect Secure user simply needs to click the Instant Meeting button or Support Meeting button in the end-user console and click Start Meeting. The system then starts the meeting.

When creating instant meetings and support meetings, the system expedites the process by skipping certain scheduling steps. For instance, the system does not prompt the meeting creator to add the e-mail addresses of other invitees. Instead, it makes the meeting creator the only meeting invitee. The meeting creator can then provide other invitees with the information they need to join the meeting, such as the meeting URL, ID, and password.

The system also expedites the scheduling process by making certain assumptions about what the meeting attendees want to do. For instance, in addition to making the meeting creator the

only meeting invitee, the system also assumes that he wants to run the meeting and therefore makes him the meeting host. (In fact, since other attendees are probably joining the meeting through the meeting URL instead of the end-user console, the meeting creator is the only user who can conduct the meeting. Additionally, the system automatically assigns a meeting name ("Pulse Collaboration (MeetingID)" for instant meetings and "Support Meeting (MeetingID)" for support meetings), a meeting start time and date (immediately), a meeting duration (one hour), and a meeting recurrence (one-time meeting).

The system also uses default settings that correspond to the meeting type:

- Instant meeting—An instant meeting is basically a standard meeting that users can create more quickly. Therefore, when a user chooses to create an instant meeting, the system applies all of the user's role-level settings, such as authentication requirements, remote control, and secure chatting.

- Support meeting—A support meeting is a two-person meeting that is primarily intended to allow a system user to quickly troubleshoot another user's problem. Therefore, the system does not enable all of the user's role-level settings. Instead, the system automatically enables those options that facilitate quick troubleshooting and disables other settings, as described below:

  - Desktop sharing enabled—When the second user joins the meeting, the system automatically shares his desktop with the meeting host, enabling the host to immediately view the user's problem without having to explain what a meeting presenter is or how to share a desktop.

  - Remote control initiated—When the second user joins the meeting, the system automatically asks him whether the host can remote control his desktop. Assuming the user clicks Yes, the meeting creator can immediately start navigating through the user's computer in order to find and fix the problem. If the user clicks No, the host can gain remote control later using the standard request mechanisms.

  - Secure chatting disabled—The system does not expose the secure chatting feature during a support meeting, since users should not need to send text messages to each other. Instead, the users should talk to each directly over the phone.

Related Documentation

- Joining Meetings

- Conducting Meetings

About MyMeeting Meetings

MyMeetings, or personal meetings, are meetings that users can quickly create without going through the system or Microsoft Outlook scheduling pages. Instead, a user simply needs to click the Meeting button in the end-user console, enter the meeting subject and click Start Meeting. The system then starts the meeting.

MyMeeting meetings are different from instant meetings in that MyMeeting meetings have a fixed meeting URL for a specific meeting. You can bookmark this URL since it doesn't change. Meetings name must be unique within your personal meeting list and can be reused indefinitely until it is deleted by either the owner or the administrator. The meeting URL uses the format:

`https://YourDevice/MyMeetingRoot/userToken/MeetingID`

where:

- *YourDevice* is the name and domain of the device hosting the meeting, such as yourserver.yourcompany.com. Pulse Collaboration pulls this name from the Hostname field

in the System > Network > Overview tab, if defined. Otherwise, Pulse Collaboration pulls the device name from the meeting creator's browser.

- *MyMeetingRoot* is the root string of your personal URL. By default, the root is meeting.

- *userToken* is a string that uniquely identifies this URL. It can be the user's username, a string (with a number automatically appended for uniqueness), or an expression. For example:

```
https://my.company.com/meetings/chris/
https://my.company.com/meetings/room1
https://my.company.com/meetings/chris.andrew
```

- *MeetingID* is the meeting name token for this meeting. It is static for a particular meeting, and can be reused indefinitely until it is deleted. For example:

```
https://my.company.com/meetings/chris/weeklystaff
```

The user's Meetings page displays their personal meeting address(es). Users can send this URL to the invitees to join whenever the meeting starts.

All past meetings are listed on the user's Meetings page, making it easy to locate a specific meeting and retrieve the meeting details.

Joining MyMeeting Meetings

Attendees can join the MyMeeting meeting by entering the meeting URL in a browser.

Once they have navigated to the meeting site, authenticated system users can directly join the meeting—they do not need to enter a username or password to access the meeting site on the system since they are already authenticated.

Non-system users must enter a name and password in the meeting sign-in page, however, since they are not yet authenticated. MyMeeting authenticates the non-system users based on the meeting IDs and passwords that they enter in the sign-in page. (Note that the system does not use the invitees' names for authentication—it only uses the names for display purposes during the meeting.)

Related Documentation

- Configuring System-Level Meeting Settings

Enabling and Configuring Pulse Collaboration

To enable and configure meetings:

1. In the admin console, choose **Users > User Roles**.

2. Select a role.

3. If you have not already enabled Pulse Collaboration, in the General > Overview tab, select the **Meetings** check box and click **Save Changes**.

    ℹ️   NOTE: If you do not select the **Meetings** check box, users cannot create meetings, schedule meetings, or view the Meetings page. Note, however, that they can still attend the meetings to which they are invited by using the link provided in their invitation e-mails or by directly entering the meeting URL in to their web browsers.

4. Choose the **Meetings > Options** tab.

5. Under Meeting Types section, specify the type of meeting you want to provide users:

- **Users cannot create meetings**—Select this option to disable meeting creation and scheduling, but still provide users access to the Meetings page in order to join the meetings to which they are invited.

- **MyMeetings**—Select this option to allow users to create personal meetings without having to schedule them ahead of time.

    - **Users can create additional meeting URLs under their personal URL**—Select this check box if you want to enable users to create additional <meetingID> tokens.

    - **Users can create Support meetings**—Select this check box if you want to enable users to create two-person support meetings.

- **Standard Meetings**—Select this option to allow users to create scheduled meetings through the Meetings page.

    - **Users can create Scheduled meetings**—Select this check box to allow users to create scheduled meetings.

    - **Users can create Instant meetings**—Select this check box to allow users to create instant meetings.

    - **Users can create Support meetings**—Select this check box if you want to enable users to create two-person support meetings.

6. Under Authentication Requirements, specify the authentication restrictions that you want users to apply to the meetings that they create:

    - **Meeting password optional** (more accessible)—Select this option to allow the meeting creator to decide whether or not the meeting requires a password to join. When you choose this option, anyone who knows the meeting URL, ID number, and password (if applicable) can join the meeting, including non-system users.

    - **Require meeting password** (more secure)—Select this option to require the meeting creator to either create a meeting password or use the one generated by Pulse Collaboration. When you choose this option, anyone who knows the meeting URL, ID number, and password can join the meeting, including non-system users.

    - **Require server-generated password** (even more secure)—Select this option to require the meeting creator to use the password generated by Pulse Collaboration. When you choose this option, anyone who knows the meeting URL, ID number, and password can join the meeting, including non-system users.

    - **Require secure gateway authentication** (most secure)—Select this option to only allow invited users authenticated against the system to attend meetings. When you choose this option, the meeting creator does not need to create a meeting password, since all users must authentication through the system.

7. (MyMeeting only) Under Password Options, specify password requirements.

    - **Minimum length**—Set the minimum character length for passwords.

    - **Maximum length**—Set the maximum character length for passwords (optional). The maximum length cannot be less than the minimum length. There is no maximum limit to the length.

    - **Password must have one or more digits**—Select this option to require passwords to have at least one digit.

    - **Password must have one or more letters**—Select this option to require passwords to have at least one letter.

    - **Password must have mix of UPPERCASE and lowercase letters**—Select this option if you want all passwords to contain a mixture of upper- and lowercase letters.

- **Password must be different from username**—Select this option if the password cannot equal the username.

8. (MyMeeting only) Under Password Management, specify when passwords should be changed.

   - **Allow meeting creator to decide**—Select this option to let the meeting creator decide when to change the password

   - **Every_meetings**—Select this option to specify the number of meetings after which a password expires.

9. (Standard Meetings only) Under Password Distribution, specify the distribution method that you want meeting creators to employ:

   - **Do not display the password in the notification email** (more secure)—Select this option to require that meeting creators manually distribute the meeting password to invitees. When you select this option, Pulse Collaboration does not distribute the password in the automatic e-mail notifications it sends to invitees and Microsoft Outlook does not display the Pulse Collaboration tab (which contains the meeting password) to invitees. Omitting the password from the meeting e-mail and Microsoft Outlook calendar entry helps increase meeting security.

   - **Display the password in the notification email** (more accessible)—Select this option to automatically distribute the meeting password in the e-mail notification sent by Pulse Collaboration and to display the Pulse Collaboration tab in Microsoft Outlook calendar entries.

   - **Allow the meeting creator to decide**—Select this option to allow the meeting creator to determine whether or not Pulse Collaboration and Microsoft Outlook should automatically distribute the meeting password to meeting invitees.

     NOTE: You must enable an e-mail server in order to send meeting notification e-mails.

10. (Instant or Scheduled meetings only) Under Attendee Names, specify whether you want Pulse Collaboration to display the names of attendees during a meeting:

    - **Do not allow hiding of attendee names**—Select this option to always display the names of meeting attendees.

    - **Allow meeting creator to hide attendee names**—Select this option to allow the meeting creator to decide whether or not to display the names of meeting attendees.

    - **Hide attendee names**—Select this option to always hide the names of meeting attendees. Note that when you select this option, Pulse Collaboration still exposes the names of the meeting host and presenter to other meeting attendees.

11. (Instant or Scheduled meetings only) Under Remote Control, specify whether you want to allow meeting presenters to share control of their desktops and applications with other meeting attendees:

    - **Allow remote control of shared windows** (more functional)—Select this option to allow the meeting presenter or host to pass control of the presenter's desktop and desktop applications to any of the meeting attendees, including non-system users.

    - **Disable remote control** (more secure)—Select this option to limit control of the meeting presenter's desktop and desktop applications exclusively to the presenter.

12. Under Secure Chat, indicate whether or not you want to allow users to chat during their meetings:

- **Allow secure chat** (more functional)—Select this option to enable chatting in the meetings that are created by users who map to this role.

- **Disable secure chat** (more secure)—Select this option to disable chatting in the meetings that are created by users who map to this role.

> ℹ️ NOTE: If you change this setting while a meeting is in progress (that is, after any user has joined the meeting), Pulse Collaboration does not apply the modified setting to the in-progress meeting.

13. (Standard Meetings only) Under Pulse Collaboration for Outlook, select the **Allow users to download Pulse Collaboration for Outlook Plugin** check box if you want to allow users to schedule secure meetings through Microsoft Outlook.

14. Under Meeting Policy Settings, indicate whether or not you want to restrict the resources that are used by Pulse Collaboration users:

- **Limit number of simultaneous meetings**—Select this check box and enter a corresponding value to specify the maximum number of meetings that may be held by at any given time by members of the role.

- **Limit number of simultaneous meeting attendees**—Select this check box and enter a corresponding value to specify the maximum number of people that may simultaneously attend meetings scheduled by members of the role.

- **Limit duration of meetings (minutes)**—Select this check box and enter a corresponding value to specify a maximum duration (in minutes) that a meeting may run.

> ℹ️ NOTE: The system also limits the number of meetings users can attend. An individual user can only attend one meeting at a time per computer and cannot attend more than 10 consecutive meetings within a 3 minute period.

15. Click **Save Changes**. The system adds a Meeting link to the secure gateway home pages of the users in the specified role.

Related Documentation

- About Instant Meetings and Support Meetings

- Sending Notification E-mails

- Attending Meetings

- Conducting Meetings

- Scheduling Meetings Through Microsoft Outlook

Permissive Merge Guidelines for Pulse Collaboration

If you choose to merge roles, the system merges all options on the Users > User Roles > Select Role > Meetings > Options page to favor more accessible settings rather than more secure, except policy settings. When applying the policy settings that control the number of meetings and attendees allowed per role, Pulse Collaboration runs through the various roles trying to find one whose limit is not yet reached.

For example, you might specify that the following roles can schedule the following number of meetings:

- Engineering: 25 meetings

- Management: 50 meetings

- Sales: 200 meetings

If Joe maps to all of these roles (in the order listed), and tries to schedule a meeting, Pulse Collaboration first checks whether the scheduled meeting limit for Engineering has been met. If it has, Pulse Collaboration then checks the Management meeting quota. If that limit has been met, Pulse Collaboration checks the limit for the Sales role. Only when the limit for all of these roles has been reached does display a message to Joe telling him that the scheduled meeting limit has been reached and he cannot create a meeting. You cannot limit the number of meetings or meeting users at the realm level.

Related Documentation

- User Roles Overview

Specifying Authentication Servers That Meeting Creators Can Access

You can specify which authentication servers meeting creators may access and search when inviting other system users to meetings. When specifying servers, you can select any authentication server that you have enabled through the Authentication > Auth. Servers page of the admin console.

When you enable servers for meeting creators, Pulse Collaboration displays the following tabs to them in the Add Invitees dialog box:

- **Local**—Using the Local tab, the meeting creator may access and search for users from any enabled authentication server (including LDAP servers). The meeting creator may access and search all users that are managed through a local system authentication server in addition to all users that are managed by other types of authentication servers and cached in the system's memory. The meeting creator cannot view or search for users who are included in a non-system server's database but have not yet signed in to the system and created persistent data (such user bookmarks or password modifications).

- **LDAP**—If you enable an LDAP server, Pulse Collaboration displays the LDAP tab in the Add Invitees dialog box. The meeting creator may use this tab to access and search for all users in the enabled LDAP server(s)—not just those users who are cached in the system's memory. When a meeting creator adds a user through the LDAP tab, Pulse Collaboration also uses the e-mail attribute defined in the LDAP server to populate the invitee's e-mail address in his notification e-mail.

When adding local and LDAP users, the meeting creator's ability to access and search the servers is dependent on options you specify in the Auth Servers tab of the admin console. This tab contains two options that you may use to control access to each authentication server:

- **Access**—Select this option to allow the meeting creator to add and validate users from the corresponding authentication server. If you enable this option, Pulse Collaboration validates any users that the meeting creator tries to add from this server. If the meeting creator enters the name of a user that does not exist, Pulse Collaboration displays a warning to the creator when he finishes configuring the meeting and removes the invalid user from the list of invitees. If you disable this option, the meeting creator must use e-mail addresses instead of system usernames to invite any users from this server to a meeting. Pulse Collaboration then treats the specified users as non-system invitees.

- **Search**—Select this option to allow the meeting creator to search user entries in the corresponding authentication server. If you enable this option, Pulse Collaboration displays information about all available users who match the search criteria entered by the meeting creator. If you disable this option, the meeting creator must know the exact username and authentication server of the system users that he wants to invite to the meeting.

NOTE: If you enable an LDAP server, note that it must be searchable. Also note that you may use options in the Authentication > Auth. Servers > Select LDAP

Server > Meetings tab to specify individual LDAP attributes that Pulse Collaboration should display to meeting creators when they search an LDAP database.

To specify which authentication servers users may access and search when scheduling a meeting:

1. In the admin console, choose **Users > User Roles**.

2. Select a role.

3. If you have not already enabled Pulse Collaboration, in the General > Overview tab, select the **Meetings** check box and click **Save Changes**.

4. Choose the **Meetings > Auth Servers** tab.

5. In the User's Authentication Server section, indicate whether the members of this role may access and search the authentication servers that they are currently authenticated against.

6. In the Authentication Servers section, indicate additional authentication servers that members of this role may access and search.

7. Click **Save Changes**.

Related Documentation

- Configuring System-Level Meeting Settings

Configuring System-Level Meeting Settings

Unlike other access features, Pulse Collaboration does not have a resource policy. Instead, you configure system-level settings that apply to all roles for which this feature is enabled. You can:

- Specify session lifetime limits for meetings.

- Enable daylight savings adjustments to scheduled meetings.

- Specify the maximum color depth of meeting presentations.

- Enable automatic e-mail notifications for users who are invited to meetings scheduled through the end user console.

- Define the MyMeeting URL.

To configure Pulse Collaboration:

1. In the admin console, choose **System > Configuration > Pulse Collaboration**.

2. In the Session lifetime section, specify values for:

   - **Idle Timeout**—Use this field to specify the number of minutes a meeting session may remain idle before ending.

   - **Max. Session Length**—Use this field to specify the number of minutes a meeting session may remain open before ending.

     i NOTE: The values entered here apply to the meeting session, not the system session. For example, you may enter lower session lifetime values in the Users > User Roles > Select Role > General > Session Options page of the admin console. If the user reaches one of the role-level values before joining a meeting, he must sign back in to the device in order to access the meeting through the end user console. If the user reaches these role-level values after joining a meeting, however, they are not applied to the meeting. He may continue attending the meeting uninterrupted until he reaches the resource policy-level limits specified here.

3. In the Upload logs section, select **Enable Upload Logs** to allow non-system users to upload meeting logs.

   ℹ️ NOTE: If you select the Upload Logs option, you must also use settings in the System > Log/Monitoring > Client Logs > Settings page of the admin console to enable client-side logging.

4. In the MyMeeting section, specify values for:

   - **Root meeting URL**—Select the meeting URL you want associated with MyMeeting meetings. Meeting URLs are created in the Authentication > Signing In > Sign-In Policies page.

   - **Meeting name**—Specify the token to append to the meeting URL to uniquely identify this URL. You can use:

     - **Username**—Append the user's username to the meeting URL.

     - **Sequential room number with prefix**—Specify a string to append to the meeting URL, such as a "meeting". Numbers will be appended to the string to ensure uniqueness. For example, meeting_room1, meeting_room2, etc.

     - **Expression**—Append an expression, such as <userAttr.lname>, to the meeting URL. If the attribute is not valid, username is appended to the meeting URL instead.

       ℹ️ NOTE: Changing this token affects only users who have not created meetings. Users who have already created MyMeetings retain their existing token setting.

       To view a list of MyMeeting URLs users have already created, see System > Status > Meeting Schedule. Choose MyMeeting URLs from the View drop-down menu.

5. In the Email meeting notifications section, select **Enabled** to enable an SMTP e-mail server. Then:

   - In the SMTP Server field, enter the IP address or hostname of an SMTP server that can route e-mail traffic from the appliance to the meeting invitees.

   - In the SMTP Login and SMTP Password fields, enter a valid login name and password for the specified SMTP e-mail server (if required by the SMTP server).

   - In the SMTP Email field, enter your e-mail address or the address of another administrator. Pulse Collaboration uses the specified address as the sender's e-mail if the e-mail creator does not configure his own e-mail address on the system.

     ℹ️ NOTE: If you enable an SMTP server for use with Pulse Collaboration, you should also define a virtual hostname for your system in the Hostname field of the System > Network > Overview tab. Pulse Collaboration uses the name you specify when populating notification e-mails with meeting URLs and when making SMTP calls. If your system maps to multiple names and you do not define a virtual hostname, you may need to restrict which name users sign in to before creating a meeting. For example, if your system maps to an internal name (such as sales.acmegizmo.com) that is only accessible from inside your company's firewall and another name (such as partners.acmegizmo.com) that is accessible from anywhere, users should sign in to partners.acmegizmo.com before creating meetings. Otherwise, non-system invitees will receive e-mail notifications containing links to the system to which they cannot connect.

- Select whether to enable iCal attachments. If you disable this option, meeting invites are sent as an HTML text message, not as invites that can be imported into your calendar tool. By default, iCal attachments are enabled.

6. In the Options section, configure daylight savings and color-depth options:

- From the Observe DST rules of this country list, specify the country whose daylight savings time rules the system should observe. The client uses this setting as a baseline and then adjusts meeting times for individual users as necessary based on browser settings and client-side DST preference settings.

    > ⓘ NOTE: When a user signs into the device, Pulse Collaboration determines his time zone by running an ActiveX component called "Timezone Grabber" on his machine.

- Select **Enable 32-bit (True Color) Presentations** to allow users to present in true color. By default, Pulse Collaboration presents applications to users using the same color-depth as the presenter's desktop (up to 32-bit color). If you do not select this option and a user presents an application in 32-bit color, however, Pulse Collaboration changes the image to 16-bit to improve performance.

7. Click **Save Changes**.

8. Configure Pulse Collaboration settings for individual roles.

Related Documentation

- Enabling and Configuring Pulse Collaboration

Configuring a Pulse Collaboration Meeting Server

If your environment includes Pulse clients, you can set up your device as a meeting server so that your end users can start meetings from Pulse instead of having to log in to their secure gateway web interface.

Pulse clients can list more than one Pulse Collaboration server.

To set up your device as a meeting server for Pulse clients:

1. From the admin console, select **Pulse > Connections**.

2. Click the connection set containing the server you want to define as a meeting server.

3. Under Connections, click the connection server.

4. Under Options, select:

- **Support Pulse Collaboration integration on this connection**—Define this server as a meeting server and make it available in your end-user's Pulse Secure.

- **Support Remote Access on this connection**—Allow users to log in to this server using the web interface. If you want this server to be only a meeting server (users cannot log in to this server through the web interface), uncheck this option. However, make sure the **Support Pulse Collaboration integration on this connection** is selected.

5. Click **Save Changes**.

End users can start meetings when the Pulse client is not connected to a server. With this scenario, the Pulse client attempts to log in to the meeting server using saved credentials, if available. If saved credentials do not exist, or the server login fails, the end user is prompted to enter their credentials.

If the Pulse client is connected to a Pulse server when the user starts a meeting, the client first checks to see if the meeting server is the same server that Pulse is currently connected to. If

they are the same server, then the meeting starts immediately. If they are different servers, the client then checks if the end user is in the same authentication group on both servers. If it is the same authentication group, the client can use the existing user session to single sign-on (SSO) into the meeting server. This SSO to a meeting server feature is configurable through the admin console.

To configure Pulse Collaboration SSO:

1. From the admin console, select **Users > User Realms >**.

2. Select an existing realm, or create a new Realm.

3. On the General page, select the **Session Migration and Sharing** check box. Additional options appear.

4. Choose one of the following options:

   - **Enable Session Migration**—Copy the user session from the currently connected Pulse server to the Pulse Collaboration meeting server and then terminate the user session on the Pulse server.

   - **Enable Session Sharing**—Create a new user session on the Pulse Collaboration meeting server and retain the user's current session on the Pulse server.

5. Click **Save Changes**.

For more information on session migration, see the *Pulse Secure Client Administration Guide* located on the Pulse Secure support site.

Pulse Collaboration Meeting Server Use Cases

This topic describes several scenarios of starting a meeting where the Pulse server and Pulse Collaboration meeting server may or may not be the same. For these use cases, assume the following server configurations:

Server Name	Server Type	Additional Configuration
Server 1	Pulse Collaboration meeting server and allows active VPN sessions	Server 1 is configured for session sharing with Server 2, and the user authentication realms are the same between Server 1 and Server 2.
Server 2	Pulse Collaboration meeting server only	Server 1 is configured for session sharing with Server 2, and the user authentication realms are the same between Server 1 and Server 2. Server 2 is configured for session sharing with Server 3 using IF-MAP, and the authentication realms for the user are same between Server 2 and Server 3.
Server 3	Allows only active VPN sessions (is not a Pulse Collaboration meeting server)	Server 2 is configured for session sharing with Server 3 using IF-MAP, and the authentication realms for the user are same between Server 2 and Server 3.
Server 4	Pulse Collaboration meeting server and allows active VPN sessions	No session sharing with any other server.

Use Case: User has an active VPN Session on Server 1 and starts a meeting on Server 1

Pulse uses the user's existing VPN session credentials to SSO the user in to the Pulse Collaboration meeting server. The Meeting web page appears (same as clicking the Meeting button on the end-user's secure gateway home page) and the user starts a meeting. The user is not logged out of the Pulse Collaboration meeting server when they end the meeting since this is a shared server. The meeting server sign out option on the Pulse tray menu is disabled.

If the user attempts to disconnect from Pulse while the meeting is active, a warning message alerts the user that disconnecting from the Pulse server will end their meeting.

Use Case: User has an active VPN Session on Server 1 and starts a meeting on Server 2

When the user starts the meeting on Server 2, Pulse uses the user's credentials on Server 1 to SSO in to Server 2. The Meeting web page appears (same as clicking the Meeting button on the end-user's secure gateway home page) and the user starts a meeting. The user can also sign out of Server 1 while their meeting is still active and the meeting will not disconnect. The user is not logged out of Server 2 when they end the meeting. The user can manually disconnect Server 1 or from Server 2 through the Pulse tray menu. Signing out on one server has no effect on the other server.

Use Case: User has an active VPN Session on Server 3 and starts a meeting on Server 1

Since Server 3 has session sharing only with Server 2, the option to start a meeting on Server 1 is disabled. In this case, the user can start a meeting only on Server 2.

Use Case: User has an active VPN Session on Server 3 and starts a meeting on Server 2

When the user starts the meeting on Server 2, Pulse uses the user's credentials on Server 3 to SSO in to Server 2. The Meeting web page appears (same as clicking the Meeting button on the end-user's secure gateway home page) and the user starts a meeting. The user can also sign out of Server 3 while their meeting is still active and the meeting will not disconnect. The user is not logged out of Server 2 when they end the meeting. The user can manually disconnect Server 3 or from Server 2 through the Pulse tray menu. Signing out on one server has no effect on the other server.

Use Case: User has no active VPN session and starts a meeting on Server 1

When the user starts the meeting on Server 1, they are prompted to log in. The Meeting web page appears (same as clicking the Meeting button on the end-user's secure gateway home page) and the user starts a meeting. The user is not logged out of the Pulse Collaboration meeting server when they end the meeting since this is a shared server. The meeting server sign out option on the Pulse tray menu is disabled. If the user attempts to disconnect from Pulse while the meeting is active, a warning message alerts the user that disconnecting from the Pulse server will end their meeting.

Use Case: User has no active VPN session and starts a meeting on Server 2

When the user starts the meeting on Server 2, they are prompted to log in. The Meeting web page appears (same as clicking the Meeting button on the end-user's secure gateway home page) and the user starts a meeting. The user have to re-enter their credentials if they try to log in to Sever 1, Server 3 or Server 4. The user is not logged out of the Server 2 when they end the meeting. They can, however, manually disconnect from Server 2 through the Pulse tray menu.

Troubleshooting Pulse Collaboration

If you or your end users encounter problems with Pulse Collaboration and the admin console pages described above do not help you solve the problem, we recommend that you following the guidelines below.

Troubleshooting methods include:

- Uninstall the Pulse Collaboration client from your system—If you are having a problem launching Pulse Collaboration, click the Joining a Meeting: Troubleshooting link on the Join Meeting page, and then click Uninstall. Click Return to Join Meeting and try to launch the

meeting again. The next time you try to join a meeting, Pulse Collaboration updates your client with the latest version. For information about where Pulse Collaboration installs files and which files it leaves behind after uninstallation, see the *Client-side Changes Guide* on the Pulse Secure Customer Support Center.

- Check your system's compatibility—You might encounter problems joining or presenting at a meeting if your system configuration is not compatible with Pulse Collaboration. To determine if your system is compatible, navigate to the meeting sign-in page at any time or accept the meeting invitation e-mail and click Check Meeting Compatibility. Pulse Collaboration determines your compatibility level to achieve full compatibility if required. Note, however, that the Pulse Collaboration compatibility checker does not check all factors that can affect your meeting experience.

  For a comprehensive list of about the operating systems and browsers that are supported, as well as system requirements such as CPU, memory, monitor resolutions, and screen depths, see the *Supported Platforms Document* posted on the Pulse Secure Global Support Center (PSGSC) Center.

- Determine if you are using unsupported functionality—Pulse Collaboration does not support the sharing of streaming media applications. Pulse Collaboration also does not support graphic intensive applications that dynamically change the screen resolution or screen depth.

- Install a production-level certificate on your system—We recommend that you install a production-level certificate on the Pulse Collaboration server when using Pulse Collaboration in conjunction with an SSL certificate. If you install a self-signed SSL certificate, Pulse Collaboration users might encounter difficulties signing in to meetings. If you choose to use a self-signed certificate, instruct meeting attendees to install the certificate before joining the meeting. (Through Internet Explorer, users should click View Certificate and then Install Certificate when they see the error message.)

- Refer to the *Pulse Collaboration Error Messages* PDF—The *Pulse Collaboration Error Messages* PDF on the Pulse Secure Global Support Center (PSGSC) Center lists errors that you might encounter when configuring or using Pulse Collaboration and explains how to handle them.

- Contact Pulse Secure Support—If you encounter an error and cannot solve it using the solutions described above, send a clear description of the problem to Pulse Secure Support with detailed steps explaining how to reproduce the problem, the error message text, your system operating system and build number, and your system administrator log files, installation log files, and client-side log files.

Known Issues with Pulse Collaboration

Launching Pulse Collaboration Using the Java Client

When using the Java client to launch a Pulse Collaboration, if the user clicks No on the certificate warning presented by the JVM, the meeting client does not launch, but it appears to the user as though the applet is still loading.

Toolbars on Macintosh and Linux Platforms

Even if the viewers are set to full screen, the toolbar is still visible on the Macintosh and Linux platforms.

Joining Meetings from a Cluster

When using two devices in a Pulse Collaboration cluster, users should always connect to the VIP (Virtual IP) address to join the Pulse Collaboration, not the IP address of the physical machine.

Clock Synchronization in Clusters

Pulse Collaboration may function erratically if the time clocks on devices in a cluster are not synchronized. We recommend you use the same NTP server for each node within a cluster to keep the system times synchronized.

Number Attendee Limitation with Safari

When creating a Pulse Collaboration using the Safari Web browser, you cannot add more than 250 attendees.

Dial-Up Bandwidth

When presenting, the presenter should consider which access methods are being used by attendees. Dial-up attendees may have bandwidth issues for presentations that redraw the screen or update the screen too frequently. If the presentation saturates the dial-up attendees' bandwidth, remote control and chat functions may not work, as they require sending data back to the system over the same, saturated, dial-up link over which they are receiving data.

Creating Clusters

In progress Pulse Collaboration meetings are stopped if a cluster is created during the meeting.

Related Documentation

- Using Device Certificates

- Monitoring Pulse Collaboration

Monitoring Pulse Collaboration

You can use the following pages in the admin console to monitor Pulse Collaboration performance and users:

- System > Status > Overview—Use this page to view system capacity utilization on the device.

- System > Status > Meeting Schedule—Use this page to view which users are currently signed in to a meeting and expel them from meetings if required.

Related Documentation

- Troubleshooting Pulse Collaboration

CHAPTER 28 E-Mail Client

About the E-Mail Client

The e-mail support provided by your system depends on the optional features licensed for your system:

- Secure Email Client option—If you have the Secure Email Client option, the system supports the Internet Mail Application Protocol (IMAP4), the Post Office Protocol (POP3), and the Simple Mail Transfer Protocol (SMTP).

  > NOTE: Secure Email Client does not support the roaming session options you can select in Users > User Roles > Role > General > Session Options.

- Secure Application Manager option—If you have the Secure Application Manager option, the system supports the native Microsoft Exchange MAPI protocol and the native Lotus Notes protocol.

  > NOTE: If your system is licensed with the Secure Application Manager option, which supports the native Microsoft Exchange MAPI protocol and the native Lotus Notes protocol, this section does not apply.

The Secure Email Client option enables users to use standards-based e-mail clients to access corporate e-mail securely from remote locations without the need for any additional software, such as a VPN client. The system works with any mail server that supports Internet Mail Application Protocol (IMAP4), Post Office Protocol (POP3), and Simple Mail Transfer Protocol (SMTP), including the Microsoft Exchange Server and Lotus Notes Mail server, which provide IMAP4/POP3/SMTP interfaces.

The system sits between the remote client and your mail server, serving as a secure e-mail proxy. The remote client uses the system as a (virtual) mail server and sends mail using the SSL protocol. The system terminates SSL connections from the client and forwards the decrypted mail traffic within your LAN to your mail server. The system then converts unencrypted traffic from the mail server into S-IMAP (Secure IMAP), S-POP (Secure POP), and S-SMTP (Secure SMTP) traffic, respectively, and transports it over SSL to the e-mail client.

> NOTE: If you have configured multiple sessions per user, note the following regarding e-mail sessions. If a user has concurrent sessions and starts the e-mail client from all sessions, the e-mail client from the last session is the only one that

can access the backend e-mail server through the system. For example, if a user has two concurrent sessions and starts the e-mail client from both sessions, only the second session can access the e-mail server.

Related Documentation

- Choosing an E-Mail Client

- Working with a Standards-Based Mail Server

- Working with the Microsoft Exchange Server

- About Lotus Notes and the Lotus Notes Mail Server

- Enabling the E-Mail Client at the Role Level

- Writing the E-Mail Client Resource Policy

Choosing an E-Mail Client

The system supports the following e-mail clients:

- Outlook 2000 and 2002

- Outlook Express 5.5 and 6.x

- Netscape Messenger 4.7x and Netscape Mail 6.2

Users who need remote access to e-mail typically fall into one of two categories:

- Corporate laptop users—These users use the same laptop when in the office and traveling.

- Home machine users—These users use a different machine at home than their office machine.

Before recommending an e-mail client to your users, read the following sections about how the supported clients interact with:

- Standards-based mail servers, including the Lotus Notes Mail Server

- Microsoft Exchange Server

  ⓘ NOTE: You can find instructions for configuring each of the supported e-mail clients on the Integration Guides and How-Tos page of the Pulse Secure Global Support Center (PSGSC) Center.

Related Documentation

- About the E-Mail Client

- Working with a Standards-Based Mail Server

- Working with the Microsoft Exchange Server

- About Lotus Notes and the Lotus Notes Mail Server

- Enabling the E-Mail Client at the Role Level

- Writing the E-Mail Client Resource Policy

Working with a Standards-Based Mail Server

The system works with mail servers that support IMAP4, POP3, and SMTP.

- IMAP Mail Servers

- Corporate laptop users—May use any of the six supported e-mail clients. We recommend that users use the same client—configured to point to the system—both while in the office and while traveling to ensure a seamless experience.

- Home machine users—May use any of the six supported e-mail clients for remote access to the IMAP server via the system.

- POP Mail Servers

  - Corporate laptop users—May use any of the four Outlook e-mail clients*. We recommend that users use the same client—configured to point to the system—both while in the office and while traveling to ensure a seamless experience.

  - Home machine users—May use any of the four Outlook e-mail clients* for remote access to the POP server via the system.

*The Netscape e-mail clients cannot be used in POP mode for remote access, because they do not support S-POP, which is required by the system for secure data transmission.

Related Documentation

- About the E-Mail Client

- Choosing an E-Mail Client

- Working with the Microsoft Exchange Server

- About Lotus Notes and the Lotus Notes Mail Server

- Enabling the E-Mail Client at the Role Level

- Writing the E-Mail Client Resource Policy

Working with the Microsoft Exchange Server

The Microsoft Exchange Server supports:

- Native MAPI (Messaging Application Programming Interface) clients

- IMAP clients

- POP clients

- Outlook Web Access (OWA)

The system provides access to the Microsoft Exchange Server through IMAP and POP clients using the Secure Email Client option and through OWA using the secure Web browsing feature.

Exchange Server and IMAP Clients

If your corporate mail server is an Exchange Server, then we presume that an employee's office machine is configured to use the Outlook 2000 or 2002 e-mail client in native MAPI mode.

- Corporate laptop users: May use either of the Outlook Express or Netscape clients for remote access to the Exchange Server via the system.

  NOTE: The Outlook 2000 client only supports one mail server configuration, which in this case would be the native MAPI mode, thus preventing users from using the same client for remote access. The Outlook 2002 client provides support for simultaneous MAPI and IMAP server configurations but does not support IMAP access when the MAPI account is off-line, preventing remote users from retrieving e-mail.

- Home machine users: May use any of the six supported e-mail clients for remote access to the Exchange Server via the system, assuming no MAPI account is configured on the remote machine.

When users run the Outlook Express or Netscape clients in IMAP mode, please note the following folder management behavior:

- When using Outlook Express mail clients—Deleted e-mails appear in the Outlook Express Inbox with a strike through them; they are not moved to the Deleted Items folder on the Exchange Server, which is the behavior when using the Outlook 2000 or 2002 client. When a user purges deleted e-mails in an Outlook Express client, the e-mails are gone forever. We recommend that Outlook Express users either:

  - Manually drag e-mails they wish to delete to the Deleted Items folder that appears under Local Folders (these are default folders that appear). This folder syncs with the Deleted Items folder on the Exchange Server, enabling users to retrieve deleted e-mails later.

  - Leave deleted e-mails in the Outlook Express Inbox, and then the next time they log in to their Outlook 2000 or 2002 program, move the deleted e-mails to the Deleted Items folder.

- When using Netscape mail clients—Deleted e-mails are moved to the Netscape Trash folder and no longer appear in the Netscape Inbox, but they do not disappear from the Outlook 2000 or 2002 Inbox unless users:

  - Configure the Netscape program to move deleted messages to the Trash folder and check the option to expunge the Inbox upon exiting.

  - Run only one program at a time and exit when finished so that the other program's Inbox synchronizes with the server and displays the same messages.

Also, sent e-mails are moved to the Netscape Sent folder (or other user-defined folder). If users wants sent messages to appear in the Microsoft Exchange Server Sent Items folder, then they need to manually drag them from the Netscape Sent folder to the Sent Items folder.

Exchange Server and POP Clients

If your corporate mail server is an Exchange Server, then we presume that an employee's office machine is configured to use the Outlook 2000 or 2002 e-mail client in native MAPI mode.

- Corporate laptop users: May use either of the supported Outlook Express clients for remote access to the Exchange Server via the system.

- Home machine users: May use any of the four Outlook clients for remote access to the Exchange Server via the system, assuming no MAPI account is configured on the remote machine.

  i NOTE: The Netscape e-mail clients cannot be used in POP mode for remote access, because they do not support S-POP, which is required by the system for secure data transmission.

Exchange Server and Outlook Web Access

To provide OWA access to your Exchange Server and enable users to access the Exchange Server through the system Web browsing feature, simply deploy OWA as a Web-based application on your intranet. You do not need to perform any additional setup to deploy an OWA implementation outside of your network.

ℹ️   NOTE: Using the system to access Outlook Web Access protects the Outlook Web Access IIS Web Server from standard attacks, such as Nimda, and thus is much more secure than putting Outlook Web Access directly on the Internet.

Related Documentation

- About the E-Mail Client

- Choosing an E-Mail Client

- Working with a Standards-Based Mail Server

- About Lotus Notes and the Lotus Notes Mail Server

- Enabling the E-Mail Client at the Role Level

- Writing the E-Mail Client Resource Policy

About Lotus Notes and the Lotus Notes Mail Server

The Lotus Notes Mail Server provides POP3 and IMAP4 interfaces, enabling users to retrieve e-mail from a Lotus Notes mail configuration through the system.

To enable access to:

- Corporate IMAP/POP/SMTP mail servers—Specify mail server, e-mail session, and authentication information in the Users > Resource Policies > Email Settings page of the admin console.

- Microsoft Exchange Servers and Lotus Notes Servers—Use settings in the Users > User Roles > SAM > Applications page of the admin console.

Related Documentation

- About the E-Mail Client

- Choosing an E-Mail Client

- Working with a Standards-Based Mail Server

- Working with the Microsoft Exchange Server

- Enabling the E-Mail Client at the Role Level

- Writing the E-Mail Client Resource Policy

Enabling the E-Mail Client at the Role Level

To use the Email Client feature, you must first enable it at the role level and then create a resource policy that specifies mail server settings.

To enable the Email Client feature at the role level:

1. In the admin console, choose **Users > User Roles >** *RoleName* **> General > Overview**.

2. In the Access features section, select the **Email Client** check box.

3. Click **Save Changes**.

4. Create a resource policy that specifies mail server settings.

Related Documentation

- About the E-Mail Client

- Choosing an E-Mail Client

- Working with a Standards-Based Mail Server

- Working with the Microsoft Exchange Server
- About Lotus Notes and the Lotus Notes Mail Server
- Writing the E-Mail Client Resource Policy

Writing the E-Mail Client Resource Policy

When you enable the Email Client access feature for a role, you need to create a resource policy that specifies mail server settings. Unlike other access features, Secure Email Client has only one resource policy that applies to all roles for which this feature is enabled. If you choose to enable the e-mail client service for users, you must specify IMAP/POP/SMTP mail server information and user authentication settings. The system serves as the e-mail proxy for the specified server(s).

The system supports multiple mail servers. You can require all users to use a default mail server or you can enable users to specify a custom SMTP and IMAP or POP mail server. If you allow users to specify a custom mail server, the user must specify the server settings through the system. The system manages e-mail usernames to avoid name conflicts.

> NOTE: SMTP e-mail proxy is not supported on the MAG Series Pulse Secure Gateways.

To write an Email Client mail server resource policy:

1. Enable the Email Client feature at the role-level.

2. In the admin console, choose **Users > Resource Policies > Email Client**.

3. Under Email Client Support, click **Enabled**.

4. Under Email Authentication Mode, select an option:

   - **Web-based email session**—Users must complete a one-time e-mail setup for the system. Then, users configure their e-mail client to use the username and password that are generated by the system e-mail setup. It is recommended that users sign into the device to start an e-mail session. (default)

   - **Combined Connect Secure and mail server authentication**—Users configure their e-mail client to use the following credentials:

     - Username—The user's normal mail server username or a username that is generated by the system e-mail setup if one of the following are true: (a) the user has multiple mail server usernames, or (b) the username on the system mail server are different

     - Password—The user's system password followed by a customizable credential separator character followed by the user's mail server password.

       Users do not have to sign in to the device to use e-mail.

   - **Mail server authentication only**—Users configure their e-mail client to use their normal mail server username and password. Users do not have to sign in to the device to configure or use e-mail.

     > NOTE: Your users can easily determine their username and password for e-mail by going to the Email Setup page.

5. Under Default Server Information, specify your mail server information. The system serves as the e-mail proxy for this server.

NOTE: You can specify only one default mail server. If users need to retrieve e-mail from more than one SMTP and POP or IMAP server, then allow users to define additional mail servers by clicking the appropriate check box. If you allow users to specify custom servers, they need to enter that server information one time in their Email Setup page.

6. Under Email Session Information, specify the:

- **Idle Timeout** value, which controls how long a user's e-mail session may remain idle before ending the e-mail client session.

- **Max. Session Length** value, which controls how long a user's e-mail session may remain active before ending the e-mail client session.

7. Click **Save Changes**.

Related Documentation

- About the E-Mail Client

- Choosing an E-Mail Client

- Working with a Standards-Based Mail Server

- Working with the Microsoft Exchange Server

- About Lotus Notes and the Lotus Notes Mail Server

- Enabling the E-Mail Client at the Role Level

CHAPTER 29 VPN Tunneling

About VPN Tunneling

The VPN tunneling access option (formerly called Network Connect) provides a VPN user experience, serving as an additional remote access mechanism to corporate resources using Connect Secure. This feature supports all Internet-access modes, including dial-up, broadband, and LAN scenarios, from the client machine and works through client-side proxies and firewalls that allow SSL traffic.

When a user launches VPN tunneling, the system transmits all traffic to and from the client over the secure VPN tunnel. The only exception is for traffic initiated by other system-enabled features, such as Web browsing, file browsing, and telnet/SSH. If you do not want to enable other system features for certain users, create a user role for which only the VPN tunneling option is enabled and make sure that users mapped to this role are not also mapped to other roles that enable other system features.

With VPN tunneling, the client's machine effectively becomes a node on the remote (corporate) LAN and becomes invisible on the user's local LAN; the system serves as the Domain Name Service (DNS) gateway for the client and knows nothing about the user's local LAN. Users may define static routes on their PCs, however, to continue to access the local LAN while simultaneously connecting to the remote LAN. Since PC traffic goes through the VPN tunnel to your internal corporate resources, make sure that other hosts within a user's local network cannot connect to the PC through the VPN tunnel.

In the event of broken network connectivity, only the Windows and Macintosh versions of VPN tunneling try (indefinitely) to reconnect.

You can ensure that other hosts in a remote user's LAN cannot reach internal corporate resources by denying the user access to the local subnet (configured on the Users > User Roles > Select Role > VPN Tunneling tab). If you do not allow access to a local subnet, then the system terminates VPN tunneling sessions initiated by clients on which static routes are defined. You may also require clients to run endpoint security solutions, such as a personal firewall, before launching a network-level remote access session. Host Checker, which performs endpoint security checks on hosts that connect to a device, can verify that clients use endpoint security software.

NOTE: A Hosts file entry is added by VPN tunneling to support the following case:

- If, when VPN Tunneling connects, split tunneling is disabled and the original externally resolved hostname (the hostname the user initially connected to prior to the VPN tunnel launch) resolves to another IP address against the internal DNS, the browser will redirect to a "Server not found" page, because no route is defined within the client system.

- At a graceful termination (sign-out or timeout) of the VPN tunnel client connection, the Hosts file is restored. If the Hosts file was not restored in a prior case due to an ungraceful termination, the Hosts file will be restored the next time the user launches VPN tunneling.

For VPN tunneling to communicate, the following ports must be open:

- UDP port 4242 on loopback address

- TCP port 443

- If using ESP mode, the UDP port configured on the device (default is UDP 4500).

The VPN tunneling option provides secure, SSL-based network-level remote access to all enterprise application resources using the device over port 443. Port 4242 is used for IPC communication between the VPN tunneling service and the VPN tunnel executable on the client PC. Typically endpoint products do not block this type of IPC communication. However, if you have an endpoint product that does block this communication, you must allow it for VPN tunneling to work properly.

> **NOTE:** If you enable the multiple sessions per user feature, VPN tunnel clients may not be assigned the same IP address. For example, VPN tunnel client may be assigned a different VIP address each time they connect to a device when the system is obtaining the DHCP addresses from a DHCP server.

Related Documentation

- Task Summary: Configuring VPN Tunneling

- Defining VPN Tunneling Role Settings

- About VPN Tunneling Resource Policies

- Network Connect Launcher (NC Launcher) Overview

- Troubleshooting Network Connect Errors

VPN Tunneling on 64-Bit Linux Platforms

A native 64-bit VPN Tunneling client is not yet available. Instead, changes in the existing 32-bit client were made so that it can be run on 64-bit platforms. Because of this, VPN Tunneling has dependencies with 32-bit Java and 32-bit standard libraries even when running on a 64-bit platform.

See the *Pulse Connect Secure Supported Platforms Guide* for a list of supported browsers, platforms and plug-ins for VPN Tunneling.

To run VPN Tunneling on a 64-bit Linux platform, you must perform the following tasks on your Linux system:

- Install a 64-bit web browser and configure the Java plug-in.

- Download and install the 32-bit Java for Linux.

- Update the Java alternatives links. If you install the 32-bit Java using package managers like "apt-get" and "yum" and so forth, the Java alternatives links are updated automatically and you can skip this step.

  - **sudo update-alternatives --install /usr/bin/java java *32-bit-Java-path priority***.

  - Check that 64-bit Java is the default. To see which is the default Java, use the **update-alternatives –-display java** command.

  - If necessary, use the **sudo update-alternatives –-config java** command to change the default Java.

- Install the standard 32-bit libraries and components. For example see Table 93.

Table93: Installing Standard 32-Bit Libraries and Components

Linux Distribution Method	Command
Ubuntu NOTE: This command installs Java 7 64-bit, the plugin for your browser, and the 32-bit Java with all related 32-bit libraries.	**apt-get install icedtea-7-plugin openjdk-7-jre:i386**
Fedora	**yum –y install xterm** **yum –y ld-linux.so.2** **yum –y libstdc++.so.6** **yum –y libz.so.1** **yum –y libXext.so.6** **yum –y libXrender.so.1** **yum –y libXtst.so.6**
OpenSUSE	**zypper install libXi.so.6**

The syntax for launching VPN Tunneling from the command line is:

32-bit_Java_path -cp NC.jar NC -h *ivehostname* **-u** *username* **-p** *password* **[-r** *realm***] -f** *sa_certificate_in_der_format* **[-l** *gui_log_level* **[-L** *ncsvc_log_level***] [-y** *proxy* **-z** *proxy_port* **[-s** *proxy_username* **-a** *proxy_password* **[-d** *proxy_domain***]]]**

There are no changes in the Connect Secure admin GUI or in the VPN Tunneling client user interface to run on a 64-bit Linux platform, however you should note the following:

- The 32-bit Java path must be used when launching VPN Tunneling from the command line.

- If the VPN Tunneling launcher cannot find the 32-bit Java path in the alternatives links, the "Setup Failed. Please install 32-bit Java and update alternatives links using update-alternatives command. For more details, please refer KB article KB25230." error appears.

- Agentless host checker is supported when VPN Tunneling is launched from a browser but is not supported when VPN Tunneling is launched from command line.

Related Documentation

- About VPN Tunneling

Task Summary: Configuring VPN Tunneling

The following steps do not account for preliminary configuration steps such as specifying the system's network identity or adding user IDs.

To configure Connect Secure for VPN tunneling:

1. Enable access to VPN tunneling at the role-level using settings in the Users > User Roles > Role > General > Overview page of the admin console.

2. Create VPN tunneling resource policies using the settings in the Users > Resource Policies > VPN Tunneling tabs:

1. Specify general access settings and detailed access rules for VPN tunneling in the Access Control tab of the admin console.

2. Specify Connection Profiles to assign to remote users in the Connection Profiles tab of the admin console.

3. (Optional) Specify split tunneling behavior for VPN tunneling in the Split Tunneling tab of the admin console.

4. Specify whether or not to enable GINA/Credential Provider installation, employ split tunneling, and/or auto-launch behavior in the Users > User Roles > *Role* > VPN Tunneling page of the admin console.

> **NOTE:** If you choose to activate split tunneling behavior in this page, you must first create at least one split-tunneling resource profile, as described above.
>
> You must enable VPN tunneling for a given role if you want a user mapped to that role to be able to use GINA/Credential Provider during Windows logon.

5. Specify an IP address for the VPN tunneling server-side process to use for all VPN tunneling user sessions on the System > Network > VPN Tunneling page in the admin console.

6. Ensure that an appropriate version of VPN tunneling is available to remote clients.

7. If you want to enable or disable client-side logging for VPN tunneling, configure the appropriate options in the System > Log/Monitoring > Client Logs > Settings page of the admin console.

To install VPN tunneling, users must have appropriate privileges, as described in the *Connect Secure Client-Side Changes Guide*. If the user does not have these privileges, use the Juniper Installer Service available from the Maintenance > System > Installers page of the admin console to bypass this requirement.

VPN tunneling requires signed ActiveX or signed Java applets to be enabled within the browser to download, install, and launch the client applications.

By default, Vista Advanced firewall blocks all inbound traffic and allows all outbound traffic. For VPN tunneling to work in conjunction with Vista Advanced firewall, configure the following settings:

- Change the Vista Advance firewall default settings to block all inbound and outbound traffic

- Create the following outbound rules in the appropriate firewall profile:

  - Create a port rule to allow any to any IP and TCP any port to 443

  - Create a custom rule to allow 127.0.0.1 to 127.0.0.1 TCP any to any

- Allow iExplorer.exe

In prior releases you could specify whether the system compiles packet logs for specific VPN tunneling users. This option is no longer available as it impacts performance.

Related Documentation

- Defining Default Options for User Roles

- Defining VPN Tunneling Access Control Policies

- Creating VPN Tunneling Connection Profiles

- Defining Split Tunneling Network Policies

- Downloading Client Installer Files

- Enabling Client-Side Logging

VPN Tunneling Execution

The VPN tunneling agent executes as follows:

1. If Graphical Identification and Authorization (GINA) is installed and registered on the remote client, the client automatically initiates a VPN tunnel to the device when the user signs into Windows; otherwise, the user needs to sign into the device and click on the VPN Tunneling link on the end-user home page (if you have not configured VPN tunneling to launch automatically).

    NOTE: SSO is supported only when VPN tunneling GINA is the only GINA installed on the client's system.

2. If the user does not have the latest version of the VPN tunneling installer, the system attempts to download an ActiveX control (Windows) or a Java applet (Macintosh and Linux) to the client machine that then downloads the VPN tunneling software and performs installation functions. If the system fails to download or upgrade the ActiveX control to a Windows client due to restricted access privileges or browser restrictions, the system uses a Java applet to deliver the VPN tunneling software to the client.

    NOTE: If Microsoft Vista is running on the user's system, the user must click the setup link that appears during the installation process to continue installing the setup client and VPN tunneling. On all other Microsoft operating systems, the setup client and VPN tunneling install automatically.

   Whether the system downloads an ActiveX control or a Java applet, both components attempt to identify the presence and version of existing VPN tunneling software on the client before determining which of the following installation functions to perform:

   - If the client machine has no VPN tunneling software, install the latest version.

   - If the client machine has an earlier version of VPN tunneling software, upgrade the shared VPN tunneling components to the newer version and install the most current UI version.

      NOTE: For information about valid Java applets, installation files and logs, and the operating system directories in which delivery mechanisms run, see the *Connect Secure Client-Side Changes Guide*.

3. Once installed, the VPN tunneling agent sends a request to the system to initialize the connection with an IP address from the pre-provisioned IP pool (as defined by the VPN Tunneling Connection Profiles resource policies applicable to the user's role).

4. The VPN tunneling system tray icon starts running in the taskbar on a Windows client or in the Dock on a Mac client.

5. The system allocates an IP address (from a VPN Tunneling Connection Profiles resource policy) and assigns a unique IP to the VPN tunneling service running on the client.

6. The client-side VPN tunneling service uses the assigned IP address to communicate with the VPN tunneling process running on the system.

7. After the system allocates an IP address to the client, it opens a direct channel of communication between the client and all enterprise resources to which the user's resource policy allows access. The internal application server sees the source IP as the client's IP address.

The client-side VPN tunneling agent communicates with the device, which, in turn, forwards client requests to enterprise resources.

> ℹ️ NOTE: If you use Host Checker to validate the presence of client-side security components based on policies you define on the system and the client cannot conform to the security policies at any point during a VPN tunneling session, Host Checker terminates the session.

Related Documentation

- Task Summary: Configuring VPN Tunneling

Automatically Signing into VPN Tunneling Using GINA

The Graphical Identification and Authorization (GINA) sign-in function is an automated sign-in method you can install and enable on Windows clients signing in to a Windows NT domain. You can require VPN tunneling to install GINA on the client machine, or you can allow users to decide whether or not to install GINA when they launch VPN tunneling.

> ℹ️ NOTE: You cannot install more than one GINA automatic sign-in function on a client's system. If another application on the client's system uses a GINA function, VPN tunneling cannot install and activate the GINA component.

If GINA is installed on the client, it automatically prompts the user to choose whether or not to launch VPN tunneling each time he/she signs in to Windows. If you choose to make GINA installation optional, the user can activate GINA using the Auto connect when login to Windows option in the VPN tunneling window. This option is only available during an open VPN tunneling session.

The option to enable GINA installation on client systems is available when you define role attributes in the Users > User Roles > *Role* > VPN Tunneling page. See Figure 114.

Figure114: GINA Installation Process
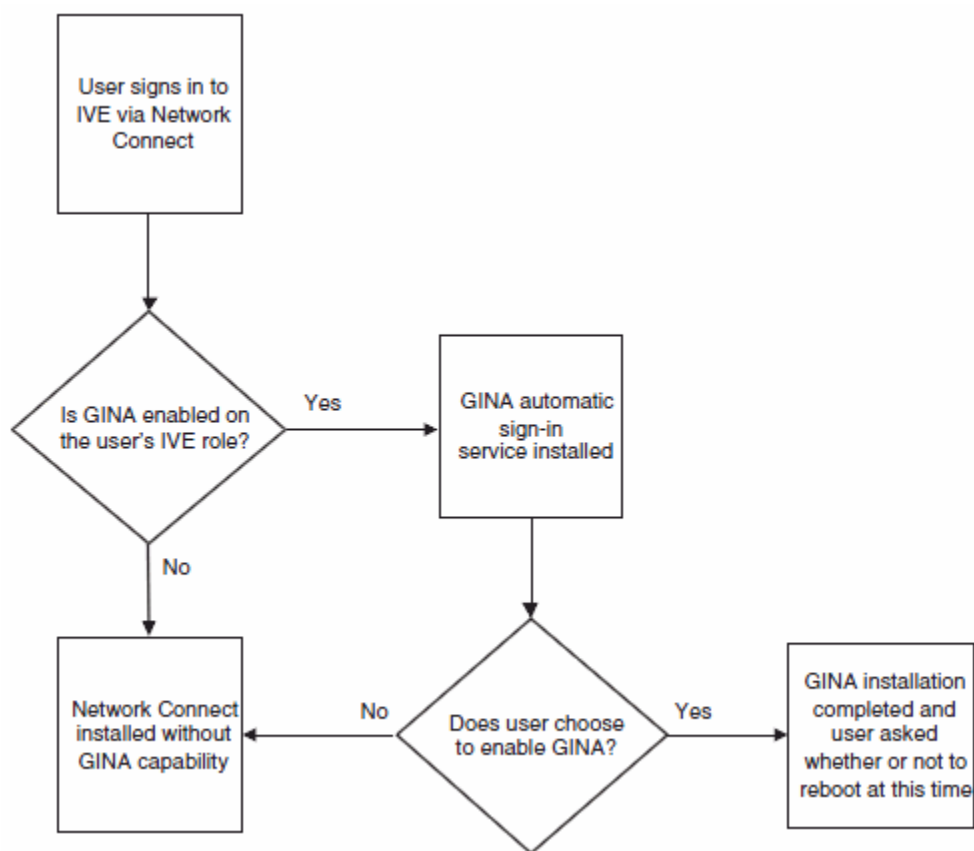
The GINA installation process takes place one time and requires the user to perform a system reboot in order to enable GINA sign-in capability. From that session forward, GINA prompts the user to decide whether or not to launch VPN tunneling at each Windows sign-in. When the user signs in, unless otherwise specified, GINA passes the user's Windows sign-in credentials to the system for authentication before establishing the VPN tunneling tunnel.

> NOTE: end users cannot modify their Windows user password through VPN tunneling GINA.

When a user logs in to the device through the Juniper GINA, if the version of the VPN tunneling client on the user's computer matches that on Connect Secure, the Pulse Secure GINAestablishes a VPN tunneling connection. If the VPN tunneling versions do not match, the Pulse Secure GINAdoes not establish a VPN tunneling connection to Connect Secure. Prior to release 5.4, the Pulse Secure GINAdisplays a version mismatch warning and allows users to log in to the Windows desktop using their cached credentials. With release 5.4 and later, the Pulse Secure GINAallows the users to log in to the Windows desktop using their cached credentials and then launches a standalone VPN tunneling. Users log in to the device and the appropriate VPN tunneling client automatically downloads to the user's computer and launches.

If you use Host Checker to validate the presence of client-side security components (pre-authorization), Host Checker starts after VPN tunneling is launched. This is sometimes called a

system-mode check. Host Checker exists after successful validation and is later restarted once the user is logs in to their desktop (called user-mode).

Using GINA Chaining

VPN tunneling supports GINA chaining. GINA chaining means that one GinaDLL calls another GinaDLL. By default, enabling VPN tunneling GINA also enables VPN tunneling GINA chaining. The VPN tunneling client detects any currently installed GINA component on top of the existing GINA chain. If the GINA component is compatible, VPN tunneling GINA is placed in front of the current GINA components. Currently, VPN tunneling supports the following GINA components:

- Cisco VPN client (CSGina.dll)

- Microsoft GINA (msgina.dll)

- Nortel Networks VPN client (nngina.dll)

- RSA SecurID (AceGina.dll)

- Novell GINA (NWGINA.dll)

If an installed GINA component is not supported (that is, not in the above list), a warning message appears and the VPN tunneling GINA is not installed.

If you uninstall a GINA component after VPN tunneling adds its information to the GINA chain, the VPN tunneling GINA removes the saved GINA information and does not call the removed GINA component the next time it goes through GINA chaining.

> **NOTE:** If the VPN tunneling GINA is installed at the top of the GINA chain (meaning, it is the last one installed), the VPN tunneling GINA is uninstalled when you uninstall the VPN tunneling client. However, if the VPN tunneling GINA is in the middle of the chain, you must remove all GINAs higher in the chain than the VPN tunneling GINA prior to removing the VPN tunneling GINA.

Credential Provider for Windows Vista and Later

In releases prior to Windows Vista, the customization of interactive user logon was done by creating a custom GINA. Users entered their authentication credentials in the logon UI and GINA passed this information to Winlogon for authentication. However, because GINAs do more than pass authentication information, they are typically difficult to implement.

Windows Vista introduced a new authentication model where the logon UI and Winlogon talk directly with each other. A credential provider is a module that plugs into the logon UI and describes the credential information required for the login UI to render and to communicate with an external authentication provider. After the credential provider gathers the credential information, it passes the final credentials to Winlogon.

There are two basic types of credential providers: standard authentication and Pre-Logon Access Providers (PLAP). Standard authentication includes password-based or certificate-based credentials. A PLAP is a special type of credential provider that allows users to make a network connection before logging in to their system. Another difference between these two

types of providers is timeout. PLAP credentials have no timeout where standard credentials typically have a 120 second timeout.

The VPN tunneling credential provider is a PLAP provider. This provider is visible only if the system is configured as part of a domain. The VPN tunneling provider creates a network connection. If the user's credentials are the same as the domain credential (SSO) then the credential information is entered only once. If the user's credentials are not the same as the domain credentials, the users selects another credential provider for domain authentication.

After a user logs in through VPN tunneling credential providers, the user has 5 minutes to log in to Vista either through single sign-on or through another credential provider. After the user logs into Vista, VPN tunneling attaches to the tunnel. If the user does not log in to Vista within 5 minutes, the VPN tunneling tunnel is disconnected.

To install the VPN tunneling credential provider,

1. Make sure your client user is part of a Windows domain.

2. In the Admin console, go to User Roles > VPN tunneling and select the Require VPN tunneling to start when logging into Windows option.

3. When installing VPN tunneling on the client system (running Windows Vista), you are prompted by the GINA/Credential Provider window to configure the GINA/Credential Provider authentication. Click OK.

4. Once the VPN tunnel is established on the client system, open the VPN tunneling window. Go to the Advanced View and select the Information tab. In the Results section, ensure that the GINA/Credential Provider plug-in is configured. You should see something similar to GINA Plug-In: Configured.

To use credential provider:

1. Log out of Windows and press Ctrl+Alt+Delete.

   You should see the Network logon icon. If you see only the Windows user standard tiles, click the Switch user option under the standard Windows credential tiles to see the Network logon icon.

2. Click the Network login icon and then click the Connect Secure logon icon.

3. Enter your Windows domain credential and click the right arrow button. For your username, use the format domain\username or user@domain.

   VPN tunneling signs the user in to the default URL and proxy server in config.ini.

   > NOTE: If your Connect Secure credential is not the same as your Windows domain credential, an alert box appears. Click OK and enter your Connect Secure credentials in the login window that appears. The window also contains an option button to launch another window to enter a URL, proxy server, and so forth.

VPN tunneling credential provider supports the following authentication provider: local authentication, LDAP, RADIUS (UN/PWD only), NIS, ADS and Dial-up connection. In additional, smart card credential provider supports certificate login.

Smart Card Credential Provider

Windows Vista also supports smart card credential provider—passing user credentials upon a smart card being inserted. If there is smart card present, a VPN tunneling Smart Card Credential Provider DLL tile shows on the PLAP layer. Click the tile and enter your smart card PIN to login.

To install the smart card VPN tunneling credential provider,

1. Make sure your client user is part of a Windows domain.

2. In the Admin console, select **User Roles >** *Role* **> VPN Tunneling** and select the **Require client to start when logging into Windows** option.

3. When installing VPN tunneling on the client system (running Windows Vista), you are prompted by the GINA window to configure the GINA authentication. Click **OK**.

   Use the smart card to log in to the device from a browser so the config.ini file will contain the smart card login URL which can then be used by the smart card DLL.

4. Once the VPN tunnel is established on the client system, open the VPN tunneling window. Go to the Advanced View and select the Information tab. In the Results section, ensure that the GINA plug-in is configured. You should see something similar to GINA Plug-In: Configured.

To use the smart card credential provider:

1. Log out of Windows and press Ctrl+Alt+Delete.

   You should see the Network logon icon located in the lower right corner of your screen. If you see only the Windows user standard tiles, click the Switch user option under the standard Windows credential tiles to see the Network logon icon.

2. Click the Network login icon and then click the smart card icon.

3. Enter your PIN number or password and click the right arrow button.

   VPN tunneling uses the PIN to retrieve the stored certificate and to log in to Connect Secure. After a successful login, the PIN is passed to Winlogon to log in to Vista.

   > NOTE: If your Connect Secure credential is not the same as your Windows domain credential, an alert box appears. Click OK. If a connection icon appears in the lower right corner of your screen, switch to the standard credential login tiles and log in to Vista. Otherwise, enter your Windows credential in the login box.

   > VPN tunneling retrieves the user principal name (UPN) from the smart card and compares them with the login user and domain names. If they do not match, the tunnel is disabled. The UPN typically has the format user@domain.

Related Documentation

- Task Summary: Configuring VPN Tunneling

Credential Provider Authentication for Connect Secure

The Pulse credential provider integration enables connectivity to a network that is required for the user to log on to the Windows domain. For example, the domain controller might reside behind a firewall and the endpoint uses credential provider login to connect to Connect Secure prior to domain login. Pulse integrates with Microsoft credential providers to enable password based login and smart card login. A credential provider interface appears as a tile on a Windows (Vista or later) login screen. See Figure 115.

Figure115: Pulse Logon Tile
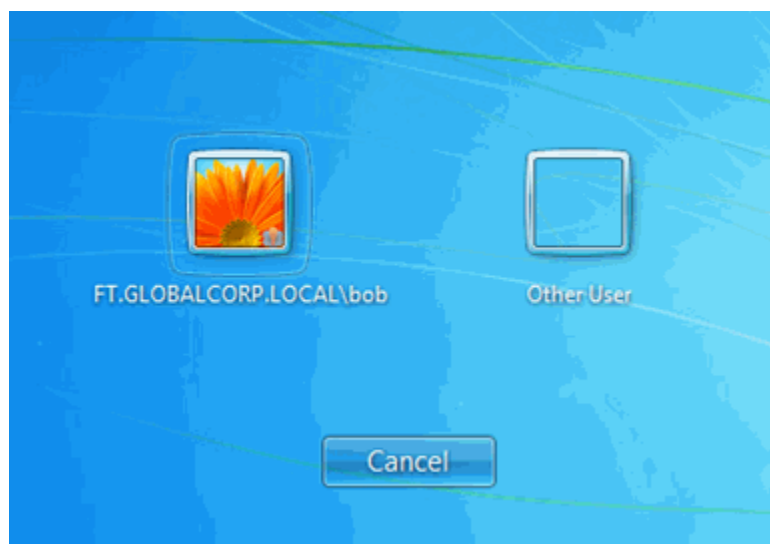
You enable Pulse credential provider support on a Pulse connection. After the connection has been downloaded to the endpoint through the normal Pulse distribution methods, a Pulse logon tile appears on the endpoint's desktop. When the user initiates the logon process, Pulse establishes the connection.

Pulse supports the following credential provider types:

- user-at-credprov—The connection is established before the user login using credentials collected at the selected credential tile, which provides single-sign-on functionality. The connection is maintained as an active connection on the user's desktop.

- machine-then-user-at-credprov—The connection is established using machine credentials when no user is logged in. When a user clicks a logon tile and provides user credentials, the machine connection is disconnected and a new connection is established. When the user logs off, the user connection is disconnected and the machine connection is reestablished. In one typical machine-then-user-at-cred prov implementation, the machine connection and the user connection are mapped to different VLANs.

Pulse credential provider support usage notes:

- If the endpoint includes more than one Pulse Layer 2 connection, Windows determines which connection to use:

1. If a network cable is attached to the endpoint, Layer 2 wired connections are attempted, and then wireless connections. If there are more than one wireless network available, the order is determined by the scan list specified as a Pulse connection option.

2. After all Layer 2 options are attempted, Pulse runs location awareness rules to find one or more eligible Layer 3 connections that are configured for credential provider login. If more than one Layer 3 connection is found, Pulse prompts the user to select a connection. A user can cancel the network connection attempt by clicking the cancel button.

3. After Pulse evaluates all configured connection options, Pulse returns control to Windows, which enables the user login operation.

- For connections that use user credentials, the Pulse connection may be configured so that prompts are presented during the login process, for example, prompts for realm or role selection or a server certificate trust prompt. For connections that use machine credentials, Pulse prompts cause the connection to fail because there is no interface to allow a response to the prompts. You can suppress any potential realm and role choice by specifying a preferred realm and role for the connection.

- Pulse upgrade notifications and actions are disabled during credential provider login and postponed until the user connection is established. Host Checker remediation notifications are displayed.

- To allow users to login using either a smart card *or* a password, you can create different authentication realms for each use case and then specify a preferred smart card logon realm and a preferred password logon realm as part of the connection properties.

To enable user-at-credprov credential provider support for a Pulse connection:

1. Create a Pulse connection set for the role (**Users > Pulse > Connections**), and then create a new Pulse connection. You can select **Connect Secure or Policy Secure (L3), Policy Secure (802.1X), or SRX** for the connection type.

2. In the Connection is established section, select one of the following options:

   - **Automatically at user login**—The user credentials are used to establish the authenticated Pulse connection to the network, log in to the endpoint, and log in to the domain server. The Pulse connection may be configured so that prompts are presented during the login process, for example, prompts for realm or role selection or a server certificate trust prompt.

   - **Automatically when the machine starts. Connection is authenticated again at user login**—Machine credentials are used to establish the authenticated Pulse connection to the network when the endpoint is started. When a user clicks the login tile and provides user credentials, the connection is authenticated again and the original connection is dropped. When the user logs off, the user connection is ended and the machine connection is established again. In one typical use case, the machine credentials provide access to one VLAN and the user credentials provide access to a different VLAN. Be sure that the Pulse connection does not result in Pulse prompts, for example, prompts for realm or role selection or a server certificate trust prompt, because the machine credential login does not present an interface to respond to the prompts.

3. For a Layer 2 connection that uses machine certificate authentication, make sure that the connection has an entry in the Trusted Server List. To allow any server certificate, type **Any**

as the Server certificate DN. To allow only one server certificate, specify the server certificate's full DN for example, C=US; ST=NH; L=Kingston; O=My Company; OU=Engineering; CN=c4k1.stnh.mycompany.net; E=ausername@mycompany.com.

4. Specify **Realm and Role Preferences** to suppress realm or role selection dialogs during the logon process:

- **Preferred User Realm**—Specify the realm that for this connection. The connection ignores any other realm available for the specific logon credentials

  The following options enable you to allow the user to login using a smart card or a password:

  - **Preferred Smartcard Logon Realm**—Preferred realm to be used when user logs in with a smart card.

  - **Preferred Password Logon Realm**—Preferred realm to be used when user logs in with a password.

    > NOTE: Be sure that the authentication realms you specify exist, and that they support the appropriate login credential option.

- **Preferred User Role Set**—Specify the preferred role or the name of rule for the role set to be used for user authentication. The role or rule name used must be a member of the preferred user realm.

To enable machine-then-user-at-credprov credential provider support for a Pulse connection:

1. Create a Pulse connection set for the role (**Users > Pulse > Connections**), and then create a new Pulse connection. You can select **Connect Secure or Policy Secure (L3), Policy Secure (802.1X), or SRX** for the connection type.

2. In the Connection is established section, select one of the following options:

- **Automatically at user login**—The user credentials are used to establish the authenticated Pulse connection to the network, login to the endpoint, and login to the domain server. The Pulse connection may be configured so that prompts are presented during the login process, for example, prompts for realm or role selection or a server certificate trust prompt.

- **Automatically when the machine starts. Connection is authenticated again at user login**—Machine credentials are used to establish the authenticated Pulse connection to the network when the endpoint is started. When a user clicks the login tile and provides user credentials, the connection is authenticated again and the original connection is dropped. When the user logs off, the user connection is ended and the machine connection is established again. In one typical use case, the machine credentials provide access to one VLAN and the user credentials provide access to a different VLAN. Be sure that the Pulse connection does not result in Pulse prompts, for example, prompts for realm or role selection or a server certificate trust prompt, because the machine credential login does not present an interface to respond to the prompts.

3. For a Layer 2 connection that uses machine certificate authentication, make sure that the connection has an entry in the Trusted Server List. To allow any server certificate, type **Any** as the Server certificate DN. To allow only one server certificate, specify the server

certificate's full DN for example, C=US; ST=NH; L=Kingston; O=My Company; OU=Engineering; CN=c4k1.stnh.mycompany.net; E=ausername@mycompany.com.

4. Specify **Realm and Role Preferences** to suppress realm or role selection dialogs during the logon process for both machine logon and user logon:

   - **Preferred Machine Realm**—Specify the realm that this connection uses when establishing the machine connection. The connection ignores any other realm available for the specific logon credentials

   - **Preferred Machine Role Set**—Specify the role or the name of rule for the role set that this connection uses when establishing the machine connection. The role or rule name used must be a member of the preferred machine realm.

   - **Preferred User Realm**—Specify the realm that for this connection that is used when a user logs onto the endpoint. The connection ignores any other realm available for the user's logon credentials.

   - **Preferred User Role Set**—Specify the preferred role or the name of rule for the role set to be used for user authentication. The role or rule name used must be a member of the preferred user realm.

5. Optionally specify pre-login preferences:

   - **Pre-login maximum delay**—The time period (seconds) that a Windows client waits for an 802.1x connection to succeed during the login attempt. The range 1 to 120 seconds.

   - **Pre-login user based virtual LAN**—If you are using VLANs for the machine login and the user login, you can enable this check box to allow the system to make the VLAN change.

6. Click **Save Changes** and then distribute the Pulse connection to Pulse client endpoints.

The Pulse tile appears on the login page the next time the end users log in.

> *i* NOTE: The user account must exist on both the Windows PC and on Connect Secure with the same login name.

Check the user logs for credential provider log-in information. See Figure 116.

Figure116: Credential Provider Log Information

Related Documentation

- *Client Connection Set Options for Pulse Policy Secure*

Launching VPN Tunneling During a Windows Secure Application Manager Session

Users can launch VPN tunneling while signed in to Connect Secure via Windows Secure Application Manager (WSAM). When a user launches VPN tunneling in this scenario, however, the VPN tunneling installer automatically terminates the WSAM session prior to launching VPN Tunneling.

During the process, the user is prompted with a warning message informing them that they are about to terminate their WSAM session in favor of launching VPN tunneling. We recommend that you configure users' VPN tunneling resource policies to feature as much access to network resources as they would have in their WSAM sessions. This way, when users choose to launch VPN tunneling (simultaneously terminating WSAM) they will still be able to access the same network resources.

> NOTE: If users choose not to launch VPN tunneling, the VPN tunneling installer still automatically installs the client application on their computer, but does not launch VPN tunneling. After the client application has been installed, users can choose uninstall it manually via their secure gateway home page or the folder options available in the Windows Start menu.

Related Documentation

- Task Summary: Configuring VPN Tunneling

Logging In To Windows Through a Secure Tunnel

Use the Logoff On Connect feature for users to log in to their Windows environment through an existing VPN tunnel. This feature lets them authenticate against a Windows Domain server in real time, as opposed to authenticating with the locally cached credentials. When this feature is enabled, they are automatically logged off Windows after the VPN tunneling session starts. The standard Windows login screen re-appears and they log in using their Windows credentials. Their Windows environment is now established through the VPN tunnel.

> NOTE: Users must log in to Windows within 5 minutes of the login screen re-appearing or before the Host Checker policy evaluate period ends, whichever is shorter. If they do not, their VPN tunnel connection may time out and they will not be logged in to Windows through a secure tunnel. An error appears if the VPN tunnel connection times out.
>
> The Logoff On Connect feature is not supported within SVW.

To use the Logoff On Connect feature:

1. Users log on to their local machine using their domain cached credentials. Their machine must be part of a Windows domain.

2. Users launch VPN tunneling and click **Tools** from the login page.

3. Select the **Logoff on Connect** option and click **OK**.

4. Users enter their username and password credentials in the login page.

A tunnel is established and logs them off of their local machine. The Windows login page appears.

5. Users enter their username and password credentials to sign in to their Windows Domain using the VPN tunnel.

Related Documentation

- Task Summary: Configuring VPN Tunneling

VPN Tunneling Connection Profiles with Support for Multiple DNS Settings

To ensure remote users are able to perform DNS searches as efficiently or as securely as possible, you can configure the system to allow multiple DNS settings during VPN tunneling sessions, based on a user's role membership.

When the system launches a user's VPN tunneling session, it uses a matching profile based on the user's role membership containing IP address, DNS, and WINS settings.

If you enable split-tunneling, the DNS search order setting allows you to define which DNS setting takes precedence—for example, search for a DNS server on the client's LAN before the system's DNS server, or vice-versa. VPN tunneling makes a backup of the client's DNS settings/search order preference before establishing a connection. After the session terminates, VPN tunneling restores the client to the original DNS settings. If you disable split-tunneling, all DNS requests go to the system's DNS server and your setting for the DNS search order preference does not apply.

> NOTE: After stopping and restarting a DNS client, the client may not pick up the search order of multiple DNS addresses in a timely manner, resulting in an incorrect lookup order when launching VPN tunneling. The rules governing DNS name resolution and failover are complex and often specific to the particular client operating system. You or the end user can attempt to run the ipconfig /registerdns commands from a command window on the client machine. This may reset the search order to the correct order. To understand the search resolution order for DNS servers, refer to the appropriate Microsoft DNS documentation for your operating system platform.

When employing a multi-site cluster of Connect Secure devices, the IP pool and DNS settings may be unique to each device residing at a different site. For this reason, the system allows the VPN Tunneling Connection Profile policy to be node-specific. That is, the resource policy enables the client to connect to the same device in the cluster each time a new session is established.

Related Documentation

- Deploying an IVS

- Configuring VPN Tunneling for Use on a Virtualized Connect Secure

VPN Tunneling Incompatibility with Other VPN Client Applications

Third-party vendor VPN client applications may be incompatible with VPN tunneling. Table Table 94 lists known VPN client vendors and VPN tunneling's relative compatibility with those vendors' VPN client applications.

Table94: VPN tunneling Compatibility with Third-Party VPN Clients

Vendor	Compatible?
Cisco	Yes
Nortel	Yes
NS Remote	Yes
Intel	Yes
Checkpoint	Yes

If you want to install VPN tunneling on a client featuring an incompatible VPN client application, you must uninstall the incompatible application before you install or launch VPN tunneling on the client.

Related Documentation

- Task Summary: Configuring VPN Tunneling

Linux Client Requirements

Linux clients signing in to VPN tunneling via Mozilla Firefox must ensure that the OpenSSL libraries are installed on the client. Most Linux versions come pre-packaged with OpenSSL. If you encounter a Linux user that does not have the required OpenSSL libraries, you can direct them to the following resource where they can be obtained and installed for free:

See http://www.openssl.org/related/binaries.html for details. (You can also advise users to compile their own version by directing them to the source at http://www.openssl.org/source/.) The required version is libssl.so.0.9.6b.

Related Documentation

- Task Summary: Configuring VPN Tunneling

Client Side Logging

VPN tunneling client-side logs are files that reside on the remote client containing sign-in, debug, and other statistical information you can use to troubleshoot potential issues with VPN tunneling. When you enable client-side logging for VPN tunneling users, the client records VPN tunneling events in a series of log files, continually appending entries each time a feature is invoked during subsequent user sessions. The resulting log files are useful when working with the support team to debug problems with VPN tunneling.

If VPN tunneling users turn client-side logging off, (even if logging is enabled on the system) the client does not record any new client-side log information. If the user turns on the logging function and the system is then configured to disable client-side logging, the client does not record any new client-side log information.

Related Documentation

- Task Summary: Configuring VPN Tunneling

- Enabling Client-Side Logging

VPN Tunneling Proxy Support

VPN tunneling provides support for remote clients using a proxy server to access the Internet (and Connect Secure via the Internet), as well as clients who do not need a proxy to access the Internet, but who access resources on an internal network through a proxy. VPN tunneling also provides support for clients accessing a Proxy Automatic Configuration (PAC) file that specifies client and system proxy settings enabling access to Web applications.

> ℹ️ NOTE: The VPN tunneling client does not support the use of the MS Winsock proxy client. Please disable the MS Winsock proxy client before running the VPN tunneling client. For more information, see http://www.microsoft.com/windowsxp/using/mobility/expert/vpns.mspx.

To address these varying methods of proxy implementation, VPN tunneling temporarily changes the proxy settings of the browser so that only traffic intended for the VPN tunneling session uses the temporary proxy settings. All traffic not intended for the VPN tunneling session uses the existing proxy settings.

> ℹ️ NOTE: The VPN tunneling client does not support the option to automatically detect proxy settings. You must choose to use either an automatic configuration script (PAC) or specify a proxy server. You cannot use both a proxy server and an automatic configuration script, together. You can define one or the other under the Proxy section in Users > Resource Policies > VPN Tunneling > Connection Profiles > *Profile*.

Whether split-tunneling is enabled or disabled, the system supports the following proxy scenarios:

- Using an explicit proxy to access Connect Secure

- Using an explicit proxy to access internal Web applications

- Using a PAC file to access Connect Secure

- Using a PAC file to access internal Web applications

Please note the following exceptions:

- The system does not support redirect downloads and therefore does not support the redirecting of the internal PAC file download.

- The system's dsinet client does not support SSL; you cannot obtain the internal PAC file from the SSL server.

- The system does not support "auto detect proxy". If both static proxy and "auto proxy script (pac)" are defined, it uses the static proxy configuration.

- The VPN tunneling profile does not have a static proxy exception field for internal proxy. If you require proxy exceptions, you can use a PAC file with proxy exception logic.

- The VPN tunneling client supports "auto proxy script (pac)" only when the configuration is the PAC file URL. If the URL is a redirect URL or IE proxy configuration script it is not supported.

When split-tunneling is enabled, VPN tunneling manages proxy settings in one of the following ways, depending on the method with which the proxy is implemented:

- For remote clients using a proxy server to access the Internet, all HTTP requests generated by the browser and intended for the system go through either an explicit proxy or a PAC file accessed by the remote client. Because the presence of an explicit proxy or access to a PAC file is already provisioned on the client-side, the client sets up the local, temporary proxy before attempting to establish a VPN tunnel.

- For remote clients using a proxy server to access the Internet, all HTTP requests generated by the browser and intended for the system go through either an explicit proxy or a PAC file accessed by the remote client. Because the presence of an explicit proxy or access to a PAC file is already provisioned on the client-side, the client sets up the local, temporary proxy before attempting to establish a VPN tunnel.

- When a remote client accesses a preconfigured HTTP-based PAC file, the client cannot access the PAC file until after a VPN tunnel is established. After a connection is established, the client accesses the PAC file, includes the PAC file contents in the local temporary proxy, and then refreshes the browser proxy setting.

Related Documentation

- Task Summary: Configuring VPN Tunneling

VPN Tunneling Quality of Service

To support quality of service (QoS) on your internal network via VPN tunneling, the system translates the "inner" IP packet header (for Application-layer packet encapsulation, for example) to the "outer" packet header, thus enabling Network layer-level packet prioritization. Routers in the network are then able to identify, prioritize, and appropriately forward VPN tunneling IPsec packets across the network. This feature helps ensure that you are able to support time-sensitive IP packet transmission and reception like IP video streams, for example.

NOTE: VPN tunneling QoS applies to UDP (IPsec) packets only. SSL packet encapsulation and forwarding behavior remains unchanged when you employ the QoS feature.

Related Documentation

- Task Summary: Configuring VPN Tunneling

VPN Tunneling Multicast Support

To enable streaming IP video broadcasts over the internal network, VPN tunneling features Internet Group Management Protocol (IGMP) gateway multicast proxy support.

ℹ️ NOTE: VPN tunneling does not support IGMP v2. If you are using VPN tunneling multicast support, and you are using L2 switches, make sure the switches support IGMP v3.

When users initiate a request to join a multicast group, the system initiates an IGMP join message to the local multicast router or switch on the client's behalf. In addition, the system stores the IGMP group request queries in its cache so that whenever a multicast router in the network polls the system for IGMP group information, it responds with its current collection of multicast user and group requests. If a router or switch does not receive a response from the system , the system's multicast group information is removed from the router or switch's forwarding table.

ℹ️ NOTE: VPN tunneling supports streaming media at up to 2 mbps on a single tunnel (megabits per second).

Related Documentation

- Task Summary: Configuring VPN Tunneling

About Split Tunneling Role Options

Figure 117 shows an unprotected network that contains the endpoint client and other unprotected resources, and a protected network that contains networks that can only be accessed through a VPN tunnel through Connect Secure.

Figure117: Example Network Scenario
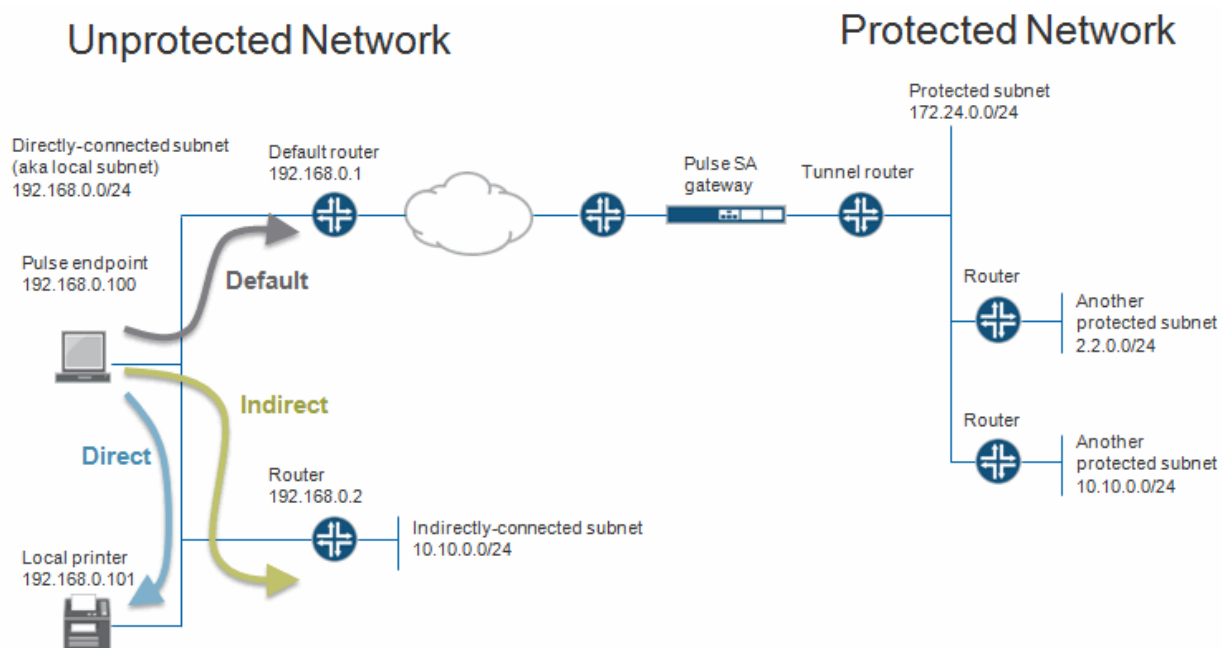
Before a VPN tunnel is created, there are three types of endpoint routes.

- Directly-connected subnet routes—Hosts on the directly-connected subnet can be reached without forwarding through a router. The ability to access these hosts is defined as *local subnet access*. 192.168.0.0/24 is an example of a directly-connected subnet route.

- Indirectly-connected subnet routes—These routes that have an explicit non-direct entry in the route table. Traffic must go through a router. 10.10.0.0/24 is an example of an indirectly-connected subnet route.

- Default route—this route if the destination is neither a direct-connected or indirect-connected subnet route. 0.0.0.0 is an example of a default route.

Enabling Split Tunneling

Options on the Users > User Roles > *Role Name* >VPN Tunneling page determine how the endpoint routes are modified when the VPN tunneling is established. See Figure 118.

Figure118: Split Tunnel User Role Options

Split Tunneling options are:

- **Enable**—Adds or modifies routes for specific subnets to go to the tunnel, allowing access to the protected subnets. Subnets are defined in the Users > Resource Policies > VPN Tunneling > Split-tunneling Networks window. In the case of subnet overlap (the specified split-tunnel subnet conflicts with an existing endpoint route), the Route Precedence option is used. For example, 2.2.2.0/24 goes through the tunnel. 10.10.0.0/24 is both a split-tunnel subnet and an indirectly-connected subnet. The Routing Table, defined below, defines how 10.10.0.0/24 is handled.

- **Disable**—Modifies the default route to go through the tunnel, allowing access to the protected network. For example, 0.0.0.0 now goes through the tunnel.

Defining the Route Precedence Options

The Route Precedence option determines how the directly-connected subnet routes and the indirectly-connected subnet routes are modified. This depends on whether split-tunneling is enabled.

If split tunneling is disabled:

- Tunnel routes have precedence—Both directly-connected subnet routes and indirectly-connected subnet routes go through the tunnel. Endpoints lose access to the unprotected subnets. See Figure 119.

Figure119: Tunnel Route Precedence Example
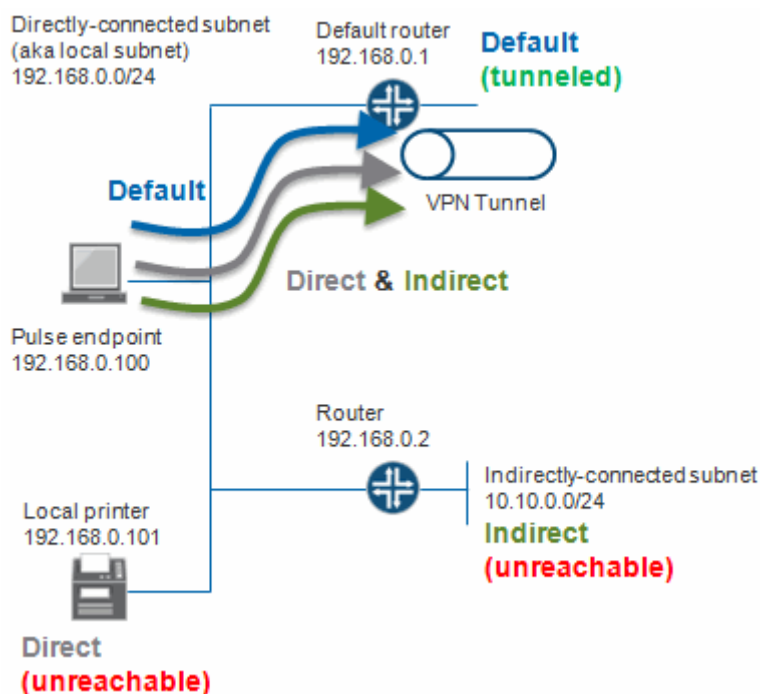
- Endpoint routes have precedence—Neither directly-connected subnet routes nor indirectly-connected subnet routes go through the tunnel. See Figure 120.

Figure120: Endpoint Route Precedence Example
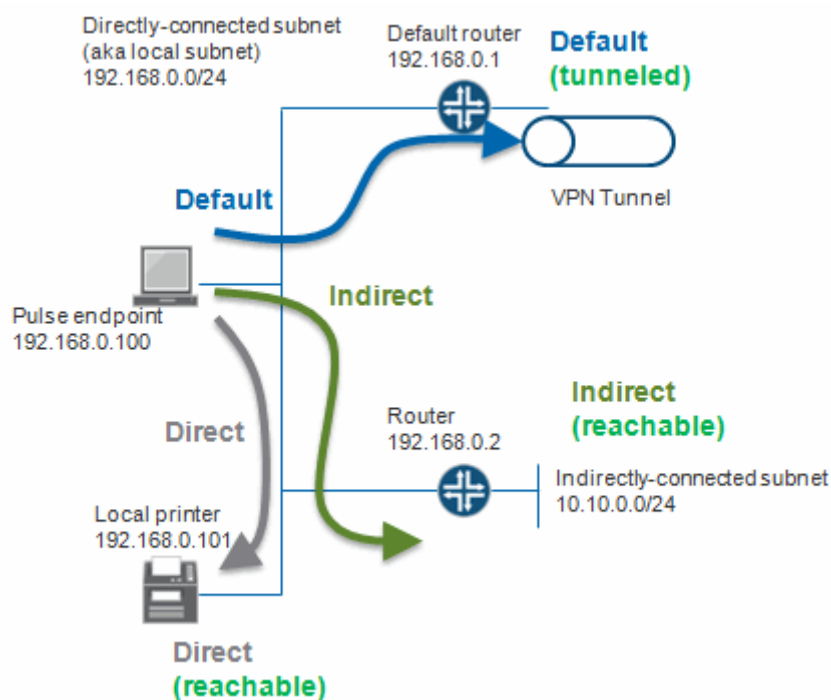
- Tunnel routes with local subnet access—Similar to the tunnel routes have precedence option except that directly-connected subnet routes do not go through the tunnel. Stated another way, tunnel all traffic except for traffic that is destined for the local subnet. This allows endpoints to access the local subnet. See Figure 121.

Figure121: Tunnel Routes with Local Subnet Access Example
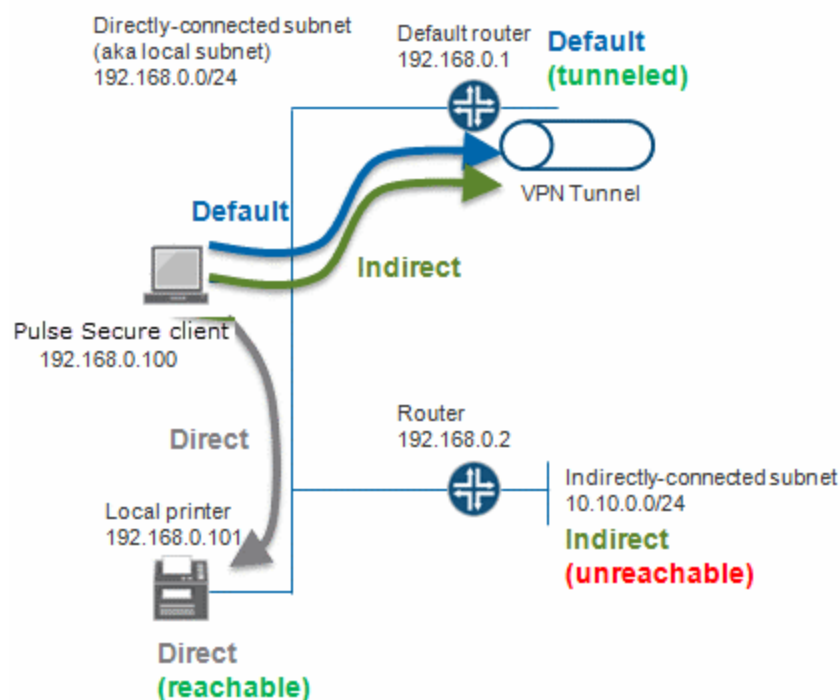
If split tunneling is enabled:

- Non-overlapped subnets are added to the route table. For example, 2.2.0.0/24 goes through the tunnel. See Figure 117.

- Overlapped subnets have 2 options:

  - Tunnel routes have precedence—These routes are modified to go through the tunnel and endpoints lose access to the unprotected overlapped subnet. For example, 10.10.0.0/24 goes through the tunnel. See Figure 117.

  - Endpoint routes have precedence—These routes do not go through the tunnel. For example, 10.10.0.0/24 does not go through the tunnel. See Figure 117.

Table 95 summarizes how split tunneling and route precedence works with directly-connected subnet routes, indirectly-connected subnet routes, and default routes.

Table95: Split Tunnel and Route Precedence

Split Tunnel	Route Precedence	Direct Endpoint Route	Indirect Endpoint Route	Default Endpoint Route	Split Tunnel Policy Routes
Disabled	Tunnel	Goes through tunnel.	Goes through tunnel.	Goes through tunnel.	NA
Disabled	Tunnel with local subnet access	Does not go through tunnel.	Goes through tunnel.	Goes through tunnel.	NA
Disabled	Endpoint	Does not go through tunnel.	Does not go through tunnel.	Goes through tunnel.	NA
Enabled	Tunnel	Routes are modified to go through the tunnel only if they overlap with the split-tunneling subnets. Otherwise the routes do not go through the tunnel.	Routes are modified to go through the tunnel only if they overlap with the split-tunneling subnets. Otherwise the routes do not go through the tunnel.	Does not go through tunnel	New routes are added to go through the tunnel.
Enabled	Tunnel with local subnet access	Does not go through tunnel.	Routes are modified to go through the tunnel only if they overlap with the split-tunneling subnets. Otherwise the routes do not go through the tunnel.	Does not go through tunnel.	New routes are added to go through the tunnel.
Enabled	Endpoint	Does not go through tunnel.	Does not go through tunnel.	Does not go through tunnel.	New routes are added to go through the tunnel.

Related Documentation

- Defining VPN Tunneling Role Settings

- Task Summary: Configuring VPN Tunneling

- Defining Split Tunneling Network Policies

- Creating VPN Tunneling Connection Profiles

Defining VPN Tunneling Role Settings

Use role-level settings to specify split-tunneling, auto-launch, auto-uninstall, Graphical Identification and Authentication (GINA) options, and the Riverbed Steelhead Mobile Controller (if any).

To specify VPN tunneling split-tunneling, auto-launch, auto-uninstall, and GINA installation options:

1. In the admin console, choose **Users > User Roles >** *Role Name* **> VPN Tunneling**.

2. Under General VPN Options, select one of the following options:

   - **Enable Split Tunneling**—This option activates split-tunneling and adds (or modifies) routes for specific subnets to go to the tunnel, allowing access to the protected subnets. The subnets are specified in the Users > Resource Policies > VPN Tunneling > Split-tunneling Networks window. In the case of subnet overlap (for example, the specified

split-tunnel subnet conflicts with an existing endpoint route), the Route Precedence option (described below) is used.

- **Disable Split Tunneling**—All network traffic from the client goes through the VPN tunnel, allowing access to the protected network. When the session is established, predefined local subnet and host-to-host routes that might cause split-tunneling behavior are removed, and all network traffic from the client goes through the VPN tunnel. With split tunneling disabled, users cannot access local LAN resources during an active VPN session.

3. Under **Pulse Secure client options**, select:

- **Route precedence** – This option defines how the directly-connected subnet routes and the indirectly-connected subnet routes are modified. The exact effect depends on whether split-tunneling is enabled.

  - **Tunnel Routes**—The route table associated with the Pulse virtual adapter take precedence. Pulse overwrites the physical interface routes if there is conflict between the Pulse virtual adapter and the physical adapters. Pulse restores the original routes when the connection is ended.

  - **Tunnel Routes with local subnet access (Pulse on Windows and Mac OS X only)**— Network traffic addressed to the networks defined in the split tunnel resource policies goes through the VPN tunnel. Network traffic that is addressed to the directly-connected (local) subnet goes to the local subnet. The default route is set to the local subnet so all other network traffic is subject to the original endpoint routing table.

  - **Endpoint Routes**—The route table associated with the endpoint's physical adapter take precedence.

    > ⓘ NOTE: Setting route precedence to **Endpoint Routes** allows users to access the local subnet regardless of whether split tunneling is enabled or disabled.

- **Route Monitor** – Specify whether you want route monitoring enabled.

  - **Yes** – VPN tunneling ends the connection only if the route change affects the VPN tunnel traffic. For example, if the route metric is changed higher, it should not disconnect VPN tunneling.

  - **No** – Route tables are allowed to change on the client endpoint.

- **Traffic Enforcement**—When Traffic Enforcement is enabled, Pulse creates rules on the endpoint's firewall (Mac and Win) that ensure that all traffic conforms to the split tunneling configuration. For example, a local program might bypass the routing tables and bind traffic to the physical interface instead of allowing it to go through the Pulse virtual interface. If you enable traffic enforcement, you ensure that all traffic is bound by the split tunneling configuration.

  Yes – All traffic is bound by the split tunneling configuration.

  No – Network traffic that bypasses the local routing tables is not bound the split tunneling configuration.

- **Enable TOS Bits Copy**—Select this option to control the client behavior in networks that employ quality of service (QoS) protocols. When you enable this check box, the Pulse Secure client copies IP Type of Service (TOS) bits from the inner IP header to outer the IP Header. Note that enabling this option might require a reboot of the client endpoint when the client software is installed for the first time on Windows endpoints. Pulse Secure clients support TOS bit copy only for IPsec transport and not for SSL transport.

- **Multitask**—Select this option if you want VPN tunneling to operate in multicast mode.

- **Auto-launch**—Select this option to activate VPN tunneling automatically when the endpoint is started.

4. Under **Options for Pulse Secure client on Windows**, select:

- **Launch client during Windows Interactive User Logon**—When this option is enabled, the Pulse Secure client starts when the user logs into Windows. Note that this setting is not the same as the Pulse connection settings that control machine authentication and credential provider authentication. Choose one of the following options:

  Require client to start when logging into Windows

  Allow user to decide whether to start client when logging into Windows

5. For Session Scripts, specify the following:

- **Windows: Session start script**—Specify a script (.bat, .cmd, or .exe) to run for users assigned to the role after Pulse connects with Connect Secure. For example, you can specify a script that maps network drives on an endpoint to shares on protected resources.

- **Windows: Session end script**—Specify a script (.bat, .cmd, or .exe) to run for users assigned to the role after Pulse disconnects from Connect Secure. For example, you can specify a script that disconnects mapped network drives. If there is no start script defined, or the start script has not been run, the end script does not run.

Options for Pulse Secure client on Mac apply only to Pulse and Network Connect on Apple OS X endpoints:

- **Mac: Session start script**—Specify a script (.bat, .cmd, or .exe) to run for users assigned to the role after Pulse connects with Connect Secure. For example, you can specify a script that maps network drives on an endpoint to shares on protected resources.

- **Mac: Session end script**—Specify a script (.bat, .cmd, or .exe) to run for users assigned to the role after Pulse disconnects from Connect Secure. For example, you can specify a script that disconnects mapped network drives. If there is no start script defined, or the start script has not been run, the end script does not run.

When VPN tunneling launches, start and end scripts are copied to the client and, upon session termination, are removed from the client. Scripts can be accessed locally or remotely via file share or other permanently-available local network resource. Macintosh clients only support running start and end script located on the local machine.

> NOTE: The client should be a member of the same domain as the remote server to allow VPN tunneling to copy start and end scripts. If the client

credentials are unknown to the server, the script copy fails, and VPN tunneling does not prompt the user to enter username and password.

Windows only supports scripts with the .bat or .cmd extension (referring to batch files, not the .cmd applications within MSDOS). To run a .vbs script, the user must have a batch file to call the .vbs script. Similarly, to run an .exe application (like C:\WINDOWS\system32\mstsc.exe), the user must have a batch file to call the .exe application.

The client makes a copy of the end script after the tunnel has been set up and stores the script in a temporary directory to ensure that, if the network connection were to fail, the end script can still be used to terminate the VPN tunnel session.

- Select the **Skip if Windows Interactive User Logon Enabled** option to bypass the specified Windows session start script.

  If the client signs in to their Windows Domain via the GINA/Credential Provider automatic sign-in function, a script is executed by the Windows client. In this case, the sign-in script may be identical to the specified VPN Tunneling start script. You can use this option, therefore, as a way to avoid executing the same script twice.

6. Under Advanced Options for Network Connect, select options for legacy VPN Tunneling clients. This section is not applicable for Pulse.

- **NC Client Check** – Select this option to check the integrity of your VPN Tunneling client components prior to starting a VPN tunnel. Selecting this option requires you to have either administrator rights on the client or Juniper Installer Service (JIS) running on the client. If the integrity check fails, error 23787 "Cannot start the VPN Tunneling service. Please re-install VPN Tunneling." appears. This option is applicable only when you use a VPN Tunneling client. If you are using a Pulse Secure client, this option is ignored.

- **Do not allow client proxy** – Select this option to prevent a VPN Tunneling tunnel from being created when a client proxy is present. Error 30572, "VPN Tunneling unable to connect because a client proxy is not allowed", appears if the user tries to start a VPN Tunneling tunnel. This option is applicable only when you use a VPN Tunneling client. If you are using a Pulse Secure client, this option is ignored.

- **Auto-Uninstall Network Connect** – Select this option to automatically remove Network Connect from the client when the user logs out of their session.

- **Enable FIPS compliant Network Connect** – (Windows client systems only) Applicable only to Network Connect on FIPS devices, this option allows administrators to control NC-FIPS-compliance on FIPS devices. Deselect this option to allow VPN Tunneling to connect from unmanaged endpoints to Connect Secure (running software version 7.2 or later) with a private CA issued certificate as the device certificate. This option is applicable to both SSL and ESP modes. This option is not supported on non-Windows client systems.

  NOTE: This option is available only on FIPS devices, and is enabled automatically when FIPS is enabled (see "Enabling FIPS Level 1 Support"). If this option is disabled manually, and FIPS support is disabled and then reenabled, this option remains disabled for the role.

| Tunnel Type | Enable FIPS compliant Network Connect | |
	Enabled	Disabled
SSL	SSL default	SSL default
	Unmanaged endpoints are not supported	Unmanaged endpoints are supported
ESP	Dishonors ESP, always `gives SSL tunnel	ESP default, fallback to SSL
	Unmanaged endpoints are not supported	Unmanaged endpoints are supported

This option affects the following clients:

- Network Connect launched from a browser window

- Network Connect launched from a mini-browser window

- nclauncher

- GINA/credential provider

7. Under Options for the Steelhead Mobile Controller, select the data acceleration options for Android devices:

- **Enable Steelhead Optimization** – Select this option if you have one or more Riverbed Steelhead appliances.

- **Server Location** – Enter the IP address or hostname of your Steelhead Mobile Controller.

  NOTE: The Steelhead IP address replaces the destination address in client packets. You may want to add this address to your access control lists.

8. Click **Save Changes**.

Related Documentation

- About Split Tunneling Role Options

- Task Summary: Configuring VPN Tunneling

- Defining Split Tunneling Network Policies

- Creating VPN Tunneling Connection Profiles

About VPN Tunneling Resource Policies

VPN tunneling resource policies specify a variety of session parameters you can use to determine the method of access for remote clients. You can configure the following types of resource policies and apply them to one or more user roles:

- **Access resource policies**—This policy type specifies which resources users may access when using VPN tunneling, such as Web, file, and server machines on the corporate intranet.

- **Packet logging resource policies**—This policy type allows you to compile client-side VPN tunneling packet logs on the system to help diagnose and resolve connection issues. Connection profiles resource policies—This policy type specifies which option (DHCP or system-managed IP address pool) The system uses to assign an IP address to the client-

side VPN tunneling agent. You can also use this feature to specify the transport protocol and encryption method for the VPN tunneling session.

- **Split Tunneling resource policies**—This policy type enables you to specify one or more network IP address/netmask combinations for which the system handles traffic passed between the remote client and the corporate intranet.

A few notes about specifying resources for a VPN tunneling resource policy:

- You cannot specify a hostname for a VPN tunneling resource policy. You can only specify an IP address.

- You can specify protocols (such as tcp, udp, icmp) for VPN tunneling. For all other access feature resource policies, specifying protocols is not supported.

- If the protocol is missing, all protocols are assumed. If a protocol is specified, then the delimiter ":// " is required. No special characters are allowed.

- You cannot mix port lists and port ranges, such as 80, 443, 8080-8090 for VPN tunneling resource policies.

- If you specify a port, you must specify a protocol.

- If the port number is missing, the default port * is assigned for http.

Related Documentation

- Defining VPN Tunneling Access Control Policies
- Creating VPN Tunneling Connection Profiles
- Defining Split Tunneling Network Policies
- Specifying Resources for a Resource Policy

Defining VPN Tunneling Access Control Policies

Use the VPN Tunneling Access Control tab to write a resource policy that controls resources users can connect to when using VPN tunneling.

To write a VPN tunneling access resource policy:

1. In the admin console, choose **Users > Resource Policies > VPN Tunneling > Access Control**.

2. On the Access Control page, click **New Policy**.

3. On the New Policy page, enter:

    - A name to label this policy.

    - A description of the policy. (optional)

4. In the Resources section, specify the resources to which this policy applies.

   ⓘ   NOTE: When a packet is fragmented, fragment #1 contains more information than all subsequent fragments. Fragment #1 contains the IP address, protocol, and port information. All subsequent fragmented packets contain just the IP address and protocol information. Therefore the VPN Tunneling ACL evaluates

the first packet fragment different from the subsequent packet fragments. For the subsequent packet fragments, the system applies the VPN Tunneling ACL based on just the IP address and protocol since the port number is not available.

5. In the Roles section, specify:

- **Policy applies to ALL roles**—To apply this policy to all users.

- **Policy applies to SELECTED roles**—To apply this policy only to users who are mapped to roles in the Selected roles list. Make sure to add roles to this list from the Available roles list.

- **Policy applies to all roles OTHER THAN those selected below**—To apply this policy to all users except for those who map to the roles in the Selected roles list. Make sure to add roles to this list from the Available roles list.

6. In the Action section, specify:

- **Allow access**—Select this option to grant access to the resources specified in the Resources list.

- **Deny access**—Select this option to deny access to the resources specified in the Resources list.

- **Use Detailed Rules**—Select this option to define resource policy rules that put additional restrictions on the specified resources.

7. Click **Save Changes**.

8. On the Access Policies page, order the policies according to how you want to evaluate them. Keep in mind that once the system matches the resource requested by the user to a resource in a policy's (or a detailed rule's) Resource list, it performs the specified action and stops processing policies.

Related Documentation

- Writing a Detailed Rule for Resource Policies

- Task Summary: Configuring VPN Tunneling

Creating VPN Tunneling Connection Profiles

Use the **Users > Resource Policies > VPN Tunneling > Connection Profiles** page to create VPN tunneling connection profiles. When the system receives a client request to start a VPN tunneling session, it assigns an IP address to the client-side agent. Thesystem assigns this IP address based on the DHCP Server or IP Address Pool policies that apply to a user's role. In addition, this feature allows you to specify the transport protocol, encryption method, and whether or not to employ data compression for the VPN tunneling session.

Nodes in a multi-site cluster share configuration information, which means that devices in different networks share an IP address pool. Since any node may receive the client request to start the VPN tunneling session, you need to specify an IP filter for that node that filters out only those network addresses available to that node. When the cluster node receives a request to create a VPN tunnel, it assigns the IP address for the session from the filtered IP address pool.

NOTE: Release 7.3 and Release 7.4 do not support IP filters for IPv6 addresses.

To write a VPN tunneling connection profile:

1. In the admin console, choose **Users > Resource Policies > VPN Tunneling > Connection Profiles**.

2. On the Connection Profiles page, click **New Profile** and configure the settings described in Table 96.

3. Save the configuration.

4. On the Connection Profiles page, order the profiles according to how you want to evaluate them. Keep in mind that once the system matches the resource requested by the user to a resource in a profile's (or a detailed rule's) Resource list, it performs the specified action and stops processing profiles. See Table 96.

Table96: VPN Tunneling Connection Profile Settings

Setting	Guidelines
Name	A name to label this policy.
Description	A description of the policy (optional).
IPv4 address assignment	
DHCP servers	Specify the hostname or IP address of a network Dynamic Host Configuration Protocol (DHCP) server responsible for handling client-side IP address assignment.
	You can specify up to three DHCP servers by listing each one on a separate line. When multiple DHCP servers are listed, the system sends a DHCP Discover message to all listed DHCP servers and then waits five seconds for a response. If multiple DHCP servers respond, the system chooses the one with the longest lease period.
	The system sends a DHCP release packet to the DHCP server when the VPN tunneling session ends.
	DHCP provides a framework for passing configuration information to hosts. Configuration parameters and other control information are carried in tagged data items that are stored in the options field of the DHCP message. You can specify the DHCP options to forward by entering the option number, its value and type and then clicking Add. For a complete list of DHCP options, see the "RFC2132 - DHCP Options and BOOTP Vendor Extensions" article available on the Internet. To delete an option, select the check box next to the option number then click the **Delete** button.
DHCP options	By default, the client's hostname is sent by Connect Secure to the DHCP server in the DHCP hostname option (option12.) Passing the useruid in the DHCP hostname option is no longer supported. As an alternative, you can configure the following entry in the DHCP options table. For example:
	option number=12, option value=<username><authMethod>, option type=String
	Or you can pass a value by adding an entry in the DHCP options table for hostname with whatever value you want. For example:
	option number=12, option value=foo, option type=String

IPv4 address pool	Specify IP addresses or a range of IP addresses for the system to assign to clients that run the VPN tunneling service. Use the canonical format: ip_range. The last component of the IP address is a range delimited by a hyphen (-). No special characters are allowed. The ip_range can be specified as shown in the following list: • a.b.c.d — Specifies a single IP address. • a.b.c.d-e.f.g.h — Specifies all IP addresses from the first address to the last address, inclusive. • a.b.c.d-f.g.h — An abbreviated form that specifies the range a.b.c.d through a.f.g.h • a.b.c.d-g.h — An abbreviated form that specifies the range a.b.c.d through a.b.g.h. • a.b.c.d-h — An abbreviated form that specifies the range a.b.c.d through a.b.c.h. • a.b.c.d/mask — Specifies all addresses in a network. For example, to allocate all addresses in the range 172.20.0.0 through 172.20.3.255, specify 172.20.0.0-3.255. Or, to allocate all addresses in a class C network, specify 10.20.30.0/24. NOTE: Be sure to specify a sufficient number of addresses in the IP address pool for all of the endpoints in your deployment. When all of the addresses in the pool have been assigned to endpoints, additional endpoints are unable to obtain a virtual IP address and are blocked from accessing protected resources. The system logs a message in the Event log when an IP address cannot be assigned to an endpoint. We recommend that you set up your network so that the client-side IP address pool, or the DHCP server specified in the VPN tunneling connection profile, resides on the same subnet as Connect Secure. If your network topology dictates that the system internal IP interface and the IP address pool or DHCP server reside on different subnets, you need to add static routes to your intranet's gateway router(s) to ensure that your Enterprise resources and Connect Secure can see each other on the internal network. If you are running a multi-unit cluster across a LAN, make sure that the IP address pool contains addresses that are valid for each node in the cluster. Then, configure an IP filter for each node to apply to this IP address pool. The system does not support a common IP address pool for VPN tunneling for an Active/Active cluster. In A/A VPN tunneling deployments, we recommend that you split the IP pool into node-specific subpools. Furthermore, you are advised to perform static route configuration on the backend router infrastructure in a coordinated fashion, with static routes to each subpool pointing to the internal IP address of the hosting cluster node as the next-hop gateway. IP address pool also supports attribute substitution. For example, you can enter a RADIUS role mapping attribute in this field, such as <userAttr.Framed-IP-Address>.
IPv6 address assignment	
Enable IPv6 address assignment to clients	Select this option to enable IPv6 connections. Release 7.4 does not support assignment by a DHCPv6 server. You must configure a static IPv6 address pool. NOTE: IPv6 must be enabled on internal interface for IPv6 addresses to be allocated to clients.
IPv6 address pool	Specify IPv6 address ranges for this profile, one per line. Like the IPv4 address pool, the configuration supports entering ip_range values. We recommend using the IPv6 network prefix / netmask style (such as 2001:DB8::6:0/112).
Connection settings	

Transport	Select one of the following options for transport, encryption, and compression settings:

- **ESP**—Use a UDP encapsulated ESP transfer method to securely transfer data between the client and Connect Secure. ESP uses an LZO compression algorithm. You can use the default settings or configure data transfer parameters by defining the UDP port, ESP-to-SSL fallback time-out value, and ESP encryption key lifetime values.
- **SSL**—Use the standard SSL transport method. SSL uses a deflate compression method. In SSL mode, compression is controlled by the Enable GZIP compression option on the System Maintenance Options page.

NOTE: To support IPv6 connections, be sure to set MTU greater than 1380. We recommend 1500. If the MTU value on the external interface is lower than 1380 and IPv6 address assignment is enabled, the transport setting for the connection profile is ignored. To avoid IP fragmentation, the session falls back to SSL mode for both IPv6 and IPv4 traffic.

If you select ESP mode, configure the following transport and compression settings:

- **UDP port**—Port through which you intend to direct UDP connection traffic. The default port number is 4500.

NOTE: Whether you specify a custom port number or choose to use the default port number (4500), you must also ensure that other devices along the encrypted tunnel allow UDP traffic to pass between Connect Secure and the clients. For example, if you employ an edge router and a firewall between the Internet and your corporate intranet, you must ensure that port 4500 is enabled on both the router and the firewall and that port 4500 is configured to pass UDP traffic.

IKEv2 uses port 500 exclusively. Do not configure port 500 in your VPN Tunneling profiles.

- **ESP to SSL fallback timeout**—Period of time (in seconds) to fall back to the SSL connection already established following UDP connection failure. The default is 15 seconds.

NOTE: A nonconfigurable idle timeout of 60 seconds also affects when fallback occurs. After the tunnel is established through ESP, the client sends keepalives after 60 seconds of inactivity on the ESP channel (the idle timeout). The total time to fallback is therefore the idle timeout (60 seconds) plus the fallback timeout. For example, if ESP to SSL fallback timeout is set to 25 seconds, it takes approximately 60+25 or 85 seconds for the VPN tunneling client to switch.

- **Key lifetime (time based)**—Period of time (in minutes) the system continues to employ the same ESP encryption key for this connection profile. Both the local and remote sides of the encrypted transmission tunnel use the same encryption key only for a limited period of time to help prevent unauthorized access. The default is 20 minutes.
- **Key lifetime (bytes transferred)**—Maximum amount of data that is transferred on the tunnel for an ESP encryption key. The default is 0 bytes, meaning no limit.

NOTE: When either of the key lifetime limits is reached, a new key is exchanged between Connect Secure and the client. The reason for changing keys is to help prevent unauthorized access, however, changing the encryption key too frequently can increase CPU overhead on the system.

- **Replay Protection**—Activates replay protection. When enabled, this option protects against hostile "repeat attacks" from the network. When packets arrive from the client, the system checks the IP header information to verify that a packet featuring the same IP header information has not already been received. If one has been received, the packet is rejected. This option is enabled by default.

  If you activate the Enable TOS Bits Copy option, IP packets with different TOS bits may be reordered when passing through gateway routers on your network. To ensure that any packets received out of order are not automatically dropped when they reach the system, you can disable the Replay Protection option.

NOTE: We recommend that you leave replay protection enabled if you are not expecting more than one source of packets from the client (for example, if only one application is transmitting and receiving traffic over the VPN tunnel).

- **Compression**—Use compression for the secure connection. Compression is useful for a slow link but may cause issues in extremely large deployments since extra cycles are spent compressing the data.

If you have selected ESP, select one the following encryption settings:

- **AES128/MD5 (maximize performance)**—Uses Advanced Encryption Standard (AES) 128-bit encryption on the data channel and the MD5 authentication method for VPN tunneling sessions.
- **AES128/SHA1**—Uses AES 128-bit encryption on the data channel and the SHA1 authentication method during VPN tunneling sessions.
- **AES256/MD5**—Uses AES 256-bit encryption on the data channel and the MD5 authentication method for VPN tunneling sessions.
- **AES256/SHA1 (maximize security)**—Uses AES 256-bit encryption on the data channel and the SHA1 authentication method during VPN tunneling sessions.

NOTE: The MD5 authentication algorithm creates digital signatures. The MD5 authentication method translates an input string (like a user's ID or sign-in password, for example) into a fixed, 128-bit fingerprint (also called a "message digest") before it is transmitted to or from the system.

DNS settings	

IVE DNS Settings	In the DNS Settings section, select an option that determines the settings sent to the client:
	
	• **IVE DNS Settings**—Send the system DNS settings.
	
	• **Manual DNS Settings**—Override standard DNS settings with the settings you provide:
	
	    • **Primary DNS**—Enter the IP address for the primary DNS.
	    • **Secondary DNS**—Enter the IP address for the secondary DNS.
	    • **DNS Domain(s)**—Enter the DNS domain(s), such as "yourcompany.com, yourcompany.net".
	    • **WINS**—Enter the WINS resolution name or IP address.
	
	• **DHCP DNS Settings**—Send to the client the values the DHCP server sends to Connect Secure. There is no fallback to the DNS settings if the DHCP Server does not send any values.

Auto-allow	Select **Auto-allow IP's in DNS/WINS settings (only for split-tunnel enabled mode)** if you want to create an allow rule for the DNS server, For example, if you have defined policies to allow requests from IP address 10.0.0.0 but your DNS server has an address of 172.125.125.125 the DNS server requests will be dropped. If you select this option, the system creates a rule to allow the DNS requests.

DNS search order	Select the DNS server search order. Applicable only if split tunneling is enabled:
	• Search client DNS first, then the device
	• Search the device's DNS servers first, then the client
	• Search device DNS only.

NOTE: DNS search order does not work with iOS clients. The DNS name resolution fields (located on the System > Network > Overview window) must be configured, otherwise all DNS queries will go to the client's DNS server.

Pulse Secure client 5.0 and greater supports all DNS search order options. Prior versions of Pulse Secure client support only **Search client DNS first, then the device** and **Search the device's DNS servers first, then the client**.

For the **Search client DNS first, then the device** and **Search the device's DNS servers first, then the client** options, DNS configured on the system are added to the end user's system along with the existing DNS already available on the end user's system. So either the device DNS servers or client DNS servers get precedence at the end user's systems.

When the **Search device DNS only** option is selected, DNS on the end user's system are replaced with device DNS. This option is recommended to avoid ISP's DNS hijacking. Note that this option is applicable only for Windows platforms; non-Windows clients will use the **Search the device's DNS servers first, then the client** search order if this option is selected. When using this option, you must ensure that packets to the system DNS are going through the tunnel. To do this, add the required routes to the split tunnel networks policy (Users > Resource Policies > VPN Tunneling > Split-Tunneling Networks), or select the **Auto-allow IPs in DNS/WINS settings** option.

For the **Search device DNS only** option, the client software (Pulse or Network Connect), removes the DNS information of the available adapters on the client system after the tunnel is created. Once the tunnel is created, the client does not monitor the presence of new adapters and does not monitor if changes are made to the DNS settings of existing adapters. Because of this, the **Search device DNS only** option may not work properly if any of the following occurs after the tunnel is created:

• A new interface appears with a DNS server that does DNS hijacking.

• A third-party application adds DNS to the adapters whose DNS was removed by the client as part of the tunnel set up process.

• Third-party applications change the TCP/IP option from "Use the following DNS servers" to "Obtain DNS servers automatically" for those adapters whose DNS was removed by the client software as part of the tunnel set up process.

• End users enable the interfaces that are in the disabled state during the tunnel set up process.

Proxy Server Settings

Proxy server settings	Select one of the following options: • **No proxy server**—Specifies that the new profile requires no proxy server. • **Automatic (URL for PAC file on another server)**—Specify the URL of the server on which the PAC file resides, and the frequency (in minutes) with which the client polls the server for an updated version of the PAC file. You can configure VPN tunneling to check for an updated PAC files as often as every 10 minutes. The default (and minimum) update period is 10 minutes. The PAC file should reside on a Web server, not on the local PC. The PAC file update method runs on a 10 minute interval. Specifying a frequency update period that is a multiple of 10 will get an exact result. If you specify the update frequency at a value that is not a multiple of 10, it is rounded up to the next interval. For example, if you specify the update frequency at 15 minutes, the system updates a PAC file every 20 minutes. NOTE: VPN tunneling limits the size of internal (server side) PAC files. The logical maximum size is 256 KB. The actual maximum size that can be used in your deployment might be smaller, reduced according to the size of other VPN tunneling settings in use, such as the number of split tunnel networks and DNS suffix entries. • **Manual configuration**—Specify the IP address or the hostname of the server and provide the port assignment. • **Preserve client-side proxy settings**—By default, VPN tunneling may change proxy settings when needed. For example, VPN tunneling may temporarily change the proxy settings of the browser so that traffic intended for the VPN session uses the temporary proxy settings. Select the Preserve client-side proxy settings option to prevent the client-side proxy settings from being overridden by VPN tunneling. If you select this option, HTTP and FTP traffic path can change after VPN tunneling establishing the connection. Please analyze the proxy logic and split-tunnel option, and make sure it directs the traffic as intended. • **Disable client-side proxy settings**—Disables the client's proxy settings after the VPN tunnel is established. In the use case where the client proxy configuration (proxy.pac) is hosted on a LAN server and users are outside the office network, proxy.pac is not accessible and users access the Internet directly. However, after a VPN tunnel is established, proxy.pac becomes accessible, and that causes all Internet requests to go through the tunnel to the proxy server. When you select Disable client-side proxy settings, client requests are served through the Pulse server directly. When the tunnel is disconnected, the client proxy settings are restored.
Roles	Specify one of the following options: • **Policy applies to ALL roles**—To apply this policy to all users. • **Policy applies to SELECTED roles**—To apply this policy only to users who are mapped to roles in the Selected roles list. Make sure to add roles to this list from the Available roles list. • **Policy applies to all roles OTHER THAN those selected below**—To apply this policy to all users except for those who map to the roles in the Selected roles list. Make sure to add roles to this list from the Available roles list.

Related Documentation

- Configuring VPN Tunneling for Use on a Virtualized Connect Secure

- Task Summary: Configuring VPN Tunneling

Defining Split Tunneling Network Policies

Use the Split Tunneling Network tab to write a VPN tunneling resource policy that specifies one or more network IP address/netmask combinations for which the system handles traffic passed between the remote client and the corporate intranet. You can also specify traffic that should not pass through the VPN tunnel.

When split-tunneling is used, VPN tunneling modifies routes on clients so that traffic meant for the corporate intranet networks flow through the tunnel and all other traffic goes through the

local physical adapter. The system tries to resolve all DNS requests through the physical adapter first and then routes those that fail to the VPN tunneling adapter.

For example:

- If split tunnel is disabled, all split tunnel configuration is ignored, including the exclude route. The default route goes to the tunnel allowing access to the protected network.

- Split tunneling is enabled and the included route contains 10.204.64.0/18 and the exclude traffic contains 10.204.68.0/24. In this scenario, networks from 10.204.64.0/18 to 10.204.127.0/18 will pass through the VPN tunnel with the exception of the 10.204.68.0/24 network, which will not pass through the VPN tunnel.

- If split tunneling is enabled and the include route contains 10.204.64.0/24 (subnet of the excluded route) and the exclude route contains 10.204.64.0/18 (super set of the included route) then the included network's traffic will still be routed through the VPN tunnel.

> **ⓘ** NOTE: If split tunneling is enabled and there are no include routes configured to be sent to the client, VPN tunneling adds a default route to send traffic through the tunnel.

In addition to using subnets to define the traffic flow, the split tunneling resource policy supports a dynamic per session resource list based on user attributes. The authentication server can be any type, but must able to pass user attributes during authentication or authorization. Below is one example scenario. The exact steps depend on your implementation and will vary from this example.

1. A user enters their credentials to initiate a session.

2. Connect Secure contacts the back-end authentication server, for example using a RADIUS Access-Request message.

3. The RADIUS server checks the policies and contacts other back-end resources (if configured) to authenticate the user and to retrieve parameters for the user session.

4. The back-end server manager or policy manager sends the attribute list to the RADIUS server along with a list of hosts and IP subnets. Note that Steps 3 and 4 vary depending on your deployment.

5. The RADIUS server returns the list of internal hosts and subnets to the system as part of the RADIUS Access-Accept message.

6. The system is configured with the split tunneling resource policy for that user session based on the received subnet information from the RADIUS server and returns the policy to the client. The client uses this information to make the local split tunnel decisions.

To write a split tunneling networks resource policy:

1. In the admin console, choose **Users > Resource Policies > VPN Tunneling > Split-tunneling Networks**.

2. On the Connect Split Tunneling page, click **New Policy**.

3. On the New Policy page, enter:

   - A name to label this policy.

- A description of the policy (optional).

4. In the Resources section, specify:

   - One or more network IP address/netmask combinations for which the system handles traffic passed between the remote client and the corporate intranet. You may also use the '/' (slash) notation to specify these networks.

   - A user attribute, for example **<userAttr.Framed-Route>**, to send to the back-end authentication server.

   You can specify both IP address/netmask combinations and user attributes in the Resources section.

   Please note the following:

   - If a user attribute is configured in the resource list, it is dynamically resolved from the user session data.

   - Invalid entries (including "*") are ignored.

   - If a configured user attribute is not resolved at runtime for the resource lists, the device acts as if split tunneling is disabled.

5. In the Roles section, specify:

   - **Policy applies to ALL roles**—To apply this policy to all users.

   - **Policy applies to SELECTED roles**—To apply this policy only to users who are mapped to roles in the Selected roles list. Make sure to add roles to this list from the Available roles list.

   - **Policy applies to all roles OTHER THAN those selected below**—To apply this policy to all users except for those who map to the roles in the Selected roles list. Make sure to add roles to this list from the Available roles list.

6. In the Action section:

   - **Allow access**—Network IP address/netmask combinations specified in the Resources list pass through the VPN tunnel.

   - **Exclude access**—Network IP address/netmask combinations specified in the Resources list do not pass through the VPN tunnel.

   - **Use Detailed Rules** (available after you click 'Save Changes')—Select this option to define resource policy rules that put additional restrictions on the specified resources.

7. Click **Save Changes**.

8. On the Split Tunneling Policies page, order the policies according to how you want to evaluate them. Keep in mind that once the system matches the resource requested by the user to a resource in a policy's (or a detailed rule's) Resource list, it performs the specified action and stops processing policies.

Related Documentation

- Task Summary: Configuring VPN Tunneling

- Writing a Detailed Rule for Resource Policies

VPN Tunneling Resource Policy Configuration Use Case

This topic describes a real-world VPN tunneling application and the steps necessary to configure the appropriate resource policy providing access to remote users on the network.

Large financial institutions (also called Fortune Companies) require a robust client sign-in application like VPN tunneling to help provide remote employees seamless network connection to a large range of enterprise resources at the corporate headquarters. Often, remote users need to be able to access multiple applications on their laptops/client machines beyond simple e-mail or meeting scheduling applications. These remote super users or power users require secure, encrypted access to powerful server applications like Microsoft OutlookTM, OracleTM databases, and the RemedyTM case management system.

For this scenario, let's assume the following:

- There is a small collection of remote users who will all access their financial institution's enterprise resources via the same device.

- All the users have the same user_role_remote role assigned to their user ID

- Host Checker and Cache Cleaner are configured and verifying the users' machines upon logging into a device and launching their VPN tunneling sessions

- All users require access to three large servers at the corporate headquarters with the following attributes:

  - outlook.acme.com at IP address 10.2.3.201

  - oracle.financial.acme.com at IP address 10.2.3.202

  - case.remedy.acme.com at IP address 10.2.3.99

- Because the Company wants to manage their IP address pool very strictly, each device provides IP addresses to remote users (our particular device controls the IP addresses between 10.2.3.128 and 10.2.3.192)

- The company is interested in the most secure access possible, simultaneously accepting only the least possible amount of client down-time

To configure a VPN tunneling resource policy providing appropriate access to the Fortune Company remote users:

1. Create a new VPN tunneling resource policy where you specify the three servers to which you want to grant remote users access:

   1. In the Resources section, specify the IP address ranges necessary to allow access to the three servers (outlook.acme.com, oracle.financial.acme.com, and case.remedy.acme.com) separated by carriage returns.

      udp://10.2.3.64-127:80,443
      udp://10.2.3.192-255:80,443

      i   NOTE: Configuring your resource as 10.1.1.1-128:* is not supported. Doing so will result in an error.

2. In the Roles section, select the **Policy applies to SELECTED roles** option and ensure that only the "user_role_remote" role appears in the Selected roles list.

3. In the Action section, select the **Allow access** option.

4. Create a new VPN tunneling connection profile where you define the transport and encryption method for the data tunnel between the client(s) and system:

   1. In the IP address assignment section, select the **IP address pool** option and enter 10.2.3.128-192 in the associated text field.

   2. In the Connection Settings section, select the **ESP transport** option and the AES/SHA1 encryption option.

   3. In the Roles section, select the **Policy applies to SELECTED roles** option and ensure that only the "user_role_remote" role appears in the Selected roles list.

Related Documentation

- Task Summary: Configuring VPN Tunneling

- VPN Tunneling Execution

About VPN Tunneling Bandwidth Management Policies

Bandwidth management controls the rate of traffic sent or received on a network interface. Bandwidth management discards excess packets and ensures that a user is allocated a specified amount of bandwidth. Traffic less than or equal to the specified rate is guaranteed to be sent. Traffic exceeding the rate is either dropped or delayed.

The total guaranteed bandwidth and spare bandwidth amounts are tracked and updated as users log in and out. Spare bandwidth is defined as the administrator-configured maximum minus the total guaranteed bandwidth for logged-in users.

Guaranteed bandwidth and maximum bandwidths are defined at the role level. This limit applies to each user in the role and ensures that each user receives at least the guaranteed amount of bandwidth but no more than the configured maximum amount. When users are mapped to multiple roles, the higher limit is used. If you do not define a guaranteed bandwidth to a role, users in that role can still log in, but they are not guaranteed any bandwidth. That is, their guaranteed bandwidth is set to zero.

Bandwidth management also applies to IVS. The administrator configures the total guaranteed bandwidth for each IVS and configures the limits for roles within each IVS. The sum of the total guaranteed bandwidths for every IVS must be less than or equal to the maximum bandwidth of the appliance. The sum of all VPN tunneling maximum bandwidths of all IVSs must be less than or equal to the VPN tunneling maximum bandwidth. Be sure to set the bandwidth for both Connect Secure (System > Network > Overview) and the IVS (System > Virtual Systems > root).

i NOTE: If you use the same VLAN across multiple IVS systems, Bandwidth Management is not supported.

To ensure the system does not allow more bandwidth than the total available, the ability to start VPN tunnels is restricted. Users can start a VPN tunnel only if the guaranteed bandwidth for

their role is available. Once users start a session, they are never dropped due to bandwidth restrictions. A privilege level controls this restriction as shown in Table 97.

Table97: Privilege Levels and Percent of Maximum Bandwidth

Privilege Level	Percent of Maximum Bandwidth
Low	Limited to 50%
Medium	Limited to 75%
High	Limited to 90%
Maximum	Limited to 100%

For example, users assigned to a low privilege level are able to launch a VPN tunnel if the total current bandwidth usage is less than 50% of the configured Maximum Bandwidth. Users assigned to the maximum privilege level are able to launch a VPN tunnel at any time as long as there is any system bandwidth available.

When a user attempts to launch a VPN connection, the sum of the Guaranteed Minimum Bandwidth of all open VPN connections is divided by the configured Total Bandwidth. If the resulting value is less than the configured privilege level of this user, then the user's VPN connection is established. Otherwise, the connection request is denied. For example, if the user's privilege is 75% and the calculated current consumption is 70%, the user's VPN connection is established. If the calculated current consumption is 80%, the user's connection request is denied and the user receives a 23791 error code.

> **NOTE:** We recommend that average employees be given Low or Medium privilege levels. Higher privilege employees can be assigned the Maximum privilege level to ensure intranet access as long as there is bandwidth available.

If a user does not have the bandwidth to set up any VPN tunnels, the user can still log in but is restricted in what they can do. For example, they may only be able to access web e-mail, etc.

A guaranteed minimum bandwidth is the bandwidth a user gets once a VPN connection is established. If the remaining VPN bandwidth is smaller than the guaranteed minimum bandwidth, the user's VPN connection request is denied and the user receives an 23791 error code. The Guaranteed Minimum Bandwidth must be smaller than the Maximum Bandwidth.

Maximum bandwidth is the bandwidth a user can use through the VPN connection. This is a limit on how much the user can use if there is bandwidth available. For example, if the user's maximum bandwidth is 100 kbps, the user cannot use more than 100 kbps regardless how much available bandwidth.

Statistics for bandwidth management are recorded in the system snapshots.

> **NOTE:** Before using VPN tunneling bandwidth management policies, you must specify the maximum bandwidth and VPN maximum bandwidth values for the appliance.

User is Mapped to Multiple Roles

The following decision process is made when a user is mapped to multiple roles:

- Calculate the Bandwidth management policies based on the privilege level defined.

  - The current used bandwidth percentage is calculated and compared with the privilege levels of the Bandwidth management policy of the mapped roles.

  - All bandwidth management polices with the privilege levels that disallow the user to set up VPN tunnels are discarded.

- Compare the matched bandwidth management policies and choose the one with the highest guaranteed minimum bandwidth. If more than one policy with the highest guaranteed minimum bandwidth exists, the policy with the highest maximum bandwidth wins.

For example, a user is mapped to 3 roles and the bandwidth management policy for each role is as follows:

	Role 1	Role 2	Role 3
Minimum guaranteed bandwidth	100 mbps	200 mbps	100 mbps
Maximum guaranteed bandwidth	500 mbps	400 mbps	400 mbps
Privilege level	Medium	High	Maximum

If the current total used bandwidth is at 80%:

- Since role 1's privilege is not enough to allow this user to set up a VPN tunnel, role 1's bandwidth management policy is ignored.

- Role 2's policy has higher minimum guaranteed bandwidth than role 3 so role 2 wins. The user receives a 200 mbps minimum guaranteed bandwidth and 400 mpbs maximum guaranteed bandwidth.

However, if the current total used bandwidth is 92%, only role 3's privilege allows the user to set up NC tunnel, so role 3's bandwidth management policy is used. Thus the user has a 100 mbps minimum guaranteed bandwidth and 400 mbps maximum guaranteed bandwidth.

Limitations

Bandwidth management may not operate correctly under the following conditions:

- More than one IVS uses the same system internal port.

- More than one IVS has the same VLAN IP.

- Overlapping VPN tunneling IPs across IVSs.

Related Documentation

- Task Summary: Configuring VPN Tunneling

- VPN Tunneling Execution

Writing a VPN Tunneling Bandwidth Management Resource Policy

To write a VPN tunneling bandwidth management resource policy:

1. In the admin console, choose **Users > Resource Policies > VPN Tunneling > Bandwidth Management**.

2. On the Bandwidth Management page, click **New Policy**.

3. On the New Policy page, enter:

   - A name to label this policy.

   - A description of the policy (optional).

4. In the Bandwidth Management Settings section, specify:

   - **Admission Privilege Level**—Select the percentage of the maximum bandwidth that allows users to start a VPN session. Only when the bandwidth is below this percentage can users log in.

   - **Guaranteed Minimum Bandwidth**—Specify the user's minimum bandwidth once they start a VPN session.

   - **Maximum Bandwidth**—Specify the user's maximum bandwidth once they start a VPN session.

     > *i* NOTE: The maximum bandwidth must be less than or equal to the maximum rated value for the appliance.

5. In the Roles section, specify:

   - **Policy applies to ALL roles**—To apply this policy to all users.

   - **Policy applies to SELECTED roles**—To apply this policy only to users who are mapped to roles in the Selected roles list. Make sure to add roles to this list from the Available roles list.

   - **Policy applies to all roles OTHER THAN those selected below**—To apply this policy to all users except for those who map to the roles in the Selected roles list. Make sure to add roles to this list from the Available roles list.

6. Click **Save Changes**.

Related Documentation

- VPN Tunneling Execution

Configuring the VPN Tunnel Server

You use the System > Network > VPN tunneling page to configure VPN tunnel server options. This topic includes the following information:

- Specifying IP Filters
- Specifying the VPN Tunneling Server Base IP Address

Specifying IP Filters

The VPN Tunneling Server uses the filter list to assign IP addresses to clients requesting a VPN client session. A filter is an IP address/netmask combination. For example: 10.11.0.0/255.255.0.0 or 10.11.0.0/16.

To add an IP address to the VPN tunneling filter list:

1. In the admin console, choose **System > Network > VPN tunneling**.

2. Specify an IP address/netmask combination and then click **Add**.

Specifying the VPN Tunneling Server Base IP Address

i NOTE: Only change the VPN tunneling server base IP address when instructed to do so by the Pulse Secure Support team.

To change the VPN tunneling server base IP address:

1. In the admin console, choose **System > Network > VPN tunneling**.

2. In the **VPN Tunnel Server IP Address** text box, specify the base IP address used by the VPN tunneling server to assign IP addresses to the tunnel interfaces created for VPN Tunneling sessions. If your service is deployed in a cluster, the base IP address you specify will be common to all cluster nodes. Be sure to configure a base IP address that does not encroach on the IP address pool for VPN Tunneling Connection Profiles or for or the IP addresses for the external or internal interface. Take DHCP servers into consideration.

   For IPv6 addresses, no additional configuration is required. The base IPv6 address used by the VPN tunneling server will be generated by prepending the network prefix fd00:: to IPv4 address configured for this.

3. Save the configuration.

Related Documentation

- About VPN Tunneling

- VPN Tunneling Connection Profiles with Support for Multiple DNS Settings

VPN Tunneling Installer Overview

To download the VPN tunneling application as a Windows executable file, go to Maintenance > System > Installers.

VPN tunneling Installation Process Dependencies

During installation, VPN tunneling interacts with a number of system components, performing checks and validations along the way. The following list provides the order of execution during installation, which may be helpful if you need to debug a VPN tunneling installation process.

1. Start Pre-Installation Process:

   1. Parse command line arguments.

   2. Set appropriate variables via command line.

3. Process commands, as necessary.

4. If the command line entry responds with help or version information, the VPN tunneling installation program quits, following the command line processing. Typically occurs when you run the VPN tunneling installer as a standalone installer.

5. Validate System:

   1. Check OS. If VPN Tunneling does not support this OS version, display error and abort validation process.

   2. Check Administrator privileges.

   3. 3rd-party GINA component – if GINA is to be registered, check whether there is any existing registered GINA component. If yes, abort installation.

   4. If there is an existing VPN tunneling installation, trigger the uninstall in upgrade mode of the existing VPN tunneling.

   5. Wait until the existing VPN tunneling uninstallation process completes (in upgrade mode).

   6. If the uninstallation process times-out, display error message and abort the VPN tunneling installation, otherwise, continue the VPN tunneling installation.

   7. Write logging registry keys for VPN tunneling components.

   8. Start VPN tunneling installation.

   9. Shared component installation:

      1. Check sharedDll registry value of the shared components to see if this is the first instance of shared component installation.

      2. Check if Neo_CleanInst flag is set.

      3. If steps a or b are true, ensure the sharedDll registry value is clean.

      4. Stop service if still running.

      5. Check installation and driver

         - If driver is installed and it is a clean installation, uninstall the driver.

         - If driver is installed and it is not a clean installation, compare driver versions.

         - If it is an upgrade, set the driver install flag, otherwise, do not install the driver (keep the current higher version driver).

   6. VPN tunneling component installation:

      1. If the driver install flag is set or if it is a clean install, install the driver.

      2. Call the shared component installation macro for the VPN tunneling service and GINA component. This macro performs a version comparison, ensures a proper upgrade, and increments the sharedDll registry key value.

      3. Copy other VPN tunneling binary files.

      4. Call the NCCopyFile macro for the files that might be locked by msGINA. This macro takes care of renaming old files and mark them delete on reboot.

      5. Register GINA if GINA flag is set.

6. Save locale and GINA settings in user's config.ini file.

7. Start the NCService.

8. Create program shortcut.

9. Create Uninstall registry keys.

10. Start VPN tunneling user interface.

11. End VPN tunneling installation process.

12. Start Post-Installation Process:

    1. Print product version and append the install log to admin log file

    2. Reboot, if the reboot flag was set.

VPN Tunneling Uninstallation Process Dependencies

During uninstallation, VPN tunneling interacts with a number of system components, performing checks and validations along the way. The following list provides the order of execution during uninstallation, which may be helpful if you need to debug a VPN tunneling uninstallation process.

1. Start Pre-Uninstall Process:

    1. Parse command line inputs, including:

        - Locale

        - Clean uninstall flag

        - Upgrade flag

2. Start uninstall operation.

3. Check Administrator privileges.

4. Unregister GINA if already registered.

5. If uninstalling in upgrade mode, stop the VPN tunneling service.

6. If the uninstallation is not in upgrade mode, check the current sharedDll registry key value. If the value is 1, this is the only instance using the shared components, so:

    1. Uninstall the driver.

    2. Delete the driver file.

    3. Stop and unregister the VPN tunneling service.

    4. Call the shared components macro to uninstall shared components. This macro decrements the SharedDLL registry key value and removes the source file.

        ℹ️ NOTE: If the uninstall process is in upgrade mode, this step is not executed because the uninstall is triggered from a VPN tunneling installation process and the shared component macro in the installation process will handle the shared component upgrade operations.

    5. Delete other VPN tunneling files, including:

- dsNcAdmin.dll

- dsNcDiag.dll

- versioninfo.ini

6. Call the NCDeleteFile macro to delete the files that may be locked by msGINA.

7. Delete VPN tunneling registry keys.

8. Remove VPN tunneling program file directories.

9. End the uninstall process.

10.   Print the product version and append the VPN tunneling installation log to the Admin log.

11.   Reboot, if the reboot flag was set.

Related Documentation

- Downloading Client Installer Files

Network Connect Launcher (NC Launcher) Overview

The Network Connect Launcher is a client-side command-line utility that maintains a very small footprint. You can bundle NC Launcher with other applications that need an operational Network Connect client. Bundling the NC Launcher with other applications allows you to be confident that the end user can access these applications without difficulty. The NC Launcher allows you to initiate an NC session with a script, batch file, or a call from an application, without using a graphical user interface.

> NOTE: The NC Launcher does not support the role mapping option Users must select from assigned Roles when more than one role is assigned to a user. If you use NC Launcher and more than one role can be assigned to a user, you must choose to either Merge settings for all assigned roles or you must choose the option that forces the User to select the sets of merged roles assigned by each rule.

To use the NC Launcher:

1. Write a script, batch file or application to call the NC Launcher.

2. Include a call to the NC Launcher executable.

The NC Launcher command syntax is:

```
nclauncher.exe [-version|-help|-stop|-signout] -u user -p password -url url -r
realm -ir [true | false] -t seconds -c certificate-name -d DSID
```

Argument	Action
-version	Displays the NC Launcher version information, then exits.
-help	Displays available arguments information.
-u *user*	Specifies the username.
-p *password*	Specifies the password for authentication.
-url *url*	Specifies the system URL information.

-r *realm*	Specifies the realm to which the system submits the user's credentials.
-ir [true | false]	Specifies whether to avoid client-side certificate revocation list on the client.
-t *seconds*	How long to wait for the NC tunnel to establish (in seconds).
-c *certificate-name*	Specifies the certificate to use for user authentication instead of a username and password. For *certificate-name* use the string specified in the Issued To field of the certificate. To use certificate authentication with NC Launcher, you must first configure the system to allow the user to sign in via user certificate authentication. You must also configure a trusted client CA on the system and install the corresponding client-side certificate in the Web browsers of your end users before running NC Launcher. When using the -c argument, also specify the -url and -r arguments. If the certificate is invalid, NC Launcher displays an error message on the command line and logs a message in the nclauncher.log file.
-d *DSID*	Passes a cookie to NC Launcher from another authentication mechanism when NC Launcher starts
-signout	Terminates NC tunnel and signs out current user.
-stop	Terminates NC tunnel.

> ℹ️ NOTE: nclauncher does not support secondary authentication.

For example, you might distribute a simple login application to your end users. The application might capture the end user's username and password, the system resource they are trying to reach, and the realm to which they are assigned.

In this example, you need to write a script or add logic to your application to capture the credentials the end user enters, and pass the credentials as the arguments to the nclauncher.exe, as follows:

```
nclauncher.exe -u JDoe -p my$Pass84 -url https://int-company.portal.com/usr -r User
```

Table 98 lists the possible return codes nclauncher returns when it exits.

Table98: nclauncher Return Codes

Code	Description
-1	(-Stop/-Signout) Network Connect is not running. System error occurred.
0	Network Connect started.
1	Invalid arguments.
2	Network Connect is unable to connect to the Secure Gateway.
3	Network Connect is unable to authenticate with the server.
4	The specified role is invalid or does not exist
5	Network Connect cannot run because a required pre-authentication application could not be started.
6	Network Connect installation failed.

8	Network Connect was unable to perform a required software upgrade.
10	The server to which you are trying to connection does not support this feature.
12	Network Connect failed to authenticate the client certificate.
15	Network Connect failed to authenticate the client certificate because the certificate is invalid.
16	Network Connect failed to authenticate the client certificate because the certificate has expired.
17	Network Connect failed to authenticate the client certificate because the certificate has been revoked.
18	Host Checker policy failed.

If Network Connect is launched through a browser, certificate verification is taken care of by the browser. Similarly, if Network Connect is launched through the standalone application on Windows, certificate verification is handled by the application. However, if Network Connect is launched through nclauncher on Windows, nclauncher handles the expired or revoked client certificates.

Related Documentation

- About VPN Tunneling

Launching Network Connect on Other Platforms

To launch a Network Connect command-line session on a non-Windows platform:

```
# cd ~/.juniper_networks/network_connect/
# ./ncsvc -h host -u user -p passwd -f cert_file [-r realm] [-L log_level] [-g]
[-U sign_in_url] [-y proxy] [-z proxy_port] [-s proxy_user] [-a proxy_password]
[-d proxy_domain] [-v] [-K]
```

To launch a standalone Network Connect graphical user interface session on a non-Windows platform:

```
# cd /user/local/nc/
#<javapath> -cp NC.jar NC -h host -u user -p passwd -f cert_file [-r realm] [-l
log_level] [-L log_level] [-U sign_in_url] [-y proxy] [-z proxy_port] [-s
proxy_user] [-a proxy_password] [-d proxy_domain] [-n start_script] [-t
end_script]
```

ncsvc parameters are as follows:

Parameter	Description
Sign-in Options	
-h *host*	Connect Secure hostname or IP address
-u *user*	Username
-p *passwd*	Password
-r *realm*	Connect Secure sign-in realm
-P *port*	Service port number
-f *cert_file*	Connect Secure certificate
-m *md5hash*	Connect Secure certificate MD5 hash
-U url	Connect Secure realm sign-in URL

Proxy options	
-y *proxy*	Proxy server hostname or IP address
-z *proxy_port*	Proxy server port number
-s *proxy_user*	Proxy server username
-a *proxy_pass*	Proxy server password
-d *proxy_domain*	Proxy server domain name
Logging options	
-L *log-level*	Log message level. Options are: 0—Log critical messages only 1—Log critical and error message 2—Log critical, error and warning message 3—Log critical, error, warning and info message (default) 4—Log verbose messages 5—Log all messages
Miscellaneous options	
-v	Print version number
-g	Zip and upload log file to host
-K	Kill all running ncsvc services
Script options	
-n *start_script*	Network Connect start script
-t *end_script*	Network Connect disconnect script

Related Documentation

- Network Connect Launcher (NC Launcher) Overview

Troubleshooting Network Connect Errors

Your end users may be unable to resolve some of the errors they might encounter, without your intervention. The following topics may correspond to errors listed in the end-user help.

nc.windows.app.23792

If your end user encounters this error (or nc.windows.app.1023), they have been unable to connect to the system.

Check items on both the client machine and on Connect Secure.

On the client

- The error may indicate a faulty Java installation on the client. Have the client uninstall and reinstall the JRE.

- One or more of the third-party clients on your end user's computer may be faulty or interrupting the connection. You will need to check or have your user check clients such as VPN clients, antivirus, personal firewalls, spyware, and other types of endpoint security clients.

- The \Documents and Settings folder on the end user's computer might be encrypted. Network Connect cannot install itself to an encrypted directory.

On Connect Secure:

- Ensure that the DNS name lookup for the system does not resolve to its public/private IP address.

- If the NC profile is configured to use an external DHCP server, ensure that the DHCP server is responding to the system.

- Select System > Network > Overview and make sure that you have configured DNS for the system.

Version Conflict on Downgrade

When you downgrade to a prior release, your end users might receive a version conflict error when Network Connect initiates. The problem may occur because the client contains a newer version of certain files which the older version cannot use properly.

To resolve this problem, delete the following files:

- <user_home>/.juniper_networks/ncLinuxApp.jar

- <user_home>/.juniper_networks/network_connect/*.*

If this solution does not solve the version conflict problem, contact your system administrator.

Error When Connecting to a FIPS Appliance

The error **Could not connect to Secure Gateway because the certificate is invalid or not trusted by the client system. Click OK to exit Network Connect and sign in to the Secure Gateway again. If problem persists, please contact your administrator.** occurs when the Network Connect client has received a server certificate from the FIPS appliance and:

- Connect Secure is configured with self-signed certificate then either the server certificate is not valid or is not trusted by the client system.

- Connect Secure is configured with complete certificate chain then either the server certificate is not valid or the root CA of the certificate chain is not trusted by the client.

- Connect Secure is configured with complete certificate chain, but it sent only a leaf certificate (this may happen if the system is missing some of its sub-CAs) then either the server certificate is not valid or the complete certificate chain is not available on the client.

Related Documentation

- Task Summary: Configuring VPN Tunneling

- VPN Tunneling Execution

- Configuring Enterprise Onboarding

CHAPTER 30 Enterprise Onboarding

- Configuring Enterprise Onboarding

- Managing Onboarded Devices

Configuring Enterprise Onboarding

Enterprise onboarding allows users to securely access enterprise network resources with almost any device. Wi-Fi, VPN and certificate profiles can be defined for enterprise resources and downloaded to a device during onboarding, depending on the device type.

The profiles can be defined on a single Connect Secure or Policy Secure server dedicated to onboarding or they can be defined on each server. Alternatively, the profiles can be defined on a third-party MDM server, in which case users will see a link and instructions on the onboarding page to continue onboarding using the external MDM server.

Onboarding is initiated from the browser. The supported profiles depend on the device type and whether the Pulse Secure client is installed (see Table 99).

Table99: Device Onboarding Profile Support

Device	Supported profiles
Android 4.0 or later	Supports all profiles, but the Pulse Secure client must be installed during onboarding.
iOS 6.0 or later	Supports all profiles (Safari browser).
Windows 7.0, 8.0, and 8.1	Supports Wi-Fi and certificate profiles (IE, Firefox, or Chrome browser). The Pulse Secure client onboarding application must be installed during onboarding. NOTE: Windows 8 RT and Windows 8 Phone are not supported.
MAC OS X	Supports Wi-Fi and certificate profiles (Safari browser).

Enterprise onboarding is enabled in the user role, and each profile can be applied to all user roles or specific roles. The SCEP server and CSR templates allow certificates to be generated dynamically for device and server authentication.

- Enabling Enterprise Onboarding at the Role Level

- Defining the SCEP Server

- Defining CSR Templates

- Defining VPN Profiles

- Defining Wi-Fi Profiles

- Defining Certificate Profiles

- Onboarding Devices

- Workflow for Onboarding Android Devices

Enabling Enterprise Onboarding at the Role Level

To enable enterprise onboarding for a user role:

1. In the admin console, choose **Users > User Roles >** *RoleName* **> General > Overview**.

2. In the Enterprise Device Onboarding section, select the **Enterprise Onboarding** check box.

3. Click **Save Changes**.

4. Click the **Enterprise Onboarding** tab or click **Options** next to the **Enterprise Onboarding** check box to specify the following:

   - **Auto launch** – Displays the onboarding page when the user logs in to Connect Secure or Policy Secure if enterprise onboarding is enabled for the user's role (the default). If this option is disabled, an onboarding link is displayed on the home page.

   - **Use third party MDM for onboarding** – Displays a link on the onboarding or home page where the user can download profiles from an MDM server. Enter the URL for the MDM page in the text box.

   - **Install Pulse Secure Client:** enabling this option will automatically install Pulse client during onboarding from Windows OS

5. Click **Save Changes**.

Defining the SCEP Server

The Simple Certificate Enrollment Protocol (SCEP) server configuration and CSR templates allows each client device to dynamically obtain certificates for authentication.

To define the SCEP server:

1. In the admin console, choose **Users > Enterprise Onboarding**.

2. Specify the following information:

Setting	Description
SCEP Server URL	Enter the URL for a SCEP server. The following SCEP servers are supported: • Microsoft AD 2008 • Symantec mPKI
Challenge	Specify the password required by the SCEP server.
Retries	Specify the number of attempts to access the server when the first attempt fails.
Retry Delay	Specify the number of seconds between retry attempts.
Upload Encryption Certificate	Click **Browse** to upload the certificate used to encrypt SCEP requests. To upload the certificate automatically, select the **Test Enrollment** check box, select a CSR template, and click **Test Configuration**. To create a CSR template, see "Defining CSR Templates".

3. Click **Save Changes**.

Defining CSR Templates

If the SCEP server is configured, the Certificate Signing Request (CSR) templates can be used in the VPN, Wi-Fi, and certificate profiles to allow each onboarded device to dynamically obtain certificates for authentication on all mobile devices. Up to 10 templates can be defined.

> **NOTE:** All LDAP attributes (such as **<ldap.userAttrName>**) and variables (such as **<user>**) can be used in the Subject DN, Email, and Subject Alternative Name Value fields. However, if you enter an LDAP variable with a string vector data type in the Subject Alternative Name Value field, only the first value in the string will be used.

To define CSR templates:

1. In the admin console, choose **Users > Enterprise Onboarding > CSR Templates**.

2. To add a template, click **New CSR Template** or select an existing template that you want to change, duplicate, or delete. Clicking **Duplicate** creates a copy of the selected template with **Copy of** before the template name.

3. Specify the following information:

Setting	Description
Name	Specify the template name displayed in the list of CSR templates.
Subject DN	Specify the subject distinguished name. For example: **CN=<USERNAME>,OU=Engineering=Pulse** All LDAP attributes (such as **<ldap.userAttrName>**) and variables (such as **<user>**) can be used.
Email	(Optional) Specify an email address with the <USER> variable, such as **<USER>@Pulsesecure.net**.
Subject Alternative Name Type	Select an alternative name type if the CA requires an alternative subject name. The types include RFC-822 Name (an e-mail address), DNS domain name, URI, and IP address.
Subject Alternative Name Value	Specify one or more values for the selected alternative name type. Multiple values must be separated by a comma or space. NOTE: If an LDAP variable is specified that has a string vector data type, only the first value in the string will be used.
Key Size	Select the key size used by the SCEP server.

4. Click **Save Changes**.

> **NOTES:**
>
> 1. The number of keys available in the system can be viewed at Users->Enterprise OnBoarding->CSR Templates
>
> 2. The keys are generated only if
>
>   A) The onboardign license is installed
>
>   B) A CSR template is configured for that key size

3. The max number of keys of each type is minimum of 10K keys and the number of onboard user license installed.

4. Key generation is CPU intensive and time consuming. If bulk users are going to onboard it is recommended to make sure that the number of available keys >= the numbers of users that needs to onboard.

Defining VPN Profiles

VPN profiles provide Android and iOS devices with secure access to enterprise networks. One or more VPN profiles can be assigned to specific user roles or to all roles. Up to 10 profiles can be defined.

> 🛈 NOTE: All LDAP attributes (such as **<ldap.userAttrName>**) and variables (such as **<user>**) can be used in the Username, Realm, and Role fields.

To define VPN profiles:

1. In the admin console, choose **Users > Enterprise Onboarding > VPN Profiles**.

2. To add a profile, click **New Profile** or select an existing profile that you want to change, duplicate, or delete. Clicking **Duplicate** creates a copy of the selected profile with **Copy of** before the profile name.

3. Specify the following profile information:

Setting	Description
Name	Specify the name to be displayed in the list of VPN profiles.
Description	(Optional) Enter a description of the VPN profile.
Apply to Client Types	Select the device types the profile applies to (Android and iOS only).
Server URL	Specify the URL of the VPN server (must be a Connect Secure or Policy Secure device).
Realm	Specify the realm name. The realm is required only if the sign-in URL has the **User picks from a list of authentication realms** option enabled.
Role	Specify the user role. The user role is required if the role mapping rules for the user realm specify multiple roles and the **User must select from among assigned roles** option is enabled.
Username	Specify the <USER> variable for the user name.

Authentication Method	Select **Password** or **Certificate** for the user authentication method. For certificate authentication, specify the following:
	
	• **Use CSR Template**—Select the CSR template used to obtain the certificate. To create a CSR template, see "Defining CSR Templates".
	
	• **Enable VPN On Demand**—Select this option to allow iOS devices to establish the VPN when a specific host or domain is accessed. To specify the first host or domain:
	
	  • **Match Domain or Host**—Enter a hostname or a partial domain name. For example, if you enter **example.com**, a match occurs when the user accesses any domain that ends with **example.com**, such as **www.test-example.com**.
	
	  • **On Demand Action**—When a match occurs on the specified host or domain, select whether a VPN is always established, never established, or only if the DNS look-up fails (**Establish If Needed**). Selecting **Never Establish** does not prevent an existing VPN from being used.
	
	To add another domain, click the **+** button. To remove a domain, select the check box next to the domain and click the **-** button. Up to 10 domains can be defined.

Roles	Select one of the following options: • **Policy applies to ALL roles**—To apply this profile to all users. • **Policy applies to SELECTED roles**—To apply this profile only to users who are mapped to roles in the Selected roles list. Make sure to add roles to this list from the Available roles list. • **Policy applies to all roles OTHER THAN those selected below**—To apply this profile to all users except for those who map to the roles in the Selected roles list. Make sure to add roles to this list from the Available roles list.

4. Click **Save Changes**.

Defining Wi-Fi Profiles

Wi-Fi profiles provide Android, iOS, MAC OS X, and Windows devices with secure access to wireless networks. One or more Wi-Fi profiles can be assigned to specific user roles or to all roles. Up to 10 profiles can be defined.

> *i* NOTE: All LDAP attributes (such as **<ldap.userAttrName>**) and variables (such as **<user>**) can be used in the Username and Password fields for the WPA Enterprise and WPA2 Enterprise security types.

To define Wi-Fi profiles:

1. In the admin console, choose **Users > Enterprise Onboarding > WiFi Profiles**.

2. To add a profile, click **New Profile** or select an existing profile that you want to change, duplicate, or delete. Clicking **Duplicate** creates a copy of the selected profile with **Copy of** before the profile name.

3. Specify the following profile information:

Setting	Description
Name	Specify the name to be displayed in the list of Wi-Fi profiles.
Description	(Optional) Enter a description of the profile.

Apply to Client Types	Select the device types the profile applies to (Android, iOS, MAC OS X, and Windows).
SSID	Specify the server set ID of the wireless network.
Non-Broadcast SSID	Select the check box if the wireless network does not broadcast its identity.
Auto Connect	Select the check box to connect the client automatically when the network is detected (not supported by Android clients).
Security Type	Select the type of authentication used by the network, and specify the password or enterprise settings, as required: • **None**—No authentication required. • **WEP**—Wired Equivalent Privacy used for a non-enterprise network. Enter the network shared key in the displayed text box. • **WPA Personal** or **WPA2 Personal**—Wi-Fi Protected Access used for a non-enterprise network. You can select the encryption method (AES or TKIP) and enter the network shared key in the displayed text box (applies to Windows clients only). • **WPA Enterprise** or **WPA2 Enterprise**—Wi-Fi Protected Access used for an enterprise network. Select the Extensible Authentication Protocols (EAP) supported by the network's RADIUS authentication server. For Android devices, note the following: • Android 4.3 or later is required • For the EAP-TLS protocol, the CA certificate must be configured (along with the client certificate) on Samsung devices for authentication. • An 802.1x RADIUS server certificate must be signed by a private root CA. Authentication fails if the certificate is signed by an intermediate root CA.
EAP	For the WPA Enterprise and WPA2 Enterprise security types, select the supported EAP protocols and specify the associated authentication settings: • **None**—If none of the EAP protocols is selected (Android devices only), enter the <USER> and <PASSWORD> variables in the Username and Password fields. NOTE: iOS, MAC OS X, and Windows clients require at least one of the EAP types to be selected (PEAP, EAP-TLS, or EAP-TTLS). Selecting Multiple EAP types is not supported for Android clients.

PEAP	The PEAP protocol is supported by all clients. Specify the following:
	• **Inner Authentication Method**—Select the protocol used to authenticate the username and password (None or MSCHAPv2). The None option is valid only for Android devices.
	• **Username and Password**—Enter the <USER> and <PASSWORD> variables.
	• **Outer Identity**—Specify an alternate username to be used outside the encrypted tunnel, such as **anonymous**, to conceal the user's identity in unencrypted packets.
	• **Trusted Server Name(s)**—Specify the IP address or fully qualified domain name of one or more trusted RADIUS servers used by the network. Multiple servers must be separated by a semicolon.
	• **Trusted CA Certificate**—For Windows clients, select the Trusted Root CA of the RADIUS server certificate, even if the device certificate is signed by an intermediate CA. The Trusted Root CA must be configured in a certificate profile before it can be selected here (see "Defining Certificate Profiles"). For iOS and MAC OS X clients, if the RADIUS server certificate is signed by an intermediate CA, create a certificate profile for the intermediate CA, and then select the certificate here. The certificate profile ensures that the intermediate CA is downloaded to the client.
EAP-TLS	The EAP-TLS protocol is supported by all clients. Specify the following:
	• **Username**—Enter the <USER> variable.
	• **Use CSR Template**—Select the CSR template used to obtain the certificate. To create a CSR template, see "Defining CSR Templates".
	• **Trusted Server Name(s)**—Specify the IP address or fully qualified domain name of one or more trusted RADIUS servers used by the network. Multiple servers must be separated by a semicolon.
	• **Trusted CA Certificate**—For Windows clients, select the Trusted Root CA of the RADIUS server certificate, even if the device certificate is signed by an intermediate CA. The Trusted Root CA must be configured in a certificate profile before it can be selected here (see "Defining Certificate Profiles"). NOTE: On Windows 7 clients that have multiple certificates, users are prompted to select the certificate for 802.1x connections that use EAP-TLS. For iOS and MAC OS X clients, if the RADIUS server certificate is signed by an intermediate CA, create a certificate profile for the intermediate CA, and then select the certificate here. The certificate profile ensures that the intermediate CA is downloaded to the client.

EAP-TTLS	The TTLS protocol is supported by all clients. Specify the following: • **Inner Authentication Method**—Select the protocol used to authenticate the username and password (None, PAP, or MSCHAPv2). The None option is valid only for Android devices. • **Username and Password**—Enter the <USER> and <PASSWORD> variables. • **Outer Identity**—Specify an alternate username to be used outside the encrypted tunnel, such as **anonymous**, to conceal the user's identity in unencrypted packets. • **Trusted Server Name(s)**—Specify the IP address or fully qualified domain name of one or more trusted RADIUS servers used by the network. Multiple servers must be separated by a semicolon. • **Trusted CA Certificate**—For Windows clients, select the Trusted Root CA of the RADIUS server certificate, even if the device certificate is signed by an intermediate CA. The Trusted Root CA must be configured in a certificate profile before it can be selected here (see "Defining Certificate Profiles"). Also, if the RADIUS server certificate is signed by an intermediate CA, then the public intermediate CA must be configured in a certificate profile to ensure that the intermediate CA is downloaded to the client along with the Wi-Fi TTLS profile configuration. For iOS and MAC OS X clients, if the RADIUS server certificate is signed by an intermediate CA, create a certificate profile for the intermediate CA, and then select the certificate here. The certificate profile ensures that the intermediate CA is downloaded to the client.
Roles	Select one of the following options: • **Policy applies to ALL roles**—To apply this profile to all users. • **Policy applies to SELECTED roles**—To apply this profile only to users who are mapped to roles in the Selected roles list. Make sure to add roles to this list from the Available roles list. • **Policy applies to all roles OTHER THAN those selected below**—To apply this profile to all users except for those who map to the roles in the Selected roles list. Make sure to add roles to this list from the Available roles list.

4. Click **Save Changes**.

Defining Certificate Profiles

Certificate profiles specify the device certificates sent to each client device during onboarding. Up to 10 profiles can be defined.

> NOTE: For security reasons, certificate profiles cannot be included in the XML export or import.

To define certificate profiles:

1. In the admin console, choose **Users > Enterprise Onboarding > Certificate Profiles**.

2. To add a profile, click **New Profile** or select an existing profile that you want to change, duplicate, or delete. Clicking **Duplicate** creates a copy of the selected profile with **Copy of** before the profile name.

3. Specify the following information:

Setting	Description
Client Types	Select the device types the profile applies to (Android, iOS, MAC OS X, and Windows).
Import and Use Global Certificate	Select this option to use the Connect Secure or Policy Secure global certificate to authenticate the client device. Click **Import Certificate & Key**, click **Browse** to locate the certificate file, and then click **Import**. For more information about device certificates, see "Using Device Certificates".
Import and Use CA Certificate	Select this option to import any CA certificate (public Root CA, private Root CA, public intermediate CA, or private intermediate CA). These CA's can be used in Wi-Fi profiles and must be downloaded to the client devices. Click **Import and Use CA Certificate**, click **Browse** to locate the certificate, and then click **Import CA Certificate**.
Generate per User Certificate	Select this option to use the SCEP server and a CSR template to generate a certificate for each client. Select a CSR template from the Use Certificate Template list. To create a CSR template, see "Defining CSR Templates".
Roles	Select one of the following options: • **Policy applies to ALL roles**—To apply this profile to all users. • **Policy applies to SELECTED roles**—To apply this profile only to users who are mapped to roles in the Selected roles list. Make sure to add roles to this list from the Available roles list. • **Policy applies to all roles OTHER THAN those selected below**—To apply this profile to all users except for those who map to the roles in the Selected roles list. Make sure to add roles to this list from the Available roles list.

4. Click **Save Changes**.

Onboarding Devices

Onboarding is initiated from the browser. When a user logs in, the onboarding option is displayed if VPN, Wi-Fi or certificate profiles are defined in the user's role. MAC OS X devices and iOS devices can be onboarded without installing the Pulse Secure client. For Android devices the browser displays a link to install the Pulse Secure client. For Windows devices, the browser displays a link to install the Pulse Secure Onboarding application.

Note the following:

• If the device has a VPN connection to Connect Secure, the user is warned that the connection will be closed and reestablished through the onboarding process.

• If the user onboards the device again, which may be necessary if a certificate expires or the configuration is deleted, new profiles are downloaded to the device.

• The following message IDs in the User Access Log can be used to verify the onboarding process (they include the username and device ID):

   • AUT31186—Indicates the status of an onboarding attempt (failed or successful)

   • AUT31152—Indicates onboarding failed because the maximum device limit of 10000 has been reached

   • AUT31187—Indicates the attempt to build a configuration profile failed due to an error

- AUT31188—Indicates a configuration profile was generated successfully, and lists the names of the profiles contained in the configuration profile

The following message IDs are related to device limits:

- SYS31177—Indicates the number of devices onboarded is nearing the system limit of 10000.

- SYS31178—Indicates the number of devices onboarded has exceeded the system limit of 10000 (critical error).

- SYS31193—Information message generated by a background process that attempts to delete device records when 95% of system limit (10000) is reached. It displays the number of device records deleted, the current number of onboarded devices, and the system limit.

Workflow for Onboarding Android Devices

To onboard an Android device:

1. Enter the onboarding URL in the browser.

Figure122: Onboarding Start Page
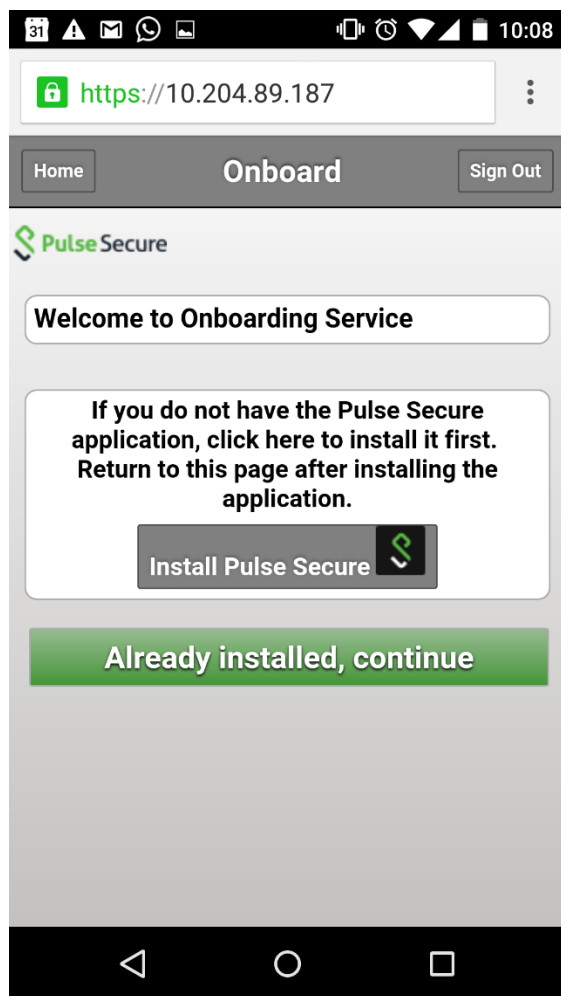
2. The first time Pulse Secure client is launched, the EULA is displayed.

Figure123: End User License Agreement
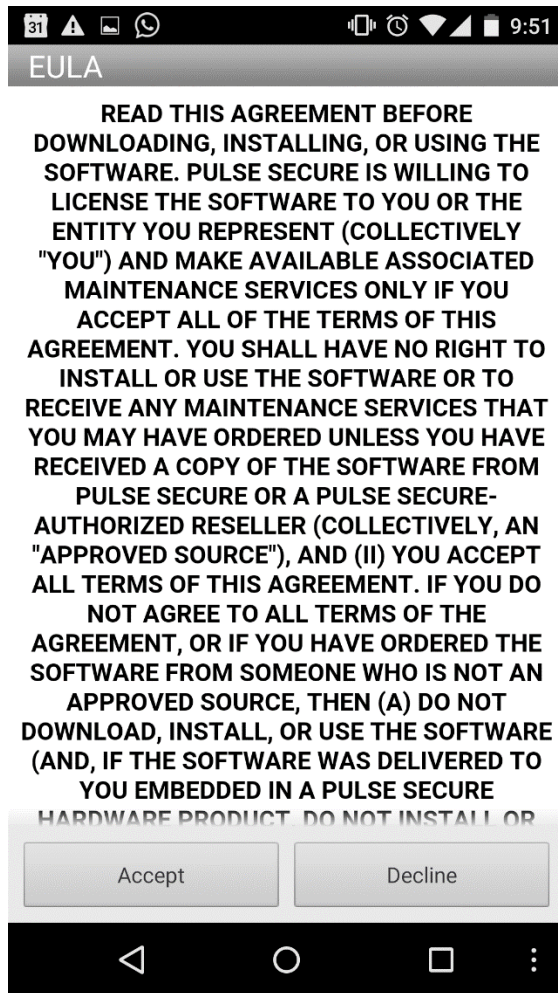
3. On the Pulse Secure client onboarding Wizard page, read the instructions carefully and tap **Next**.

Figure124: Onboarding Wizard Start Page
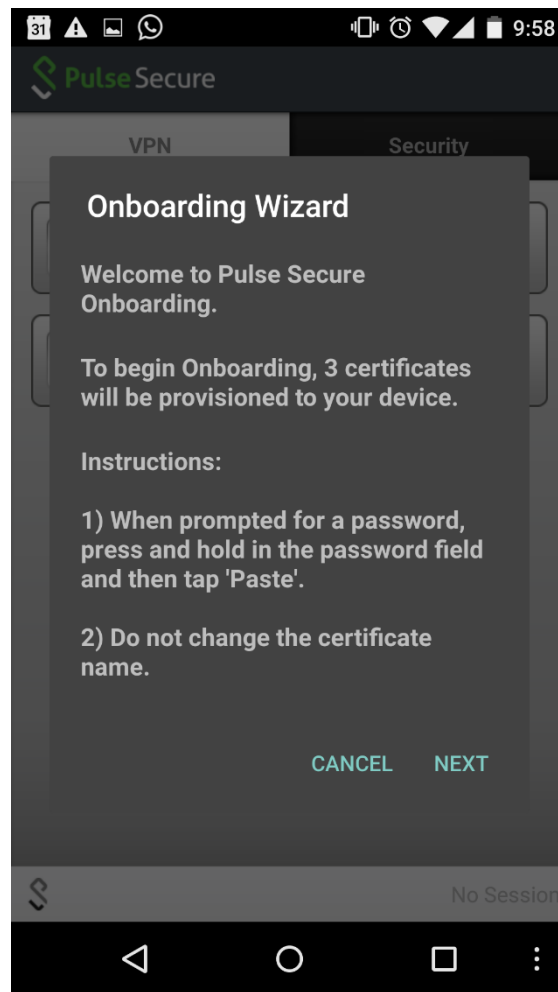
4. On the CA certificate provisioning page, tap **OK**.

Figure125: Certificate Provisioning Page
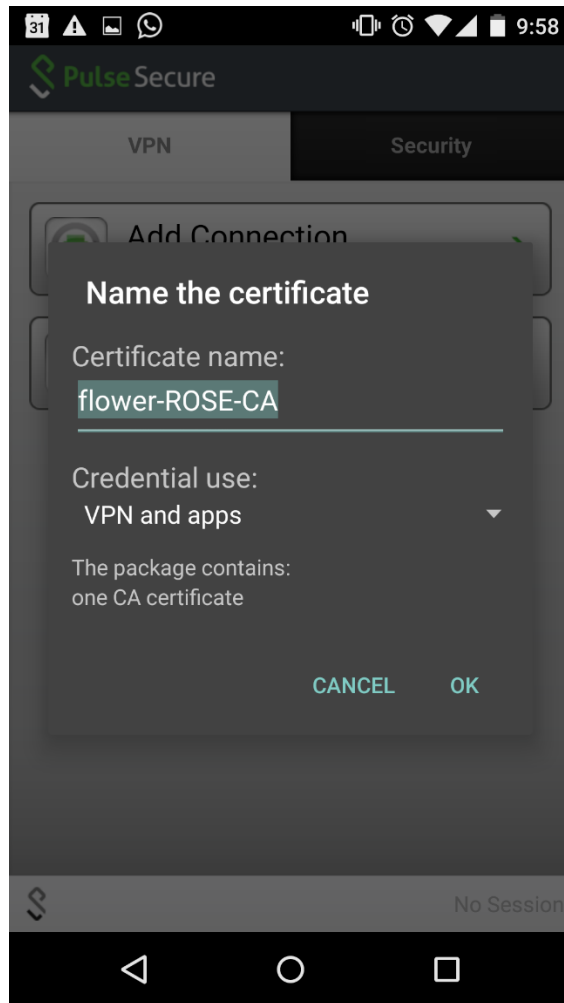
5. Paste the password from the clipboard to extract the certificates, and tap **OK**.

Figure126: Certificate Password Page
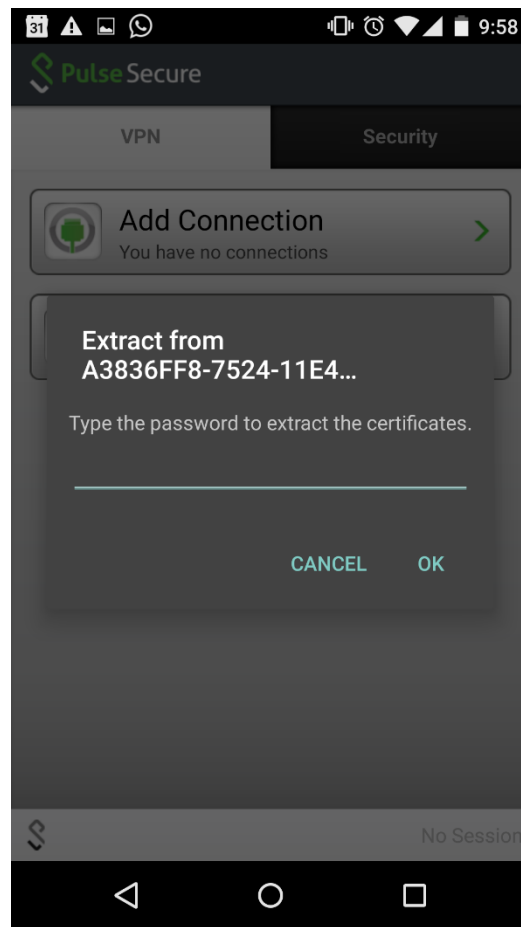
6. Tap **OK** to confirm the certificate name.

753

Figure127: Certificate Name Page
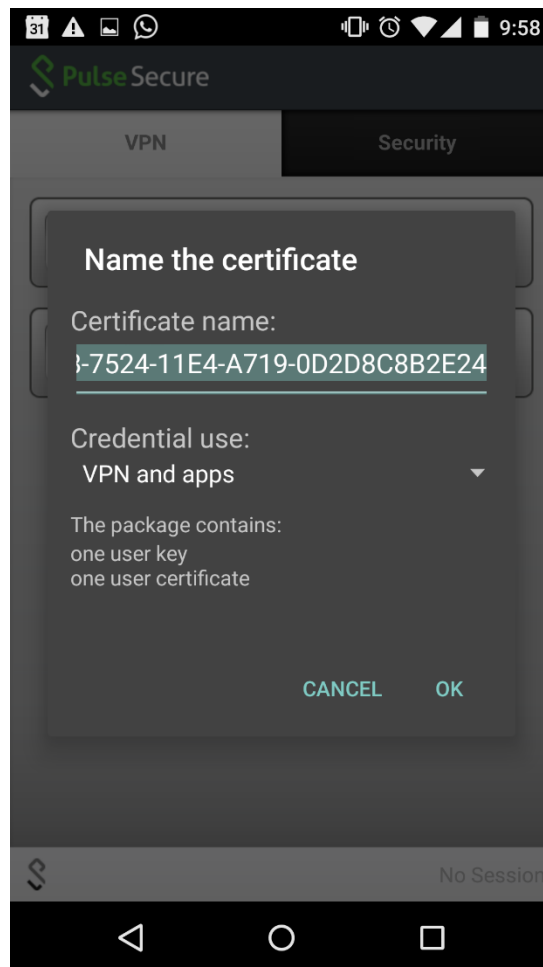
7. View the number of certificates and profiles provisioned on the client, and tap **Finish**.

Figure128: Onboarding Wizard Summary Page
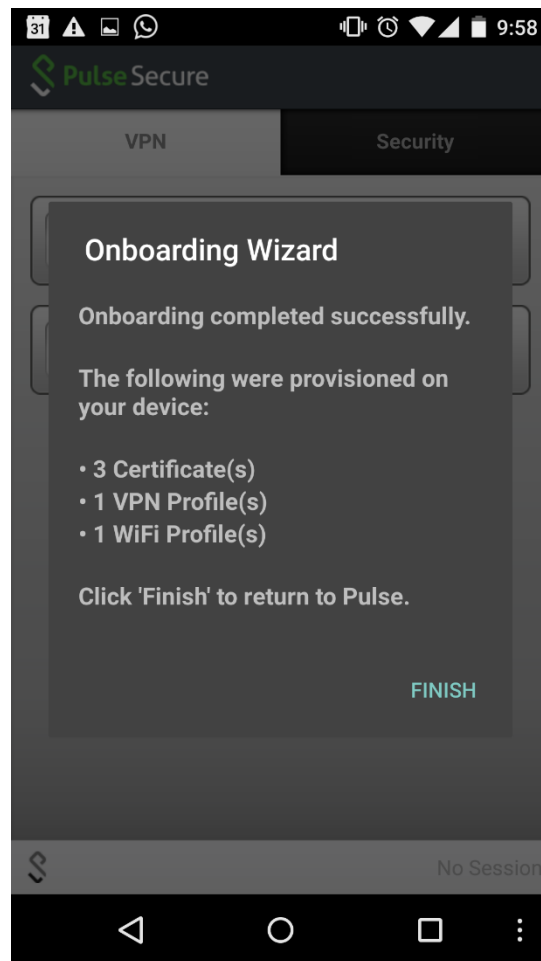
8. If onboarding fails due to an error, tap **Retry**. Users should contact their administrator if onboarding fails after three attempts.

Figure129: Error Page
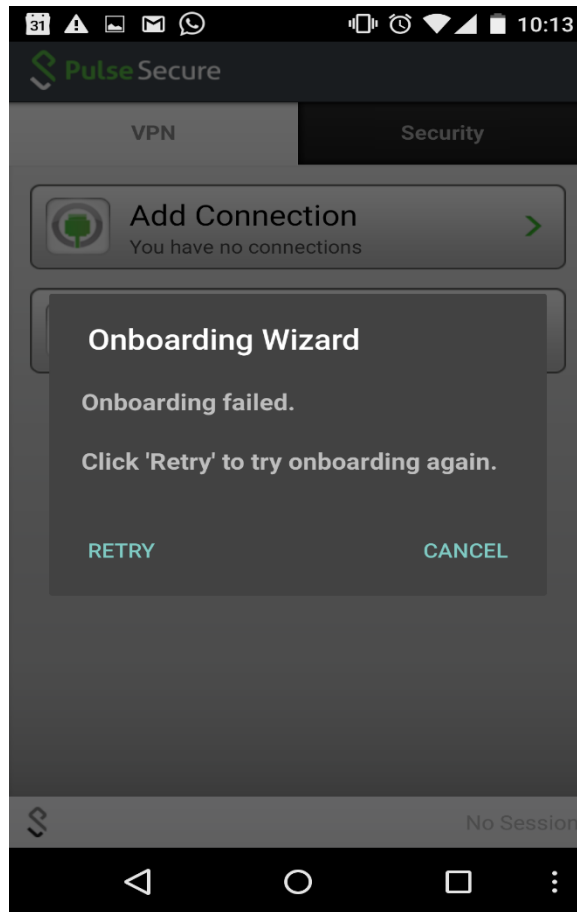
9. If onboarding is successful, tap the **VPN** tab to view the provisioned VPN connections.

Figure130: VPN Connections Page
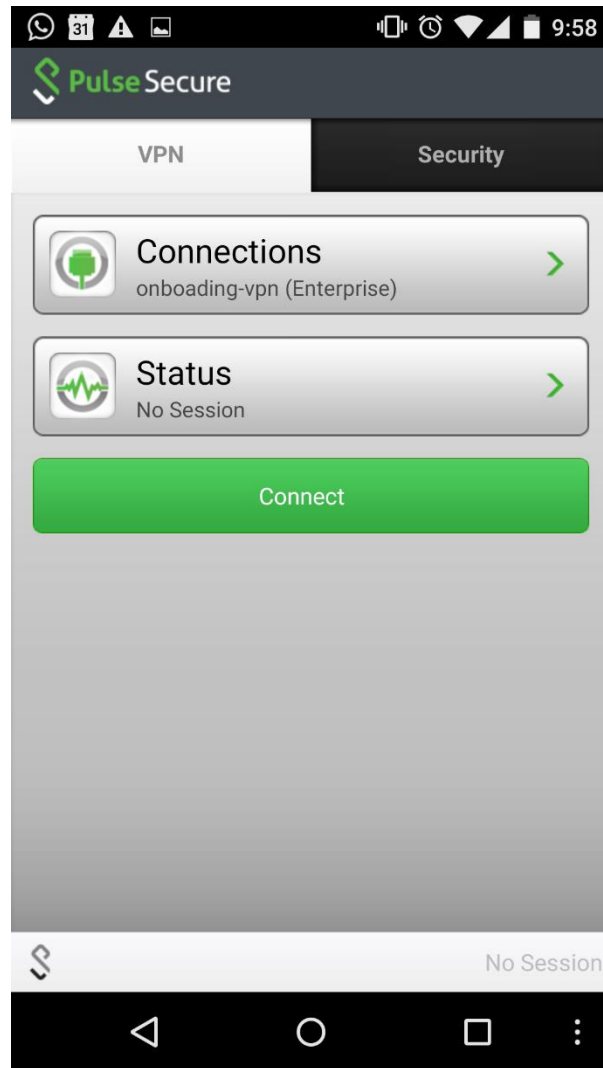
10.   Tap the Wi-Fi icon to view the provisioned Wi-Fi networks. To enable a Wi-Fi connection, select the network and tap the Connect icon.

Figure131: Wi-Fi Connections Page
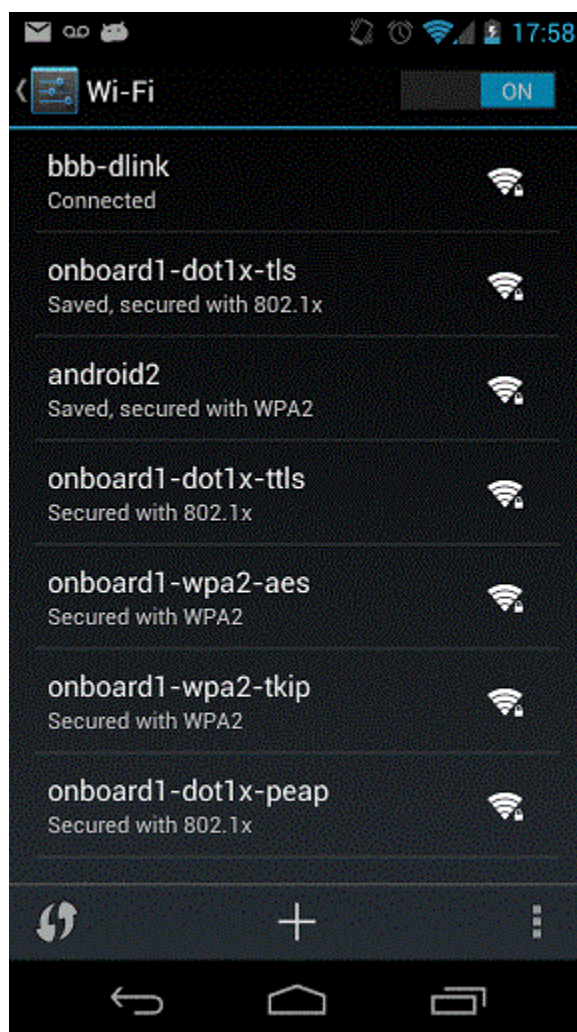
Related Documentation

- Managing Onboarded Devices

Managing Onboarded Devices

The Device Management page lists the following types of devices:

- **Onboarded devices**— Devices that have Enterprise Onboarding enabled in the user's role and have been onboarded during device registration. After a device is onboarded, it is displayed on the Device Management page until it is deleted.

The username, user roles, operating system, and registration date are shown for each device, along with the onboarded, and access status. Devices that become inactive or invalid must be deleted manually.

To view the Device Management page:

1. Select **System > Status > Devices**.

2. Use the controls described in Table 100 to view and manage the devices.

Table100: Device Management Page

Buttons	Administrative Actions
Update	• To refresh the page, click **Update**. • To view a specific user, enter the username in the Show Users Named box and click **Update**. If you do not know the exact username, use an asterisk (*) as a wildcard character. • To change the number of displayed devices, enter a number in the Show N devices box and click **Update**. TIP: To change the sort, click a column header.
Delete	To delete one or more devices, select the check box next to the appropriate devices and click **Delete**. If an onboarded device is deleted in error, the user must re-onboard the device. ActiveSync-only devices that are deleted in error are added automatically by the next ActiveSync.
Delete All	To delete the all devices, click **Delete All**.

Related Documentation

- Defining Client Types

- Displaying System Status

- Configuring Enterprise Onboarding

PART 5 System Management

CHAPTER 31 Licensing

About License Management

Pulse Connect Secure software and Pulse Policy Secure software include a license management system that lets you configure the Connect Secure device, including services modules on a MAG Series Pulse Gateway configured as Connect Secure, as a license server to allow administrators to view all configured systems and move those licenses as needed. Other devices on the network lease licenses from the central license server.

Alternatively, you can install and manage licenses directly on each device and eliminate the license server entirely. Your company needs and requirements dictate which configuration is best for you.

End User License Agreement (EULA) acceptance is mandatory and you are entitled to use the features of the software that you have licensed within the limits of your Proof of Entitlement. Contact your sales representative for more information.

Related Documentation

- About License Servers

- About License Clients

- About License Allocations

About License Servers

If you choose to use license servers, the license server software can be run on any SA*x*000 and SA*x*500 devices, the MAG and PSA NGHW models, and Pulse Secure gateway devices running the Connect Secure personality. Once you configure a device to be a license server, that appliance ceases to be anything except a license server; it will no longer accept end-user client connections. You can configure more than one license server, but each client can be associated with only one license server. A device cannot be both a license server and a license client at the same time.

Note the following about license servers:

- Only administrators can log in to a license server.

- Virtual appliances can not be configured as license servers and you cannot install licenses on virtual appliances.

- A license server can not lease licenses from another license server.

- The license server manages and leases licenses associated with a user count, such as basic concurrent user licenses, and RDP (remote desktop) licenses. It also leases concurrent meeting user licenses and Advanced Mobile Licenses – Onboarding. Figure132 displays the list of configured clients and the licenses leased by them.

Figure132: List of Configured Clients and licenses leased by them

License clients

This section displays the list of configured clients and the licenses leased by them.

New Client... | Delete...

		Client Name	Cluster	Leased by	Last Renewal Time	Concurrent Users	Premier Java Remote Desktop Applet	Concurrent Meeting Users	Advanced Mobile License - Onboarding
1.	☐	test-2		test-2	07:41:26 AM, Jul 13, 2015	50	25	0	○
2.	☐	test1		test1	11:26:54 PM, Jul 09, 2015	100	50	0	○

New Client... | Delete...

- License servers must have either an ACCESS-LICENSE-SVR or an ACCESS-SUB-SVR-*Z*YR license (for example, ACCESS-SUB-SVR-1YR) in order to be recognized as a license server. If you are using SA Series 7.0 software you must use ACCESS-SUB-SVR-*Z*YR. If you are using Pulse Connect Secure 7.1 software you can use either ACCESS-LICENSE-SVR or ACCESS-SUB-SVR-*Z*YR.

Disabled Features

Only administrators can log in to the device configured as a license server. An error message is displayed to non-administrator users attempting to log in to the license server. All existing end-user sessions are terminated when a device is configured as a license server.

Some Connect Secure features and windows are disabled on the license server's administration console.

Related Documentation

- About License Management

- Disabled Features on a License Server

About License Clients

Clients are configured with the license server to communicate with. The client then requests the licenses (over HTTPS) that are allocated to it.

An optional LICENSE-MBR license (for example: MAG2600-LICENSE-MBR, SM360-LICENSE-MBR, PSA-LICENSE-MBR) must be installed when clients need to access capacity from non-subscription licenses or if an administrator wants to surrender a client's licenses to the license server.

If you are not using a license server, all your devices are still configured as license clients. However, the steps to set up communication with the license server are not needed.

Leasing Licenses from a License Server

If the concurrent user count is greater than its leased license limit, a license client requests the license server to increase its capacity until the maximum lease limit (MUC) is met.

When the number of concurrent users drops, the client relinquishes the leases it no longer needs. When the license server receives license lease requests, it first verifies that the client has been allocated the licenses it is requesting. The license server then checks that it has sufficient licenses before granting the request.

Reserved licenses are leased for 10 days at a time. Incremental leases are leased from a configurable time of 24 hours to a maximum of 10 days. Clients can renew their licenses based on lease renewal interval as configured in the license server. The reply sent by the license server includes a new lease expiration date, which is the minimum of the current time plus the incremental lease time and the license allocation expiration date. If a client does not renew a license before the lease expires, the license server reclaims the license. The renewal interval can be 4 hours to a maximum of 24 hours.

A minimum lease interval of 24 hours is built-in. Once a client acquires an incremental license lease, it is kept for at least 24 hours even if the load diminishes on the client.

> **NOTE:** For IVS, license leasing is applicable to only the root IVS. IVS is supported on platforms (requires an IVS license) except for virtual appliances and MAG Series devices.

Related Documentation

- About License Management

Obtaining, Entering, and Upgrading Your License Keys

To take full advantage of your device you must access the Pulse Secure License Management System, provide your licensing hardware ID and authorization code(s) to obtain your license keys, and sign in to the admin console to enter the license keys you receive from Pulse Secure.

A licensing hardware ID is a unique 16-character code Pulse Secure uses to identify your particular device when generating license keys. You can find the system's licensing hardware ID above the menu options in the serial console and at the bottom of the admin console.

An authorization code is a pass key required to generate and activate license keys you or your company have purchased for your device. You receive your authorization code(s) after you purchase your device and associated product and feature licenses. Figure133 shows the license key generation and activation flow.

Figure133: License Key Generation and Activation
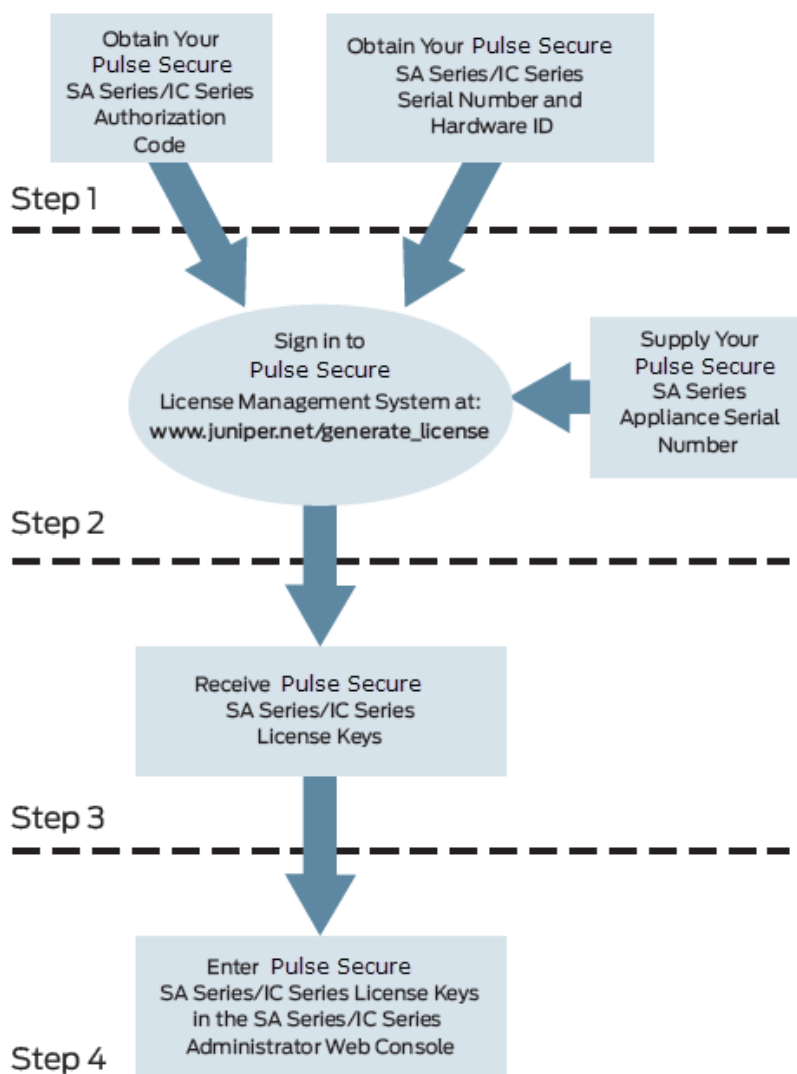
The package you download from the Juniper Networks License Management System or the e-mail message you receive from Juniper Networks might contain different types of licenses:

- **User license keys**—Enables you to host as many users as are specified in the license key code. User license keys are additive, meaning that you can expand the number of users that can access a device by simply acquiring an additional user license key and adding it to your configuration. For example, if you initially purchase a 100 user license and then purchase another 100 user license in the future, your device can accommodate up to 200 users.

- **Access feature license keys**— Allows you to enable system access methods. These keys are available for a variety of access methods including Network Connect and Secure Application Manager, Pulse Collaboration, and advanced access feature licenses.

- **In Case of Emergency (ICE) license keys**—Allows you to activate the emergency mode, and provides licenses for additional users on a device and Pulse Secure devices for up to eight weeks for periodic testing and transitioning to permanent licenses, if necessary. The ICE license does not include optional features.

- **Evaluation license keys**—Allows you to enable and roll out the latest functionality for a limited time before deciding whether or not to purchase license keys and enable the new functionality on a permanent basis. Evaluation license keys are valid for 8 weeks also (Eg: "POLSEC-EVAL-8W")

You can use the System > Configuration > Licensing tab to enter the license keys for your site, to view their expiration dates, and to delete them (if needed).

Ensure that you read the license agreement, which is accessible from the Licensing tab, before submitting your license key. The license agreement available from the Licensing tab is the same text displayed in the serial console during the initial setup.

Related Documentation

- Configuring License Options

- Activating and Deactivating Emergency Mode

Configuring License Options

Use the Juniper Networks License Management System to create and enter new license keys or to transfer license keys to replace a device. Detailed instructions are located on the Juniper Networks License Management System located at https://www.juniper.net/generate_license.

Related Documentation

- Obtaining, Entering, and Upgrading Your License Keys

- Activating and Deactivating Emergency Mode

Activating and Deactivating Emergency Mode

The emergency mode feature allows you to temporarily enable a system for a large number of users.

To activate the user license key enables you to host as many users as are specified in the license key code. User license keys are additive, meaning that you can expand the number of users that can access the device by simply acquiring an additional user license key and adding it to your configuration. For example, if you initially purchase a 100 user license and then purchase another 100 user license in the future, your device can then accommodate up to 200 users. in emergency mode, you must first install an In Case of Emergency (ICE) license using the standard system license installation procedure. Then, when the emergency occurs, you can easily activate emergency mode through the system's web console. When your emergency has passed, you should then deactivate the emergency mode.

> **(i)** NOTE: The ICE license is permanent until you activate emergency mode. Activating emergency mode switches the ICE license to a temporary license and only enables you to operate in emergency mode for 8 weeks. Once the ICE license expires, all features disappear and your users can no longer access the device using the emergency mode.

To activate or deactivate emergency mode:

1. In the Web console, select **System > Configuration > Licensing**.

2. Find the In Case of Emergency License entry in the license list. Sample ICE license names include:

   - MAGX600-ICE

   - ACCESS-ICE-25PC

3. Click the **Enable** link on the right side of the license column to activate emergency mode or click **Disable** to deactivate it.

When you enable and disable emergency mode, the system decrements the corresponding license in 5 minute intervals.

Related Documentation

- Obtaining, Entering, and Upgrading Your License Keys

- Configuring License Options

Configuring a Device as a License Server

The following outlines the steps to configuring a device as a license server. These steps assume that you have already performed the license key generation and activation steps.

> **ⓘ** NOTE: For MAG Series Pulse Secure Gateways, you must use the Connect Secure personality as the license server.
>
> Virtual appliances cannot be configured as a license server.

After you download or receive your license keys by using email:

1. In the admin console of the license server, choose **System > Configuration > Licensing > Licensing Summary**.

2. Click on the license agreement link. Read the license agreement and, if you agree to the terms, continue to the next step.

3. Enter your license key(s) and click **Add**.

4. Click the **Configure Clients** tab.

5. Select the **Enable Licensing server** checkbox.

6. (optional) Click **Advanced Settings** and enter the following values:

   - Incremental Lease Duration

   - Lease Renewal Interval

7. Click **Save Changes**.

An example of the summary section after you configure your clients is shown in Figure134.

Figure134: Configured License Server With No Clients

Summary:				
Feature	**Total**	**Reserved**	**Leased**	**Client Auto-leasing**
Concurrent Users	11010	150	150	●
Premier Java Remote Desktop Applet	1000	75	75	●

Then:

1. In the admin console of the license server, choose **System > Configuration > Licensing > Configure Clients**.

2. Click **New Client**.

3. Enter the Client ID. The ID is defined on the client device under System > Configuration > Licensing > Configure Server.

4. Enter the client password and confirm it. The password is defined on the client device under System > Configuration > Licensing > Configure Server.

5. (optional) Enter the client configuration expiration date.

6. Select the client's platform from the list. The list includes the latest PCS hardware devices like the PSA 7000. Figure 135 shows a list of the client's platforms.

Figure 135 Client Platforms
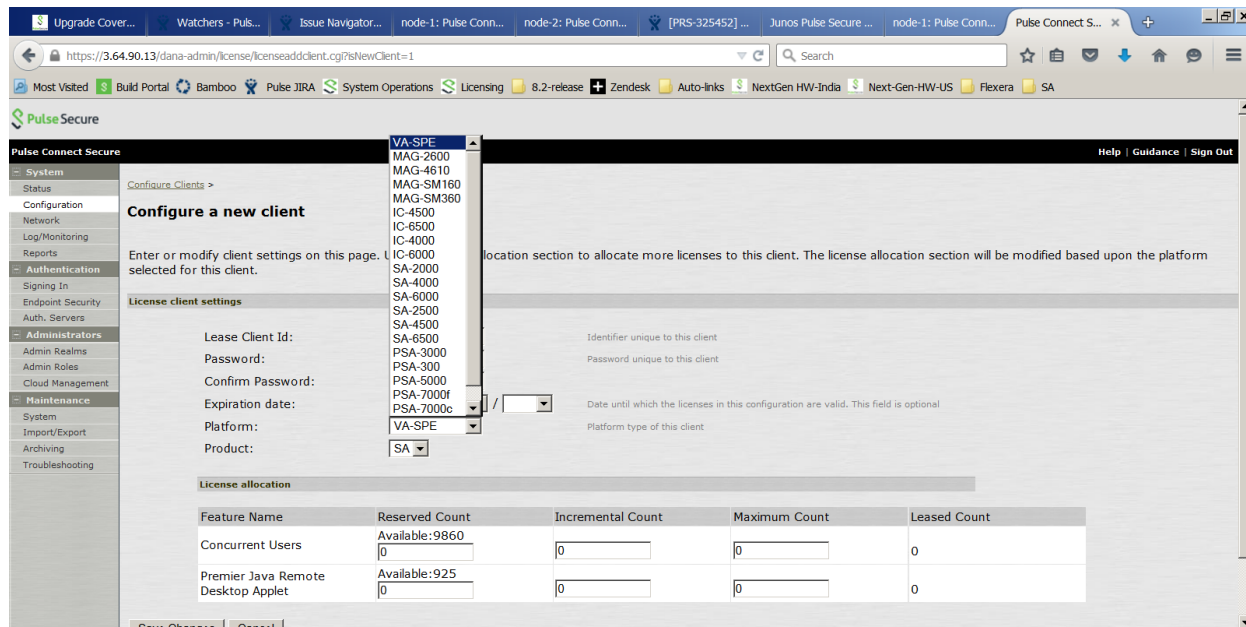
7. Select the product type to be configured to (SA).

8. For each feature you want to lease to this client, enter:

   • **Reserved Count**— the number of licenses to reserve for this client. The reserve count must be less than the available amount displayed.

   • **Incremental Count**— the incremental number of licenses to grant when the client requests more licenses. If the number of licenses on the client plus this incremental value is greater than the maximum count, no additional licenses are granted.

   • **Maximum Count**— the maximum number of licenses a client can receive for this feature. This value must be equal to or greater than the reserved count.

   Available counts are updated as you configure the client. For example, see Figure136.

9. Click **Save Changes**.

The License clients table displays the client information you entered. If the client is a member of a cluster, the cluster name is also displayed after the client contacts the server.

Figure136: Available Counts Are Updated As Clients Are Configured

Related Documentation

- Configuring a Device as a License Client

- Configuring the Automated Downloading of License Keys

- Disabled Features on a License Server

Configuring a Device as a License Client

The following procedure assumes that you have already performed the license key generation and activation steps.

The steps in this topic assume that you have already performed the license key generation and activation steps.

After you download or receive your license keys by using email:

1. In the admin console, choose **System > Configuration > Licensing > Licensing Summary**.

2. Click on the license agreement link. Under Installed License Details, read the license agreement and, if you agree to the terms, continue to the next step.

3. Enter your license key(s) and click **Add**.

4. Click **Save Changes**.

An example of the client's License Summary tab is shown in Figure137.

Figure137: Client Window After Installing Member License

To configure this device as a license server client:

1. In the admin console, choose **System > Configuration > Licensing > Configure Server**.

2. Enter the name of the license server. You can specify the IP address or hostname.

3. Enter a unique ID for this client. This ID is used to communicate and verify this client with the license server.

   IDs can contain alphanumeric characters. There is no restriction on the number of characters.

   You will need to enter this ID on the license server when adding clients.

4. Enter and confirm a password for this client.

   You will need to enter this password on the license server when adding clients.

5. Select the network to communicate with the license server from the Preferred Network menu.

   If the preferred network is configured correctly and enabled, it is used. Otherwise, the internal network is used.

6. Select the **Verify SSL Certificate** checkbox if you want the client to verify the server's SSL certificate when establishing communication with it. Figure 138 shows the list of licenses applicable for specific platform groups

Figure 138 Licenses applicable for specific platform groups

If this client is part of a cluster, you can change configuration information for this node or any node within the same cluster by selecting the node name from the pull-down menu. You can also select **Enter cluster** to update general cluster configuration information.

Related Documentation

- Configuring a Device as a License Server

- Configuring the Automated Downloading of License Keys

Configuring the Automated Downloading of License Keys

Administrators are given the option to automatically download license keys from the Juniper Networks LMS server using the Download and Install button on the Download Licenses page. This feature is applicable for all license keys (subscription-based, capacity-based, and so forth.) Communication is done via https.

The automated download process uses the following schedule for contacting the LMS server:

- Once every 3 months if there are no expiring licenses.

- Once a month if a license is expiring within the next 3 months

- Once a week if a license is expiring within the next 2 weeks

- Once a day if a license grace period is expiring within the next week

To configure the device for automated query and downloading of license keys from the Juniper Networks LMS system:

1. In the admin console, choose **System > Configuration > Licensing > Download Licenses**.

2. Select the preferred network to communicate with the Juniper Networks LMS system.

3. Select **Enabled** if you want to automate the downloading of license keys.

4. Click **Save Changes**.

You can, at any time, click **Download and Install** to immediately query the Juniper Networks LMS server and download updated license keys.

Related Documentation

- Configuring a Device as a License Server

- Configuring a Device as a License Client

Disabled Features on a License Server

The following windows and features are disabled in the administrator console when a device is configured as a license server:

- System > Status> Meeting Schedule

- System > Status > Virtual Desktop Sessions

- System > Configuration > Secure Meeting

- System > Configuration > User Record Synchronization

- System > Configuration > Sensors

- System > Configuration > Virtual Desktops

- System > Configuration > NCP

- System > Network > Network Connect

- System > Clustering

- System > Virtual Systems

- System > IF-MAP Federation

- System > Log/Monitoring > Sensors

- System > Log/Monitoring > User Access

- System > Log/Monitoring > Client Logs

- Maintenance > Push Config

- Maintenance > Troubleshooting > Monitoring > Cluster

- Maintenance > Troubleshooting > User Session

- Maintenance > Archiving > Secure Meetings

- Maintenance > Import/Export > IVS

In addition, the following services are halted on the license server:

- Mail proxy services

- Meeting processes

- Agentman daemon

- Federation server

- Federation client

CHAPTER 32 Network and Host Administration

Network and Host Administration Overview

When you install and initially set up the device, you use the serial port console to set basic network and host settings. To get started, you must use the serial console to configure these settings for the internal interface. You have the option to use the serial console to configure network and host settings for the external interface and the management interface. The network and host settings you configure with the serial port console include:

- IPv4 address
- Netmask
- Default gateway
- Speed and duplex
- MTU
- DNS

- Default domain

- WINS

Once the internal interface has been configured, you can use the admin console Network Settings pages to modify settings for the internal interface, to enable and configure the external interface and the management interface, and to configure or manage advanced networking features, including:

- Hostname

- IPv6 addresses

- VLAN ports

- Virtual ports

- Route table entries

- Host mapping table entries

- ARP cache entries

- Neighbor discovery cache entries

- System date and time (manual configuration) or NTP

Related Documentation

- Using the Serial Port

- Configuring the Internal Port

- Configuring the External Port

- Using the Internal and External Ports

- Using the Management Port

- Configuring VLAN Ports

- Using Virtual Ports

- Configuring Network Services

- Managing the Routes Table

- Managing the Hosts Table

- Managing the ARP Table

- Managing the Neighbor Discovery Table

- Configuring the System Date and Time

Configuring the Internal Port

The internal port connects to the local area network (LAN). The internal port settings are configured when you run the setup wizard from the serial console as part of the installation procedure. You can use the System > Network pages to make changes to the configuration.

To modify the internal port configuration:

1. Select **System > Network > Internal Port > Settings** to display the configuration page.

   Figure139 shows the configuration page for Policy Secure.

   Figure140 shows the configuration page for Connect Secure.

2. Complete the configuration as described in Table 101.

3. Save your changes.

Figure139: Internal Port Configuration Page

Figure140: Connect Secure Internal Port Configuration Page
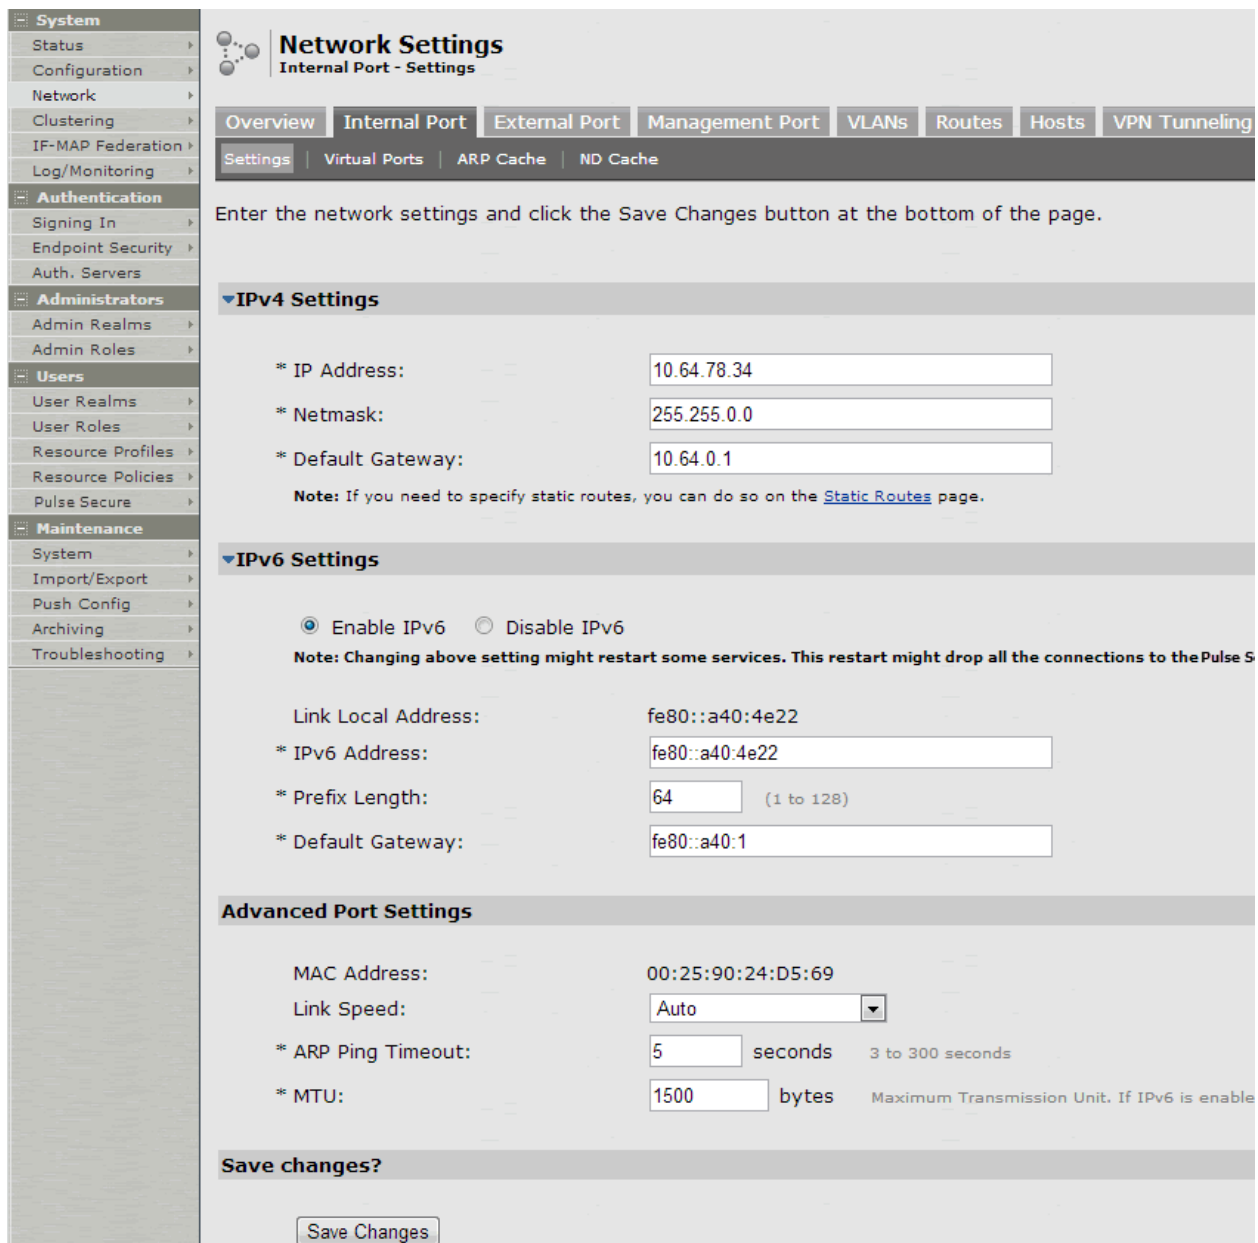
Table101: Internal Port Configuration Guidelines

Settings	Guidelines
IPv4 Settings	
IP Address	Assign an IP address. You must assign an IPv4 address to the internal interface.
	An IP address is an identifier for a computer or device on a TCP/IP network. Networks using the TCP/IP protocol route messages based on the IP address of the destination.
	The format of an IPv4 address is a 32-bit numeric address written as four numbers separated by periods. Each number can be 0 to 255.
Netmask	Assign a netmask. A netmask indicates which part of an IP address indicates network identification and which part indicates the host identification. For example, the IP address and netmask 10.20.30.1 255.255.255.0 (or 10.20.30.1/24) refer to all the hosts in the 10.20.30.0 subnet. The IP address and netmask 10.20.30.1 255.255.255.255 (or 10.20.30.1/32) refer to a single host.

Default Gateway	Specify the IPv4 address for the default gateway for the routing domain to which the device belongs.
	A gateway is the router that resides at the point of entry to the current routing domain, often called the default gateway.
IPv6 Settings	
Enable IPv6 / Disable IPv6	Disabled by default. Enable to support access from IPv6 endpoints.
	When you enable IPv6, the system acquires a link local address.
	If you switch from enabled to disabled, the system clears the link local address.
Link Local Address	Display the autoconfigured link local address (after you have enabled and saved the IPv6 configuration).
IPv6 Address	Specify a routable IPv6 address, such as a global unicast address that your network plan has provisioned for this host and interface. Automatic configuration methods are not supported. You must specify the appropriate address manually.
Prefix Length	Specify how many of the higher order contiguous bits of the IPv6 address comprise the prefix (the network portion of the IPv6 address). The default is 64.
Gateway	Specify the IPv6 address for the default gateway for the routing domain to which the device belongs.
	A gateway is the router that resides at the point of entry to the current routing domain, often called the default gateway.
Advanced Settings	
MAC Address	Display the MAC address for the interface.
Link Speed	Specify the speed and duplex combination for the interface.
	If you run SNMP_GET and then change the Link Speed value, you must wait at least 5 minutes after submitting the change before running SNMP_GET again.
ARP Ping Timeout	(IPv4 only.) Specify how long the system should wait for responses to Address Resolution Protocol (ARP) requests before timing out. Cluster nodes send ARP requests to the gateways of other nodes to determine if they are properly communicating with one another.
	If you have not deployed a cluster, the system does not use this setting. If the node belongs to a cluster, the timeout interval that you specify is synchronized across the cluster. In multisite clusters, you can override this setting for the individual nodes in the cluster using options in the System > Clustering page. Use caution when changing this setting in active/passive clusters, however, because the system also uses the ARP Ping Timeout setting on the Internal tab as a failover timer for the VIP.
MTU	Specify the maximum transmission unit.
	If IPv6 is enabled, the valid range is 1280 to 1500. If IPv6 is not enabled, the valid range is 576 to 1500.
	We recommend you retain the default MTU setting (1500) unless you must change the setting for troubleshooting purposes.

Related Documentation

- Network and Host Administration Overview

Configuring the External Port

The external port connects to the Internet. You can use the System > Network pages to configure the external port.

To configure the external port:

1. Select **System > Network > External Port > Settings** to display the configuration page.

   Figure 138 shows the configuration page for Policy Secure.

   Figure 139 shows the configuration page for Connect Secure.

2. Complete the configuration as described in Table 102.

3. Save your changes.

Figure141: Policy Secure External Port Configuration Page

Figure142: Connect Secure External Port Configuration Page

Table102: External Port Configuration Guidelines

Settings	Guidelines
Use Port?	
Use Port?	Select **Enabled** to use the port; otherwise, select **Disabled**.
IPv4 Settings	
IP Address	Specify an IP address. An IP address is an identifier for a computer or device on a TCP/IP network. Networks using the TCP/IP protocol route messages based on the IP address of the destination. The format of an IPv4 address is a 32-bit numeric address written as four numbers separated by periods. Each number can be 0 to 255.

Netmask	Specify a netmask. A netmask indicates which part of an IP address indicates network identification and which part indicates the host identification. For example, the IP address and netmask 10.20.30.1 255.255.255.0 (or 10.20.30.1/24) refer to all the hosts in the 10.20.30.0 subnet. The IP address and netmask 10.20.30.1 255.255.255.255 (or 10.20.30.1/32) refer to a single host.
Default Gateway	Specify the IPv4 address for the default gateway for the routing domain to which the device belongs. A gateway is the router that resides at the point of entry to the current routing domain, often called the default gateway.

IPv6 Settings

Enable IPv6 / Disable IPv6	Disabled by default. Enable to support access from IPv6 endpoints. When you enable IPv6, the system acquires a link local address. If you switch from enabled to disabled, the system clears the link local address.
Link Local Address	Display the autoconfigured link local address (after you have enabled and saved the IPv6 configuration).
IPv6 Address	Specify a routable IPv6 address, such as a global unicast address that your network plan has provisioned for this host and interface. Automatic configuration methods are not supported. You must specify the appropriate address manually.
Prefix Length	Specify how many of the higher order contiguous bits of the IPv6 address comprise the prefix (the network portion of the IPv6 address). The default is 64.
Gateway	Specify the IPv6 address for the default gateway for the routing domain to which the device belongs. A gateway is the router that resides at the point of entry to the current routing domain, often called the default gateway.

Advanced Settings

MAC Address	Display the MAC address for the interface.
Link Speed	Specify the speed and duplex combination for the interface. If you run SNMP_GET and then change the Link Speed value, you must wait at least 5 minutes after submitting the change before running SNMP_GET again.
ARP Ping Timeout	(IPv4 only.) Specify how long the system should wait for responses to Address Resolution Protocol (ARP) requests before timing out. Cluster nodes send ARP requests to the gateways of other nodes to determine if they are properly communicating with one another. If you have not deployed a cluster, the system does not use this setting. If the node belongs to a cluster, the timeout interval that you specify is synchronized across the cluster. In multisite clusters, you can override this setting for the individual nodes in the cluster using options in the System > Clustering page. Use caution when changing this setting in active/passive clusters, however, because the system also uses the ARP Ping Timeout setting on the Internal tab as a failover timer for the VIP.
MTU	Specify the maximum transmission unit. If IPv6 is enabled, the valid range is 1280 to 1500. If IPv6 is not enabled, the valid range is 576 to 1500. We recommend you retain the default MTU setting (1500) unless you must change the setting for troubleshooting purposes.

Related Documentation

- Network and Host Administration Overview

Using the Internal and External Ports

The internal port, also known as the internal interface, handles all LAN requests to resources, listening for Web browsing, file browsing, authentication, and outbound mail requests. You configure the internal port by providing IP address, gateway, DNS server and domain, and MTU settings during the initial setup of Pulse Connect Secure. You can also change them on the System > Network > Internal Port > Settings tab. Alternatively, you can deploy the appliance in dual-port mode to listen for incoming Web and mail proxy SSL connections on an external port.

The external port, also known as the external interface, handles all requests from users signed into Pulse Connect Secure from outside the customer LAN, for example, from the Internet. Before sending a packet, Pulse Connect Secure determines if the packet is associated with a TCP connection that was initiated by a user through the external interface. If that is the case, Pulse Connect Secure sends the packet to the external interface. All other packets go to the internal interface.

The routes that you specify for each interface apply after Pulse Connect Secure has determined whether to use the internal or external interface. No requests are initiated by Pulse Connect Secure from the external interface, and this interface does not accept any other connections (except ping and traceroute connections). All requests to any resource are issued from the internal interface.

> **NOTE:** If you enable the external port, then, by default, administrators may no longer sign in from an external location. You can open the external port for administrators on the Administrators > Admin Realms > RealmName > Authentication Policy > Source IP tab.

Using the Management Port

This topic describes how to configure the management port. It includes the following information:

- Management Port Overview
- Supported Platforms
- Configuring the Management Port
- Using the Serial Console to Configure the Management Port
- Configuring Administrator Access

Management Port Overview

You connect the management port to an Ethernet switch or router that is part of your internal local area network (LAN) and that can connect to your network management infrastructure. When the management port is enabled, the following traffic is directed out the management port: archiving (FTP/SCP), NSM, NTP, push config, SNMP, syslog. When the management port is not enabled, that traffic uses the internal port.

> **NOTE:** On Policy Secure systems, you cannot configure the user realm configuration, the RADIUS client configuration, or the Infranet Enforcer configuration to use the management port.

Supported Platforms

The following hardware platforms are equipped with a management port:

- IC6500
- SA6000, SA6500
- MAG4610
- MAG Series SM160, SM360, SM860

Configuring the Management Port

To configure the management port:

1. Select **System > Network > Management Port > Settings** to display the configuration page.

   Figure 140 shows the configuration page for Policy Secure.

   Figure 141 shows the configuration page for Connect Secure.

2. Complete the configuration as described in Table 103.

3. Save your changes.

Figure143: Policy Secure Management Port Configuration Page

Figure144: Connect Secure Management Port Configuration Page

Table103: Management Port Configuration Guidelines

Settings	Guidelines
Use Port?	
Use Port?	Select **Enabled** to use the port; otherwise, select **Disabled**.
IPv4 Settings	

IP Address	Specify an IP address. An IP address is an identifier for a computer or device on a TCP/IP network. Networks using the TCP/IP protocol route messages based on the IP address of the destination. The format of an IPv4 address is a 32-bit numeric address written as four numbers separated by periods. Each number can be 0 to 255.
Netmask	A netmask indicates which part of an IP address indicates network identification and which part indicates the host identification. For example, the IP address and netmask 10.20.30.1 255.255.255.0 (or 10.20.30.1/24) refer to all the hosts in the 10.20.30.0 subnet. The IP address and netmask 10.20.30.1 255.255.255.255 (or 10.20.30.1/32) refer to a single host.
Default Gateway	Specify the IPv4 address for the default gateway for the routing domain to which the device belongs. A gateway is the router that resides at the point of entry to the current routing domain, often called the default gateway.

IPv6 Settings

Enable IPv6 / Disable IPv6	Disabled by default. Enable to support network management traffic over IPv6 networks. When you enable IPv6, the system acquires a link local address. If you switch from enabled to disabled, the system clears the link local address.
Link Local Address	Display the autoconfigured link local address (after you have enabled and saved the IPv6 configuration).
IPv6 Address	Specify a routable IPv6 address, such as a global unicast address that your network plan has provisioned for this host and interface. Automatic configuration methods are not supported. You must specify the appropriate address manually.
Prefix Length	Specify how many of the higher-order contiguous bits of the IPv6 address comprise the prefix (the network portion of the IPv6 address). The default is 64.
Gateway	Specify the IPv6 address for the default gateway for the routing domain to which the device belongs. A gateway is the router that resides at the point of entry to the current routing domain, often called the default gateway.

Advanced Settings

MAC Address	Display the MAC address for the interface.
Link Speed	Specify the speed and duplex combination for the interface. If you run SNMP_GET and then change the Link Speed value, you must wait at least 5 minutes after submitting the change before running SNMP_GET again.
ARP Ping Timeout	(IPv4 only.) Specify how long the system should wait for responses to Address Resolution Protocol (ARP) requests before timing out. Cluster nodes send ARP requests to the gateways of other nodes to determine if they are properly communicating with one another. If you have not deployed a cluster, the system does not use this setting. If the node belongs to a cluster, the timeout interval that you specify is synchronized across the cluster. In multisite clusters, you can override this setting for the individual nodes in the cluster using options in the System > Clustering page. Use caution when changing this setting in active/passive clusters, however, because the system also uses the ARP Ping Timeout setting on the Internal tab as a failover timer for the VIP.
MTU	Specify the maximum transmission unit. If IPv6 is enabled, the valid range is 1280 to 1500. If IPv6 is not enabled, the valid range is 576 to 1500. We recommend you retain the default MTU setting (1500) unless you must change the setting for troubleshooting purposes.

Using the Serial Console to Configure the Management Port

To configure management port network settings from the serial console:

1. Start a serial console session.

2. Select item 1, **System Settings and Tools**.

3. Select item 10, **Configure Management port**. The text indicates if the option is enabled or disabled.

4. Enter the network settings for the Management Port, as prompted.

> *i* NOTE: If you enable the Management Port but neglect to configure the IP address and netmask, the port reverts to a disabled state. Also, you cannot clear Management Port settings from the serial console when the port is disabled, though you can clear them from within the admin console.

5. When prompted to accept the changes, if they are correct, enter y. Otherwise, repeat the process to correct the settings.

6. Close the serial console.

Configuring Administrator Access

You can configure the Administrators > Admin Realm > Authentication Policy > Source IP restrictions configuration to enable administrator sign-in through the management port.

You can use Administrator realms to control administrator access to system ports, including the management port.

To control administrator access to the management port:

1. Enable the management port.

2. Perform one of the following steps:

   - Select **Administrators > Admin Realms > Admin Users** to modify the default admin users realm.

   - Select **Administrators > Admin Realms**, then click **New**, to create a new administrator realm.

3. Select the **Authentication Policy > Source IP**.

4. Select one of the following options:

   - Allow users to sign in from any IP address—Allows users to sign in from any IP address to satisfy the access management requirement.

   - Allow or deny users from the following IP addresses—Specifies whether to allow or deny users access from all of the listed IP addresses, based on their settings.

     To specify access from an IP address:

     - Enter the IP address and netmask.

     - Select either Allow to allow users to sign in from the specified IP address, or Deny to prevent users from signing in from the specified IP address.

5. Select the available options to allow administrators to sign in to all available ports, to the management port or the internal port only, or to restrict them from signing in to any of the ports. In some cases you may inadvertently limit administrative access completely. If this occurs, you can reconfigure the ports by way of the serial console.

   Select from the following available options:

   - Enable administrators to sign in on the management port.

   - Enable administrators to sign in on the internal port.

   - Enable administrators to sign in on the external port.

   Figure 142 shows the configuration page for administrator access.

Figure145: Configuring Administrator Access

6. Click **Save Changes**.

Related Documentation

- Network and Host Administration Overview

- Using the Serial Port

Configuring VLAN Ports

Your network design might include VLANs to provide network segmentation. When connected to a trunk port on a VLAN-enabled switch, the system encounters traffic from all VLANs. This is useful for network designs with separate VLANs for separate classes of users or endpoints, and for making the system accessible from all VLANs. You can use RADIUS attributes to place different users in different network segments.

The system supports IEEE 802.1Q VLAN tagging. You must define a VLAN port for each VLAN. The internal port must be assigned to the root system and must be marked as the default VLAN. Routes to servers reachable from the VLAN interfaces must have the next-hop gateway set to the configured gateway for the VLAN interface, and must have the output port defined as the VLAN port.

When you save the configuration for a new VLAN port, the system creates two static routes by default:

- The default route for the VLAN pointing to the default gateway.

- The interface route to the directly connected network.

To configure a VLAN port:

1. Select **System > Network > VLANs**.

2. Click **New Port** to display the configuration page.

Figure 143 shows the configuration page for Policy Secure.

Figure 144 shows the configuration page for Connect Secure.

3. Complete the configuration as described in Table 104.

4. Save your changes.

Figure146: Policy Secure VLAN Port Configuration Page

Figure147: Connect Secure VLAN Port Configuration Page

Table104: VLAN Port Configuration Guidelines

Settings	Guidelines
Use Port?	
Use Port?	Select **Enabled** to use the port; otherwise, select **Disabled**.
VLAN Settings	
Port Name	Specify a name that is unique across all VLAN ports that you define on the system or cluster. Only alphanumeric characters, "-", or "_" are allowed.
VLAN ID	Specify a number between 1 and 4094. The VLAN ID assignment must be unique on the system.
IPv4 Settings	
IP Address	Specify an IP address and netmask combination that is from the same network as the VLAN. VLAN IP addresses must be unique. You cannot configure a VLAN to have the same network as the internal port. For example, if the internal port is 10.64.4.30/16 and you configure a VLAN as 10.64.3.30/16, you might get unpredictable results and errors. The format of an IPv4 address is a 32-bit numeric address written as four numbers separated by periods. Each number can be 0 to 255.

Netmask	Specify a netmask. A netmask indicates which part of an IP address indicates network identification and which part indicates the host identification. For example, the IP address and netmask 10.20.30.1 255.255.255.0 (or 10.20.30.1/24) refer to all the hosts in the 10.20.30.0 subnet. The IP address and netmask 10.20.30.1 255.255.255.255 (or 10.20.30.1/32) refer to a single host.
Default Gateway	Specify the IPv4 address for the default gateway for the routing domain to which the device belongs.
	A gateway is the router that resides at the point of entry to the current routing domain, often called the default gateway.

> **NOTE:** Link speed, ARP ping timeout, and MTU settings are inherited from the internal port configuration.

Related Documentation

- Configuring the Internal Port
- Network and Host Administration Overview

Using Virtual Ports

This topic describes virtual ports. It includes the following information:

- Configuring Virtual Ports
- Using Device Certificates with Virtual Ports

Configuring Virtual Ports

You can use virtual ports to provide different groups of users access to the same system using different IP aliases and domains.

Virtual ports are associated with the physical internal port and physical external port. The virtual port shares all of the network settings with the associated physical port, except for the IP address.

When you configure virtual ports, you in essence are creating name-IP address pairs. The names and IP addresses must be unique in your network. An alias can include IPv4 addresses, IPv6 addresses, or both. However, the corresponding IP protocol must be enabled on the physical port for the addresses to take effect.

To configure a virtual port:

1. Select **System > Network >** *PortName***> Virtual Ports**. *PortName* is Internal Port or External Port.

2. Click **New Port** to display the configuration page.

   Figure 145 shows the configuration page for Policy Secure.

   Figure 146 shows the configuration page for Connect Secure.

3. Complete the configuration as described in Table 105.

4. Save your changes.

Figure148: Policy Secure Virtual Port Configuration Page

Figure149: Connect Secure Virtual Port Configuration Page

Table105: Virtual Port Configuration Guidelines

Settings	Guidelines
Name	Specify a name for the virtual port. The names and IP addresses in the virtual port configuration must be unique in your network.
Physical Port	Display the name of the physical port associated with the virtual port. The virtual port inherits link speed, ARP ping timeout, and MTU settings from the physical port configuration.
IPv4 Address	Specify an IPv4 address. An alias can include IPv4 addresses, IPv6 addresses, or both. However, the corresponding IP protocol must be enabled on the physical port for the addresses to take effect.
IPv6 Address	Specify an IPv6 address. An alias can include IPv4 addresses, IPv6 addresses, or both. However, the corresponding IP protocol must be enabled on the physical port for the addresses to take effect.

Using Device Certificates with Virtual Ports

Virtual ports can be used to create multiple fully qualified domain names for user sign-in. When a user tries to sign in using the IP address defined in a virtual port, the system presents the certificate associated with the virtual port to initiate the SSL transaction.

You can approach the digital certificate security and virtual ports implementation in either of the following ways:

- Associate all hostnames with a single certificate—With this approach, you use a single wildcard certificate to validate the identity of all system hostnames, regardless of which hostname is used to sign in. A wildcard certificate includes a variable element in the domain name, making it possible for users who sign in from multiple hosts to map to the "same" domain. For example, if you create a wildcard certificate for *.yourcompany.com, the system uses the same certificate to validate its identity to users who sign in to employees.yourcompany.com as it does to users who sign into partners.yourcompany.com.

- Associate each hostname with its own certificate—With this approach, you associate different hostnames with different certificates. Create a virtual port for each hostname. A virtual port activates an IP alias on a physical port. For example, you can create two virtual ports on a single appliance, mapping the first virtual port to the IP address 10.10.10.1 (sales.yourcompany.com) and the second virtual port to the IP address 10.10.10.2 (partners.yourcompany.com). Then you can associate each of these virtual ports with its own certificate, ensuring that users authenticate through different certificates.

To associate certificates with virtual ports:

1. Create virtual ports.

2. Import the device certificates.

3. Associate the device certificates with the virtual ports:

   1. Select **System > Configuration > Certificates > Device Certificates**.

   2. Click the link of the device certificate you want to configure to display the configuration page.

      Figure 147 shows the configuration page for Policy Secure. The configuration page for Connect Secure is similar.

   3. Use the controls in the "Present certificate on these ports" section to associate ports with the certificate.

Figure150: Policy Secure Certificate Details Page

Related Documentation

- Configuring the Internal Port
- Configuring the External Port
- Using Device Certificates
- Network and Host Administration Overview

Configuring the System Date and Time

You can use the admin console to set the system date and time manually or by configuring a network time protocol (NTP) server. The system supports NTPv4, which is backwards compatible with NTPv3 and NTPv2.

> BEST PRACTICE: We recommend you use NTP to synchronize the date and time clocks on all network systems. Using NTP obviates issues that might occur with cluster synchronization, network communication that uses time-sensitive protocols, such as SAML, and implementation of time-based policies, such as local authentication server account expiration. In addition, using NTP as a standard in your network rationalizes timestamps in logs, which facilitates reporting and troubleshooting.
>
> On a VMware virtual appliance, the cockpit data may be erased each hour if the same NTP server is not defined here, on the Connect Secure license server, and on the ESXi server.

To set the system date and time:

1. Select **System > Status** to display the System Status dashboard.

2. Click the System Date and Time **Edit** link to display the configuration page.

3. Complete the configuration as described in Table 106.

4. Save the configuration.

Table106: Date and Time Configuration Guidelines

Settings	Guidelines
Time Zone	Select your time zone. Selecting the appropriate time zone enables the system to automatically adjust the time for Daylight Saving Time changes.
Use NTP Server	
NTP Server	Specify the fully qualified domain name or IP address for the NTP server.
Key	If you are using NTPv4, specify the symmetric key. The key must be pre-synchronized with the NTP server. For example, if you want to configure NIST's clock as the NTP server, you must request a key beforehand and have NIST send that key to you. The key for MD5 is in the following format: *KeyNumber* M *KeyValue* The key for SHA1 is in the following format: *KeyNumber* SHA1 *KeyValue*
Secondary NTP Server	Specify the fully qualified domain name or IP address for a secondary NTP server. If the system fails three times to reach the primary server or if the authentication fails three times, the system attempts to contact the secondary NTP server.
Key	If you are using NTPv4, specify the symmetric key.
Update Interval	Specify an update interval. The maximum interval is 999999 (enforced by the user interface).
Set Time Manually	
Date	Specify the date. You can click **Get from Browser** to automatically populate the Date and Time fields.

Time	Specify the time and select **AM** or **PM**.

Related Documentation

- Displaying System Status

- Network and Host Administration Overview

Configuring Network Services

You configure DNS and WINS services when you initially configure the system with the serial console. If necessary, you can use the System > Network > Overview page to modify the configuration. You can also use this page to configure a hostname.

The network services overview page also displays the node name (if the node belongs to a cluster), and the status and interface statistics for the internal port, external port, and management port.

To configure network services:

1. Select **System > Network > Overview** to display the configuration page.

   Figure 148 shows the configuration page for Policy Secure.

   Figure 149 shows the configuration page for Connect Secure.

2. Complete the configuration as described in Table 107.

3. Save your changes.

Figure151: Policy Secure Network Services Configuration Page

Figure152: Connect Secure Network Services Configuration Page
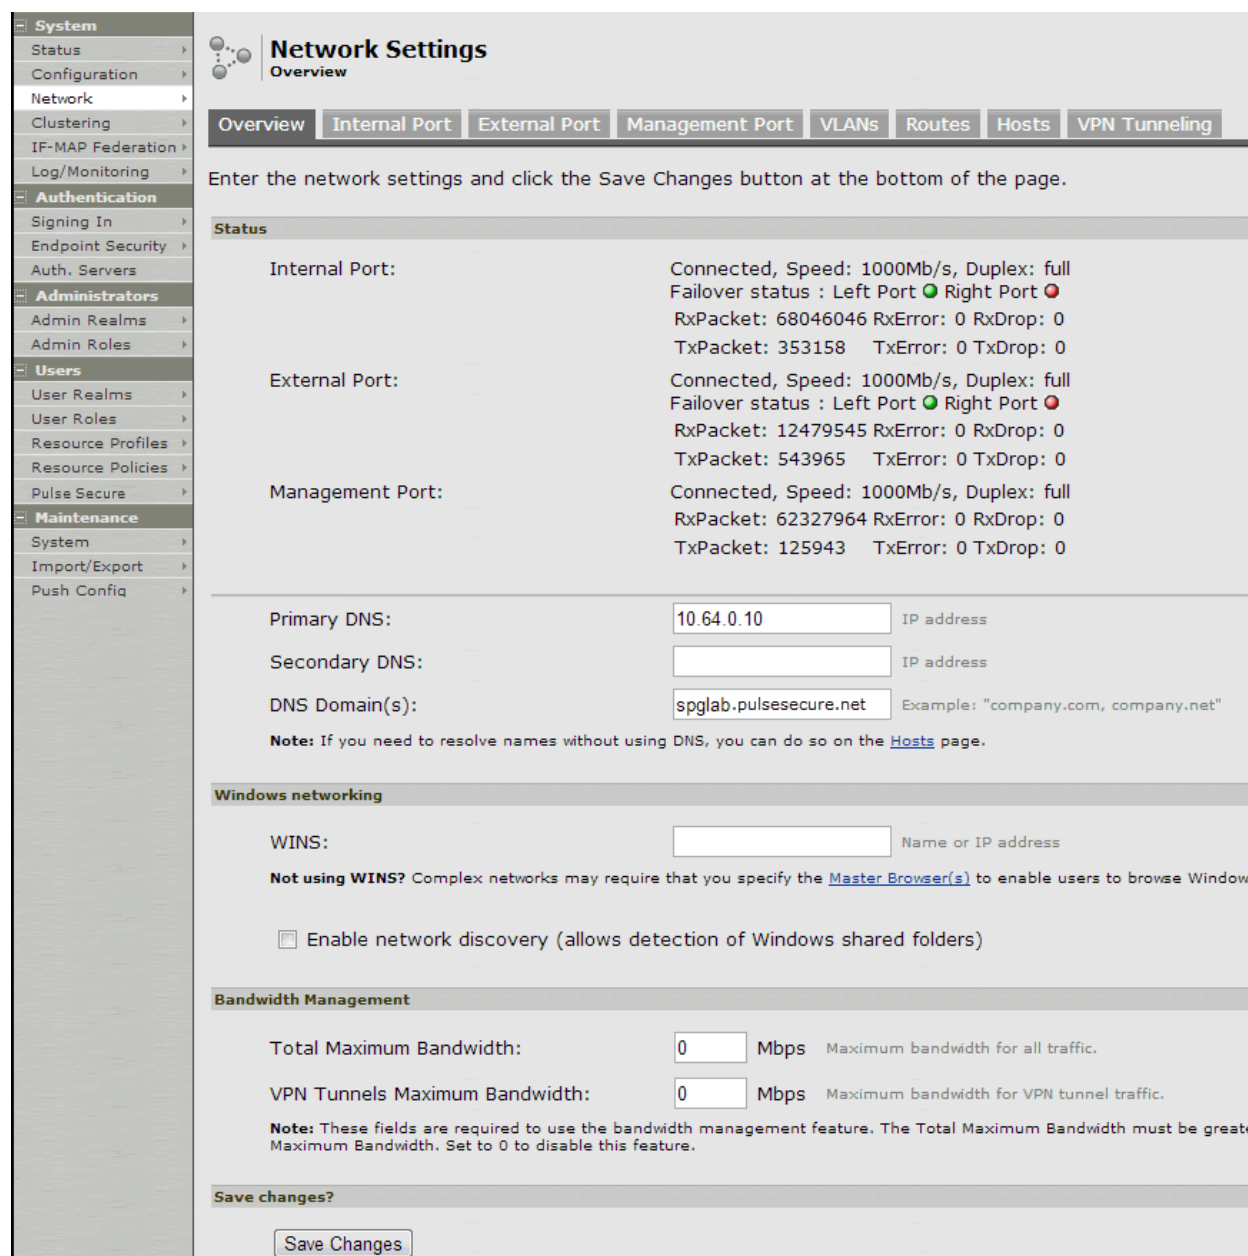
Table107: Network Services Configuration Guidelines

Settings	Guidelines
Status	
Status	Display interface statistics for the internal port, external port, and management port.
Network Identity	
Hostname	Specify a fully qualified hostname. For example, **domain.company.com**. The hostname cannot exceed 30 characters
DNS Name Resolution	
Primary DNS	Specify the IP address for the primary DNS server.

Secondary DNS	Specify the IP address for the secondary DNS server.
DNS Domain(s)	Specify a comma-separated list of default domains. The system searches the domains in the order they are listed.
Windows Networking	
WINS	Specify the hostname or IP address of a local or remote Windows Internet Naming Service (WINS) server that you use to associate workstation names and locations with IP addresses.
Bandwidth Management	This feature is available only on Connect Secure.
Total Maximum Bandwidth	Specify the maximum bandwidth for all traffic.
VPN Tunnels Maximum Bandwidth	Specify the maximum bandwidth for VPN tunnel traffic. NOTE: The value of total maximum bandwidth must be greater than the value of VPN tunnels maximum bandwidth.

Related Documentation

- Network and Host Administration Overview
- Using the Serial Port

Managing the Routes Table

The system populates the routes table with dynamic, auto-discovered routes. Many networks will not require changes to this routing table. If necessary, you can delete routes or add static routes.

To manage the routes table:

1. Select **System > Network > Routes** to display the routes table.

   Figure 150 shows the routes for Policy Secure.

   Figure 151 shows the routes for Connect Secure.

2. Use the controls described in Table 108 to manage the routes table.

Figure153: Policy Secure Routes Table

Figure154: Connect Secure Routes Table

Table108: Routes Table Controls

Controls	Description
View route table for	Use the controls to change the display to show the route table for internal, external, or management interfaces; and for IPv4 or IPv6 routes.
Delete	Select a row in the table and click **Delete** to delete a route.
New Route	Click **New Route** and complete the configuration to add a route to the table. You must specify a valid IP address, gateway, DNS address, and metric. The metric is a way of comparing multiple routes to establish precedence. Generally, the lower the number (from 1 to 15), the higher the precedence. Thus, a route with a metric of 2 is chosen over a route with a metric of 14. The metric value of zero (0) identifies the route as one that should not be used.

Related Documentation

- Network and Host Administration Overview

Managing the Hosts Table

In general, the system uses the configured DNS servers to resolve hostnames, but it also maintains a local hosts table that can be used for name resolution. The system populates some entries from host-IP address pair settings in your configuration. You can add host-IP address mappings for other hosts that might not be known to the DNS servers used by the system, or in cases where DNS is not reachable.

To manage the hosts table:

1. Select **System > Network > Hosts** to display the hosts table.

   Figure 152 shows the hosts table for Policy Secure.

   Figure 153 shows the hosts table for Connect Secure.

2. Use the controls described in Table 109 to manage the hosts table.

Figure155: Policy Secure Hosts Table

Figure156: Connect Secure Hosts Table

Table109: Hosts Table Controls

Controls	Description
Add	Specify an IP address, hostname, and comment (a description for the benefit of system administrators), and click **Add**.
Delete	Click the delete icon in the last column to delete the row from the table.

Related Documentation

- Network and Host Administration Overview

Managing the ARP Table

ARP stands for Address Resolution Protocol. In IPv4 networking, network nodes use ARP to maintain information about peer network nodes. ARP is used to associate the Layer 3 IP address with a Layer 2 MAC address of neighboring peer nodes. The system maintains an ARP table with dynamic, cached entries, and you can add static entries if necessary. The system caches dynamic entries for up to 20 minutes. Dynamic entries are deleted during a reboot. Static entries are restored after a reboot.

To manage the ARP table:

1. Select **System > Network >** *Port* **> ARP Cache**. *Port* is the Internal Port, External Port, or Management Port tab.

   Figure 154 shows the ARP table for the internal port for Policy Secure.

   Figure 155 shows the ARP table for the internal Port for Connect Secure.

2. Use the controls described in Table 110 to manage the ARP table.

Figure157: Policy Secure ARP Table

Figure158: Connect Secure ARP Cache Table

Table110: ARP Table Controls

Controls	Description
Delete	Select a row in the table and click **Delete** to delete the entry.
Delete Dynamic Entries	Delete all dynamically discovered entries.
Add	Specify an IP address, a MAC address, and click **Add** to add an entry. If you add an entry that has the same IP address as an existing entry, the system overwrites the existing entry with your new entry. Also note that the system does not verify the validity of MAC addresses.

Related Documentation

- Network and Host Administration Overview

Managing the Neighbor Discovery Table

In IPv6 networking, network nodes use the Neighbor Discovery Protocol (NDP) to determine the Layer 2 MAC addresses for neighboring hosts and routers. The system uses NDP to maintain a cache of neighboring routers that are reachable and can forward packets on its behalf.

NOTE: In the current release, you can view discovered neighbors or clear the entire cache, but you cannot add neighbors or delete individual entries.

To manage the neighbor discovery table:

1. Select **System > Network >** *Port* **> ND Cache**. *Port* is the Internal Port, External Port, or Management Port tab.

   Figure 156 shows the neighbor discovery table for the internal port for Policy Secure.

   Figure 157 shows the neighbor discovery table for the internal port for Connect Secure.

2. Use the controls described in Table 111 to manage the neighbor discovery table.

Figure159: Policy Secure Neighbor Discovery Table

Figure160: Connect Secure Neighbor Discovery Table

Table111: Neighbor Discovery Table Controls

Controls	Description
Flush NDP Entries	Delete all dynamically discovered entries.

Related Documentation

- Network and Host Administration Overview

Using IPv6

This topics describes support for using IPv6. It includes the following information:

- Understanding IPv6
- IPv6 Support Overview
- IPv6 Feature Configuration Task Summary

Understanding IPv6

IP version 6 (IPv6) is an Internet Protocol designed to succeed IP version 4 (IPv4). This topic provides an overview of IPv6. It includes the following information:

- About IPv6
- About IPv6 Address Types
- About IPv6 Address Text Representation
- About the IPv6 Unspecified Address
- About the IPv6 Loopback Address
- About IPv6 Address Prefixes
- System Normalization of IPv6 Addresses
- About Neighbor Discovery Protocol

About IPv6

The ongoing expansive growth of the Internet and the need to provide IP addresses to accommodate it is escalating the emergent use of a new IP protocol. IPv6 was designed to satisfy the current and anticipated near future requirements.

IPv4 is widely used throughout the world today for the Internet, intranets, and private networks. IPv6 builds upon the functionality and structure of IPv4 in many aspects, including:

- Larger address space—IPv6 addresses are 128 bits long instead of 32 bits. This expands the address space from 4 billion addresses to over 300 trillion trillion trillion addresses.
- New datagram format—The packet header is both simplified and enhanced to enable more secure and efficient routing.
- Improved fragmentation and reassembly—The maximum transmission unit (MTU) has been increased to 1280 bytes, for example.
- Transition mechanisms—Various network address translation (NAT) and tunneling mechanisms have been developed to support the transition to IPv6.

On February 3, 2011 Internet Assigned Numbers Authority (IANA) allotted the last block of IPv4 addresses to Regional Internet Registries (RIR), so the free pool of IPv4 addresses is now fully depleted. It is expected that in the near future Internet service providers (ISPs) will start issuing IPv6 addresses to their customers.

About IPv6 Address Types

RFC 4291, IP Version 6 Addressing Architecture, describes the following types of IPv6 addresses:

- Unicast. An identifier for a single interface. A packet sent to a unicast address is delivered to the interface identified by that address.
- Anycast. An identifier for a set of interfaces. A packet sent to an anycast address is delivered to *one* of the interfaces identified by that address.

- Multicast. An identifier for a set of interfaces. A packet sent to a multicast address is delivered to *all* interfaces identified by that address.

The *link-local address* is a special IPv6 unicast address that is used in self-traffic and essential network communication, like Neighbor Discovery Protocol (NDP). When you enable IPv6 on a Connect Secure interface, the system generates a link-local address that is derived from the interface MAC address.

When you configure IPv6 addresses for the system interfaces, you manually specify a routable address, such as *global unicast address* or an *anycast address*, depending on your routing implementation and your system deployment. A global unicast address must be globally unique so that it can be specified globally without need for modification. An anycast address represents a service rather than a specific device. An anycast address is not unique, but instead might be configured on each device in a cluster. You are not likely to use multicast addressing with Connect Secure.

About IPv6 Address Text Representation

All IPv6 addresses are 128 bits long, written as 8 sections of 16 bits each. They are expressed in hexadecimal representation, so the sections range from 0 to FFFF. Sections are delimited by colons, and leading zeroes in each section may be omitted. If two or more consecutive sections have all zeroes, they can be collapsed to a double colon.

IPv6 addresses consist of 8 groups of 16-bit hexadecimal values separated by colons (:). IPv6 addresses have the following format:

aaaa:*aaaa*:*aaaa*:*aaaa*:*aaaa*:*aaaa*:*aaaa*:*aaaa*

Each *aaaa* is a 16-bit hexadecimal value, and each *a* is a 4-bit hexadecimal value. The following is a sample IPv6 address:

2001:0DB8:0000:0000:0008:0800:200C:417A

You can omit the leading zeros of each 16-bit group, as follows:

2001:DB8:0:0:8:800:200C:417A

You can compress 16-bit groups of zeros to double colons (::) as shown in the following example, but only once per address:

2001:DB8::8:800:200C:417A

About the IPv6 Unspecified Address

In the IPv6 address space, the special "unspecified address" is 0:0:0:0:0:0:0:0. The compressed representation of the unspecified address is the double-colon (::). The unspecified address must never be assigned to a physical or virtual interface.

About the IPv6 Loopback Address

The special loopback address is the unicast address 0:0:0:0:0:0:0:1. The compressed representation of the loopback address is ::1. The loopback address may be used by a node to send an IPv6 packet to itself. It must not be assigned to a physical or virtual interface. In Release 7.3 and Release 8.0, the loopback address shown in the system Event logs is the special IPv4 loopback address (127.0.0.1). The IPv6 loopback address (::1) is not used in Event logs.

About IPv6 Address Prefixes

An IPv6 address prefix is a combination of an IPv6 prefix address and a prefix length used to represent a block of address space (or a network), similar to the use of an IPv4 subnet address and netmask combination to specify a subnet. An IPv6 address prefix takes the form *ipv6-*

prefix/prefix-length. The *ipv6-prefix* variable follows general IPv6 addressing rules. The */prefix-length* variable is a decimal value that indicates the number of contiguous, higher-order bits of the address that make up the network portion of the address. For example, 2001:DB8::/32 is an IPv6 address prefix, indicating that the first 32 bits make up the network portion of the address.

System Normalization of IPv6 Addresses

The system validates and normalizes IPv6 addresses entered by administrators. The normalized address is the address processed by the system, and it is the address that appears in logs.

Table 112 gives examples of how the system normalizes IPv6 address entries.

Table112: System Normalization of IPv6 Addresses

Example Entry	Normalized Address	Explanation
2001:DB8:1:1::3	2001:DB8:1:1::3	An address specified in compressed format is validated; the system uses the compressed form as the normalized form.
0:0:0::122	::122	Address is validated and normalized to compressed format.
FF01:0:0:0:0:0:0:101	FF01::101	Address is validated and normalized to compressed format.
2001:DB8::10.204.50.122	2001:DB8::ACC:327A	Address is validated and normalized to hexadecimal representation.
::FFFF:10.204.50.122	::FFFF:10.204.50.122	An address specified in compressed format is validated; the system uses the compressed form as the normalized form.

About Neighbor Discovery Protocol

Neighbor discovery protocol (NDP) allows different nodes on the same link to advertise their existence to their neighbors, and to learn about the existence of their neighbors.

Routers and hosts (nodes) use NDP messages to determine the link-layer addresses of neighbors that reside on attached links and to overwrite invalid cache entries. Hosts also use NDP to find neighboring routers that can forward packets on their behalf.

In addition, nodes use NDP to actively track the ability to reach neighbors. When a router (or the path to a router) fails, nodes actively search for alternatives to reach the destination.

IPv6 NDP corresponds to a number of the IPv4 protocols — ARP, ICMP Router Discovery, and ICMP Redirect. However, NDP provides many improvements over the IPv4 set of protocols. These improvements address the following:

- Router discovery—How a host locates routers residing on an attached link.

- Prefix discovery—How a host discovers address prefixes for destinations residing on an attached link. Nodes use prefixes to distinguish between destinations that reside on an attached link and those destinations that it can reach only through a router.

- Parameter discovery—How a node learns various parameters (link parameters or Internet parameters) that it places in outgoing packets.

- Address resolution—How a node uses only a destination IPv6 address to determine a link-layer address for destinations on an attached link.

- Next-hop determination—The algorithm that a node uses for mapping an IPv6 destination address into a neighbor IPv6 address (either the next router hop or the destination itself) to which it plans to send traffic for the destination.

- Neighbor unreachability detection—How a node determines that it can no longer reach a neighbor.

- Duplicate address detection—How a node determines whether an address is already in use by another node.

A router periodically multicasts a router advertisement from each of its multicast interfaces, announcing its availability. Hosts listen for these advertisements for address autoconfiguration and discovery of link-local addresses of the neighboring routers. When a host starts, it multicasts a router solicitation to ask for immediate advertisements.

The router discovery messages do not constitute a routing protocol. They enable hosts to discover the existence of neighboring routers, but they are not used to determine which router is best to reach a particular destination.

NDP uses the following Internet Control Message Protocol version 6 (ICMPv6) messages: router solicitation, router advertisement, neighbor solicitation, neighbor advertisement, and redirect.

NDP for IPv6 replaces the following IPv4 protocols: Router Discovery (RDISC), Address Resolution Protocol (ARP), and ICMPv4 redirect.

IPv6 Support Overview

This topic describes support for IP Version 6 (IPv6) networks. It includes the following information:

- Client Access Summary
- Network Topologies
- IPv6 Support and Limitations for Connect Secure Features

Client Access Summary

Connect Secure release 8.0 and later supports use of VPN Tunneling Connection Profile features to enable dual-stack endpoints to connect the Connect Secure device and access corporate network IPv4 and IPv6 resources. Table 113 summarizes supported access scenarios.

Table113: Connect Secure Release 8.0 Client Access Scenarios

Endpoint	Connect Secure Interface	Tunnel	Resource	Description of the Connection
IPv4/IPv6	IPv4	IPv4-in-IPv4	IPv4	All resource access policies are supported for access to IPv4 resources.
		IPv6-in-IPv4	IPv6	You must configure IPv4 and IPv6 address pools in the VPN Tunneling connection profile configuration. Access to IPv6 resources using VPN Tunneling connection profiles only.
IPv4/IPv6	IPv6	IPv4-in-IPv6	IPv4	You must configure IPv4 and IPv6 address pools in the VPN Tunneling connection profile configuration. All resource access policies are supported for access to IPv4 resources.
		IPv6-in-IPv6	IPv6	Access to IPv6 resources using VPN Tunneling connection profiles only.

Table 114 provides a summary of Pulse Secure client and system software requirements for IPv6 deployment types.

Table114: Pulse Secure Client Support for IPv6 Deployment Types

Connect Secure	Pulse Client	IPv4-in-IPv4	IPv4-in-IPv6	IPv6-in-IPv4	IPv6-in-IPv6
8.0	5.0	Yes	Yes	Yes	Yes
8.0	4.0	Yes	Yes	Yes	No
8.0	3.1	Yes	Yes	No	No
7.4	5.0	Yes	Yes	Yes	No
7.3	5.0	Yes	Yes	No	No

Network Topologies

Connect Secure release 8.0 and later supports Pulse Secure client access to the IPv6 corporate network using VPN Tunneling Connection Profile features.

The role-based VPN Tunneling Connection Profile determines the IP addresses assigned to the client Pulse Secure client virtual adapter. In this configuration, you must configure an IPv4 address pool. You configure an IPv6 address pool to enable access to IPv6 resources. When a client connects and is mapped to a role that includes the VPN Tunneling Connection configuration, the Pulse Secure client virtual adapter is assigned all address from each pool—both an IPv4 and IPv6 address—and a single SSL tunnel is set up. When a connection is made to the system IPv4 address, the IPv4 traffic is encapsulated in the IPv4 tunnel ("4 in 4" tunneling), and the IPv6 traffic is encapsulated in the IPv4 tunnel ("6 in 4"). When a connection is made to the system IPv6 address, the IPv4 traffic is encapsulated in the IPv6 tunnel ("4 in 6"), and the IPv6 traffic is encapsulated in the IPv6 tunnel ("6 in 6").

In this release, the DNS server used by the system must be reachable by IPv4 and must be able to resolve both A and AAAA DNS queries. Only the VPN Tunneling Connection Profile is supported for access to IPv6 resources. All other connection options and resource policies are not supported for access to IPv6 resources.

Figure 158 shows a deployment topology for dual-stack-enabled endpoints that access the system over an ISP IPv4 network.

Figure161: Dual Stack Endpoint Access Over ISP IPv4 Network
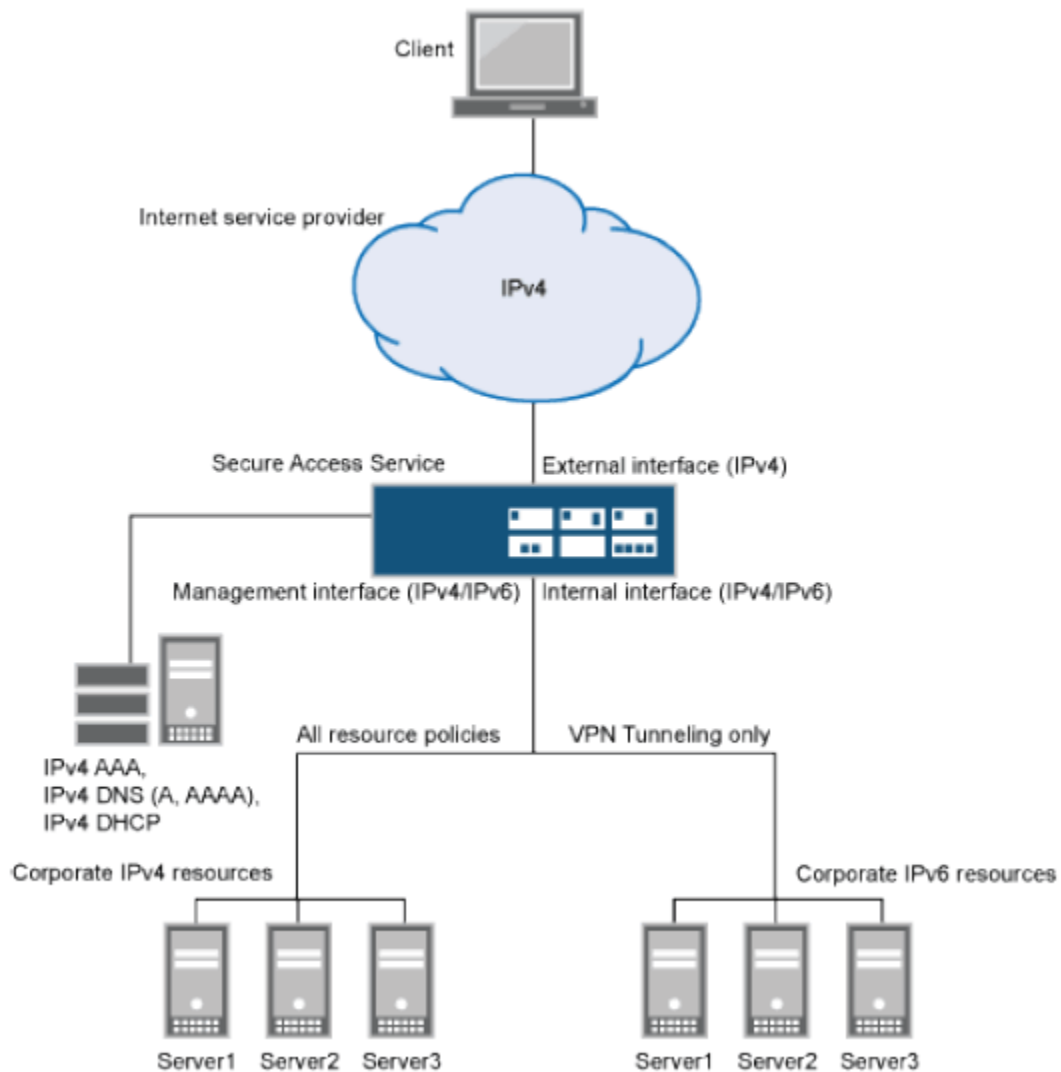
Figure 159 shows a deployment topology for dual-stack-enabled endpoints that access the system over an ISP IPv6 network.

Figure162: Dual Stack Endpoint Access Over ISP IPv6 Network
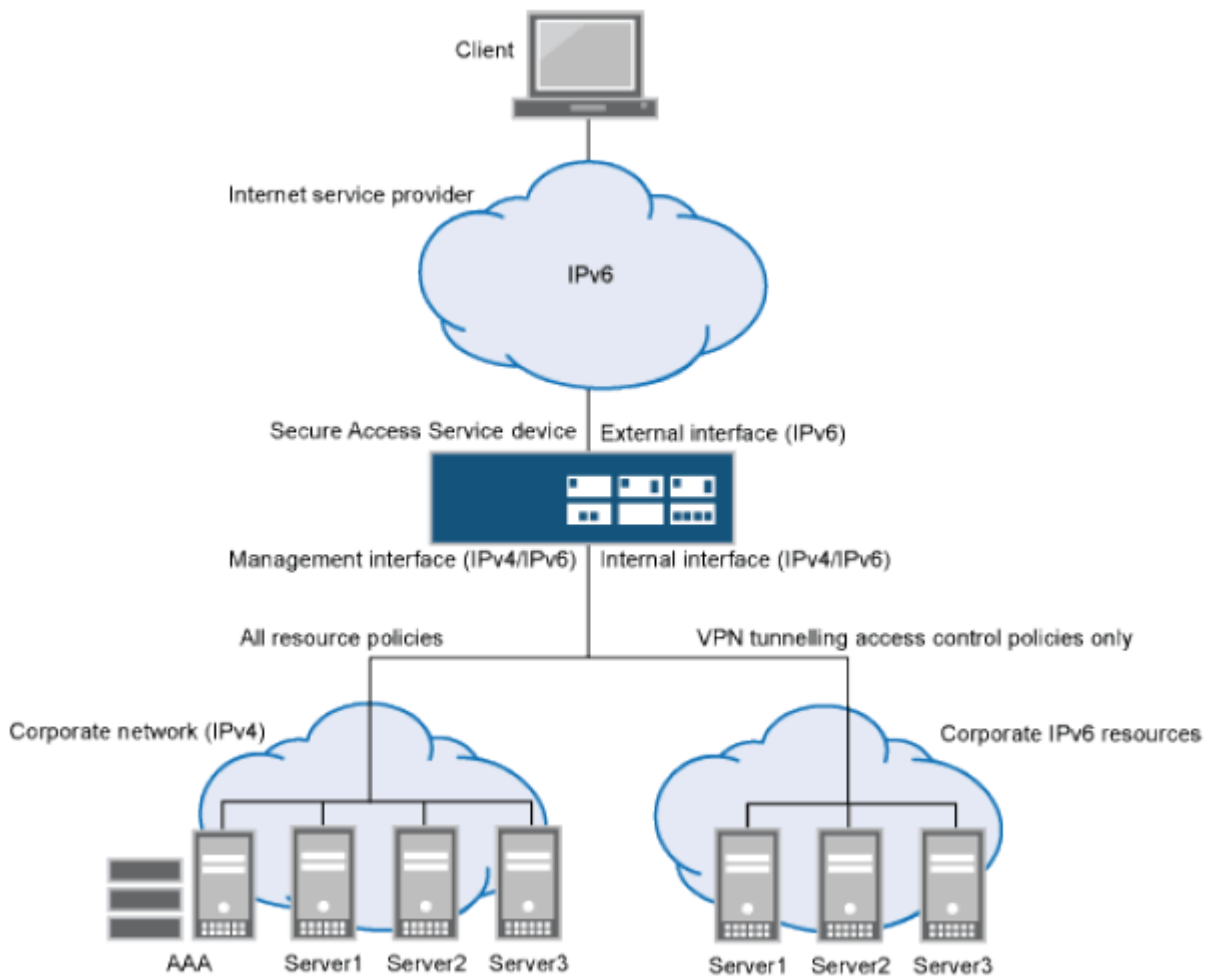
IPv6 Support and Limitations for Connect Secure Features

Table 115 summarizes IPv6 support and limitations for Connect Secure features for Release 8.0 and later.

Table115: Summary of IPv6 Support

Feature	Release 8.0 and later
Pulse Secure client access	Only the Pulse Secure client supports IPv6. The legacy Network Connect client does not support IPv6. The following behavior is expected for this release: • Endpoints must have dual-stack enabled in order to access IPv6 resources over IPv4 networks. • VPN Tunneling Connection Profiles support IPv4 and IPv6 address pools. • VPN Tunneling Connection Profiles do not support ESP mode for IPv6 resource access. If a connection is configured for ESP mode, it automatically falls back to use SSL mode. • On dual-stack endpoints, VPN Tunneling split tunneling rules are supported for IPv4-based routes only. The IPv4 traffic allowed by a split tunneling policy is forwarded to the system in an IPv4/IPv6 tunnel. • VPN Tunneling features are not supported for IPv6 in the IVS. • Pulse Secure SAM, Legacy WSAM, and Legacy JSAM do not support IPv6. Pulse Secure clients on the following platforms support VPN Tunneling connections for IPv6 resource access: • Windows 7 (32 and 64 bit) • Windows 8 (32 and 64 bit) • Mac OS/X Snow Leopard, Lion, Mountain Lion Host Checker supports IPv6. Third-party Host Checker functionality is supported to the extent that it is IPv6-capable. For example, the following third-party components might require endpoints to connect over IPv4: • Downloading antivirus signature updates from third-party vendors • Downloading Windows Patches from Microsoft download servers. Pulse Collaboration does not support IPv6.
Administrator and management access	The internal interface and management interface can be configured with an IPv4 address or dual stack (IPv4 and IPv6). The internal interface and management interface cannot be configured with only an IPv6 address because the system uses IPv4 for the connections with network services, including AAA, DHCP, and DNS. Typically, administrators access the administrator GUI through the internal interface or management interface, but you may enable administrator access through the external interface on the Authentication > Admin Realms > Admin Users > Authentication Policy > Source IP page.
Configuration through the serial console	You cannot view or configure IPv6 network settings with the serial console.
External interface configuration	IPv4, IPv6, or both is supported.
Internal interface configuration	IPv4 or both is supported. In other words, the internal interface must be configured for IPv4 connections; in addition, it may be configured for IPv6 connections. It may not be configured for IPv6 only.
Management interface configuration	IPv4 or both is supported. In other words, the internal interface must be configured for IPv4 connections; in addition, it may be configured for IPv6 connections. It may not be configured for IPv6 only.
Virtual interface configuration	An interface alias may include IPv4 addresses, IPv6 addresses, or both. However, the corresponding IP protocol must be enabled on the physical interface for the addresses to take effect.
VLAN configuration	IPv4-only.
Clustering	This release supports IPv6 configuration for active/active and active/passive clusters. The existing intracluster communication mechanism is preserved. The intracluster communication occurs over the IPv4 corporate network through the internal interfaces.
License server	IPv4 must be enabled for the "preferred network" you select for licensing protocol communication.

Web server	The implementation for IPv6 does not require reconfiguration of the system after upgrade. The Web server can listen for and accept IPv4 or IPv6 clients, and it can differentiate between them for internal purposes and for logging purposes.
ActiveSync	The implementation for IPv6 does not require reconfiguration of the system after upgrade. ActiveSync functionality is available to users connecting from IPv4 or IPv6 endpoints to an IPv4 backend server. Connection to an IPv6 backend server is not supported.
Connection profiles	After upgrading to Release 8.0, you can update your VPN Tunneling Connection Profile configuration to enable IPv6 address assignments to Pulse Secure clients. You must configure a static IPv6 address pool. DHCPv6 is not supported.
	Also note that the IP address server configuration on the System > Network > VPN Tunneling page does not support filtering for IPv6 address pools. In active/active clusters, separate connection profiles need to be created with different IPv6 address pools for each node.
	Split Tunneling policies are not supported for IPv6 resources.
	WINS is not used in IPv6 networks; therefore, WINS settings are not applicable for connection profiles used for IPv6 access.
	The server-side proxy feature does not support IPv6.
Resource policies	You can configure VPN Tunneling Connection Profiles to enable access to all IPv6 resources in your corporate network; however, you cannot configure VPN Tunneling Access Control Policies to allow or deny access to particular IPv6 resources. As a workaround for this Release 8.0 limitation, we recommend you deploy firewall security to restrict access to IPv6 resources.
	NOTE: To enable access to IPv6 resources, the DNS server used by the system must be reachable by IPv4 and must be able to resolve AAAA DNS queries.
	Core Access - Rewriter
	The implementation for IPv6 does not require reconfiguration of the system after upgrade. After upgrade, IPv6 endpoints can access internal IPv4 resources through the system. This applies to all system content rewriters: HTML, Java Script, Applets, VB Script, Flash, CSS, XML, PDF.
	You cannot configure Web Rewriting Policies for IPv6 resources.
	Core Access - Passthrough proxy
	The system passthrough proxy modes are based on hostnames or ports, not IP addresses. Therefore, the implementation for IPv6 does not require reconfiguration of the system after upgrade. Note, however, that in virtual hostname mode, your DNS server must be configured to resolve the virtual hostname to the system IP address, which can be an IPv4 or IPv6 address. Update entries in your DNS server accordingly.
	You cannot configure Passthrough Proxy Policies for IPv6 resources.
	Core Access - Hosted Java applets
	The implementation for IPv6 does not require reconfiguration of the system after upgrade. All hosted Java applets, including the premier Java RDP applet, work on IPv4 or IPv6 clients.
	You cannot configure policies that require access to hosted Java applets at IPv6 addresses.
User role VPN Tunneling options	The Split Tunneling option is not supported for IPv6. If it is enabled for a role, the system treats it as disabled.
	Route Precedence:
	• If **Tunnel Route** is selected, the client cannot access its local IPv6 network and IPv6 traffic is blocked, except DHCPv6, ICMPv6, and loopback traffic going to the physical adapter. If Route Monitoring is enabled, only IPv4 route monitoring is performed.
	• If **Endpoint Route** is selected, the client can access its local IPv6 network. Route Monitoring should be disabled.
	The Multicast option is not supported for IPv6 resources.
Role/Realm Source IP restrictions	You can specify IPv4 or IPv6 Source IP restrictions at both the role and the realm level.
	If the device is deployed behind a NAT64 device, it sees traffic coming from an IPv4 address. In this case, your Source IP restrictions should be based on the NATed IPv4 addresses.

Session roaming	You can manage session roaming across IPv6 subnets. If you enable unlimited session roaming, a session is maintained within an IPv4 network, within an IPv6 network, or from IPv4 to IPv6 and vice versa. If you configure limited session roaming, you can specify IPv4 or IPv6 subnets within which the session is maintained. However, with limited session roaming, you cannot allow sessions to roam from IPv4 to IPv6 networks, or vice versa.
Logging	The logging system can process and parse logs containing IPv6 addresses. Release 8.0 supports communication with external log systems and utilities, such as syslog, SNMP, and archiving that are reachable by IPv4 only.
Network tools	Beginning with Release 7.3, the ping6 and traceroute6 were added to the admin graphical user interface console network tools page. We do not support DNS resolution for hosts with IPv6 addresses. Hence, you cannot use hostname targets with ping6 and traceroute6.

IPv6 Feature Configuration Task Summary

IPv6-related features are not enabled by default. After you upgrade to system software release 8.0 or later, perform the tasks summarized in Table 116 to make the device ready for IPv6 traffic.

Table116: IPv6 Feature Configuration Task Summary

Action	Documentation
Enable IPv6 for the external port and configure an IPv6 address.	"Configuring the External Port"
Enable IPv6 for the internal port and configure an IPv6 address.	"Configuring the Internal Port"
Enable IPv6 for the management port and configure an IPv6 address.	"Using the Management Port"
Configure IP aliases and IPv6 addresses for virtual ports.	"Using Virtual Ports"
Re-create a cluster deployment with IPv6 configuration for external interfaces.	"Example: Creating an Active/Active Cluster That Supports IPv6 Client Access" "Example: Creating an Active/Passive Cluster that Supports IPv6 Client Access"
If you use source IP policies, configure them so that source IP restrictions are based on IPv6 addresses.	"Specifying Source IP Access Restrictions"
Configure IPv6 address assignment for VPN Tunneling Connection Profiles. DHCPv6 is not supported. Also note that the IP address server configuration on the System > Network > VPN Tunneling page does not support filtering for IPv6 address pools.	"Creating VPN Tunneling Connection Profiles"
If you permit roaming sessions but limit to roaming within the specified subnet, configure the role session option so that the subnet is defined by netmask for IPv4 and prefix length for IPv6 networks.	"Specifying Role Session Options"
View and manage the neighbor discovery cache. In Release 8.0, you can view discovered neighbors or clear the entire cache, but you cannot add neighbors or delete individual entries.	"Managing the Neighbor Discovery Table"
View IPv6 routes in the IP route table. In Release 8.0, you can view discovered IPv6 routes, but you cannot add or delete them from the route table.	"Managing the Routes Table"
Review logs. The logging infrastructure accommodates IPv6 addresses, and you can create custom filters based on IPv6 address patterns.	"Using Log Filters"
Become familiar with IPv6 network connectivity test tools, such as ping6 and traceroute6.	"Using Network Troubleshooting Commands"

Related Documentation

- Network and Host Administration Overview

- Juniper Networks Day One: Exploring IPv6

Configuring SSL Options

Use the System > Configuration > Security > SSL Options page to change the default security settings. We recommend that you use the default security settings, which provide maximum security, but you may need to modify these settings if your users cannot use certain browsers or access certain Web pages.

TLS_DHE_RSA_WITH_AES_128_CBC_SHA TLS_DHE_RSA_WITH_AES_256_CBC_SHA cipher suites are supported. Both these ciphers use RSA for server authentication and ephemeral Diffie-Hellman (DHE) for key exchange. RSA server certificate is required for these ciphers. Only TLS_DHE_RSA_WITH_AES_256_CBC_SHA is available with the **Accept 168-bit and greater** option. In the Custom SSL Cipher configuration, TLS_DHE_RSA_WITH_AES_128_CBC_SHA is available only when **AES-Medium** is selected and TLS_DHE_RSA_WITH_AES_256_CBC_SHA is available only when **AES-High** is selected. Both ciphers are lower in priority over the other widely used cipher suites.

To configure SSL options:

1. Select **System > Configuration > Security > SSL Options** to display the configuration page.

   shows the configuration page.

2. Complete the configuration as described in .

3. Save the configuration.

Figure163: SSL Options Configuration Page

Table117: SSL Options Configuration Guidelines

Settings	Guidelines
SSL FIPS Mode option	Enable FIPS mode. See the *Connect Secure FIPS Level 1 Feature Guide*.
Allowed SSL and TLS Version	Specify encryption requirements for clients. By default, the system requires SSL version 3 and TLS. The system honors this setting for all Web server traffic and all types of clients. You can require users who have older browsers that use SSL version 2 to update their browsers, or you can change this setting to allow SSL version 2, SSL version 3, and TLS.
Allowed Encryption Strength	• **Accept only 168-bit and greater**—The system gives preference to 256-bit AES over 3DES. • **Accept only 128-bit and greater**—The default. The system gives preference to RC4 ciphers. You can require users to have this level of encryption strength or change this default to an option compatible with the user base. • **Accept 40-bit and greater**—The system gives preference to RC4 ciphers. Older browsers that predate the change in the U.S. export law in year 2000 that required 40-bit cipher encryption for international export, can still use 40-bit encryption. • **Custom SSL Cipher Selection**—Specify a combination of cipher suites for the incoming connection from the user's browser. If you select the **AES/3DES** option, the system gives preference to 256-bit AES over 3DES. NOTE: When using 168-bit encryption, some Web browsers may still show 128-bit encryption (the gold lock on the browser status bar) even though the connection is 168-bit. This is typically a limitation of the browser's capability. NOTE: If you are using the IC6500 FIPS version, you can choose High, Medium, or Low security cipher suites. AES/3DES High and AES Medium are recommended for FIPS deployment.
Encryption Strength option	Normally, the allowed encryption strength is enforced after an SSL session is established, so that a user connecting with a disallowed encryption strength receives a Web page describing the problem. Enable this option to prevent a browser with a weak cipher from establishing a connection.
SSL Handshake Timeout option	Determines how many seconds elapse before the SSL handshake times out. The default is 60 seconds.
SSL Legacy Renegotiation Support option	SSL and Transport Layer Security (TLS) renegotiations can be subjected to man-in-the-middle (MITM) attacks that can lead to abuse. A new TLS extension (defined in RFC 5746) ties renegotiations to the TLS connections they are being performed over to prevent these kinds of attacks. The SSL Legacy Renegotiation Support option is enabled by default and allows renegotiation between clients and servers even if they do not support the new TLS extension. Disable this option to not allow renegotiations between clients and servers that do not support the new TLS extension. A web server restart is required when you change the value of this option.
ActiveSync Client Certificate Configuration	Use these controls to enforce client certificate requirement for activesync access on the selected ports, including virtual ports. When enabled, all ActiveSync clients must present a client authentication certificate to the system to be able to connect using ActiveSync. Non-ActiveSync access (like web browser based access to the host, NC, JSAM, WSAM, Pulse, WTS, IKEv2 and so forth) on the port/interface on which the ActiveSync client certificate is required might not work properly. We recommend you use a separate port or interface exclusively for ActiveSync access and then enable client certificate requirement for the port intended for ActiveSync access.

Related Documentation

- Network and Host Administration Overview

Configuring Miscellaneous Security Options

You can use the System > Configuration > Security > Miscellaneous page to configure the following security options:

- Persistent cookie options—Choose whether to preserve or delete persistent cookies when a session is terminated.

- Lockout options—You can configure lockout options to protect the system from denial of service (DoS), distributed denial of service (DDoS), and password-guessing attacks.

- Last login—You can choose whether to show users the time and IP address their user ID was used to sign in.

To configure cookie and lockout options:

1. Select **System > Configuration > Security > Miscellaneous** to display the configuration page.

   Figure 161 shows the configuration page.

2. Complete the configuration as described in Table 118.

3. Save the configuration.

Figure164: Miscellaneous Security Options Configuration Page

Table118: Miscellaneous Security Options Configuration Guidelines

Settings	Guidelines
Delete all cookies at session termination	
Delete / Preserve	For convenience, the system sets persistent cookies on the user's machine to support functions such as multiple sign-in, last associated realm, and the last sign-in URL. For additional security or privacy, you can choose not to set them.
Include Pulse Connect Secure's session cookie in URL	
Include / Not Include	Mozilla 1.6 and Safari may not pass cookies to the Java Virtual Machine, preventing users from running JSAM and Java applets. To support these browsers, the system can include the user session cookie in the URL that launches JSAM or a Java applet. By default, this option is enabled, but if you have concerns about exposing the cookie in the URL, you can disable this feature.
Lockout options	
Rate	Specify the number of failed sign-in attempts to allow per minute.
Attempts	Specify the maximum number of failed sign-in attempts to allow before triggering the initial lockout. The system determines the maximum initial period of time (in minutes) to allow the failed sign-in attempts to occur by dividing the specified number of attempts by the rate. For example, 180 attempts divided by a rate of 3 results in a initial period of 60 minutes. If 180 or more failed sign-in attempts occur within 60 minutes or less, the system locks out the IP address being used for the failed sign-in attempt.
Lockout period	Specify the length of time (in minutes) the system must lock out the IP address.
Last Login options	
Time / IP Address	Display the day and time and IP address the user last logged in to the system. For users, this information appears on their bookmark page. For administrators, this information appears on the System Status Overview page. These settings do not apply to the custom start page option on Role UI Options page.

The following scenario illustrates how lockout settings work. For example, assume the following settings:

- Rate = 3 failed sign-in attempts per minute

- Attempts = 180 maximum allowed in initial period of 60 minutes (180/3)

- Lockout period = 2 minutes

The following sequence illustrates the effect of these settings:

1. During a period of 3 minutes, 180 failed sign-in attempts occur from the same IP address. Because the specified value for Attempts occurs in less than the allowed initial period of 60 minutes (180/3), the system locks out the IP address for 2 minutes (fourth and fifth minutes).

2. In the sixth minute, the system removes the lock on the IP address and begins maintaining the rate of 3 failed sign-in attempts/minute. In the sixth and seventh minutes, the number of failed sign-in attempts is 2 per minute, so the system does not lock the IP address. However, when the number of failed sign-in attempts increases to 5 in the eighth minute, which is a total of 9 failed sign-in attempts within 3 minutes, the system locks out the IP address for 2 minutes again (ninth and tenth minutes).

3. In the eleventh minute, the system removes the lock on the IP address and begins maintaining the rate of 3 failed sign-in attempts per minute again. When the rate remains below an average of 3 per minute for 60 minutes, the system returns to its initial monitoring state.

Related Documentation

- Network and Host Administration Overview

Configuring NCP and JCP

The following types of internal protocols are used to communicate between Connect Secure and client applications:

- Network Communications Protocol (NCP)—Standard NCP has been replaced by oNCP. Windows client applications, including the Pulse Collaboration Windows client, WSAM, and Terminal Services fallback to NCP if oNCP fails.

- Optimized NCP (oNCP)—Optimized NCP (oNCP) significantly improves the throughput performance of the client applications over NCP because it contains improvements to protocol efficiency, connection handling, and data compression. Windows client applications, including the Pulse Collaboration Windows client, WSAM, Network Connect and Terminal Services use oNCP by default.

- Java Communications Protocol (JCP)—JCP is the Java implementation of standard NCP. The system uses JCP to communicate with Java client applications, including the Pulse Collaboration Java client, JSAM, and the Java Content Intermediation Engine.

To set NCP options:

1. In the admin console, choose **System > Configuration > NCP**.

2. (Windows clients) Under NCP Auto-Select, select:

   - **Auto-select Enabled** (recommended)—Use the oNCP by default. If you select this option, the system uses oNCP for most client/server communications and then switches to standard NCP when necessary. The system reverts to NCP if the user is running an unsupported operating system, browser type, or combination thereof, or if the client application fails to open a direct TCP connection to the device for any reason (for instance, the presence of a proxy, timeout, disconnect).

   - **Auto-select Disabled**—Always use standard NCP. This option is primarily provided for backwards compatibility.

     > NOTE: If you are using Network Connect to provide client access, we recommend that you exercise caution when employing the Auto-select Disabled option, as Mac and Linux clients cannot connect using the traditional NCP protocol. If you disable the oNCP/NCP auto-selection feature and a UDP-to-oNCP/NCP fail-over occurs, the system disconnects Macintosh and Linux clients because it fails over from UDP to NCP (instead of oNCP), which does not support these users.

3. (Java clients) Under Read Connection Timeout, set the timeout interval for Java clients (15-120 seconds). If client-side secure access methods do not receive data from the system for the specified interval, they try to reestablish a connection. Note that this value does not apply to user inactivity in client applications.

4. (Windows clients) Under Idle Connection Timeout, set the idle connection interval. This timeout interval determines how long the system maintains idle connections for client-side Windows secure access methods.

5. Click **Save Changes**.

Related Documentation

- Network and Host Administration Overview

Using the User Record Synchronization Feature

This topics describes the user record synchronization feature. It includes the following information:

- User Record Synchronization Overview
- Configuring the User Record Synchronization Authentication Server
- Configuring the User Record Synchronization Server
- Configuring the User Record Synchronization Client
- Configuring the User Record Synchronization Database
- Enabling User Record Synchronization
- Scheduling User Record Synchronization Backup

User Record Synchronization Overview

The user record synchronization feature promotes a more consistent user experience by allowing users to retain their bookmarks and individual preferences regardless of which device they log in to.

User record synchronization relies on client-server pairings. The client is the Connect Secure device that users log in to start their remote access. Each client is associated with one primary server and one backup server to store user record data. Clients can be individual appliances or a node within a cluster.

A server in this instance is the Connect secure device that stores the user data records. Each server can be configured to replicate its user record data to one or more peer servers. Servers are identified by a user-defined logical name. The same logical name can be assigned to more than one authentication server to let you associate authentication servers of different types to the same user. For example, SA1 is an ACE authentication server with user1 who creates a bookmark to www.juniper.net. SA2 is an Active Directory authentication server with the same user1. For the www.juniper.net bookmark to be transferred from SA1/ACE/user1 to SA2/AD/user1 you would assign the logical name "Logical1" to both the ACE server on SA1 and the Active Directory server on SA2.

i   NOTE: Cluster VIPs cannot be used as the IP for synchronizing between clients and peers servers.

As long as the logical name is the same, the authentication servers can be different types and different server names and still be associated with a common user. The username must be the same for user record data to be synchronized across the servers. The logical authentication server (LAS) and username combination is what uniquely identifies a user record.

The following user records are synchronized between the client and server:

- Bookmarks
  - Web
  - File
  - Terminal Services
  - JSAM
- Preferences
- Persistent cookies

- Cached passwords

User session data is not synchronized. Persistent cookies, if changed, are synchronized when the user session terminates. All other modifications to the user records are synchronized immediately. User records are stored in cache on the client node prior to being pushed to the servers.

When a user logs in to a client, their data is pulled from the associated server. The pull is performed in the background and does not delay the login process. Users using browsers that do not support JavaScript must manually refresh the index page for updated bookmarks and preferences to appear. For browsers that support JavaScript, users may see a spinning progress indicator and their home page will refresh automatically with updated bookmarks and preferences.

Clients and servers need not be installed with the same system software version as long as they are using version 7.2 or later.

> NOTE: User record synchronization uses port 17425. This port number is not configurable. If you are deploying across a firewall, configure your firewall to allow traffic on this port.

To set up user record synchronization, you perform the following tasks:

1. Enable user record synchronization for each participating client and server, identify which ones are the client and which ones are the server and assign a node name to each client and server.

2. Create a shared secret that is used to authenticate the client with the server and the server to its peer servers.

3. On each server, define which clients and peers are allowed to communicate with the server.

4. On each client, define the servers that handle records for each LAS server.

When enabling this feature, you have several options to initialize the user record database. You can:

- populate the database using user records located in the cache of the client systems.

- populate the database use user records located in the cache of the server systems.

- don't pre-populate the database but populate it as users log in and out of the client system.

If you choose the last option, users may not be able to view their saved bookmarks and preferences until the next time they log in, depending on which client they log in to.

> NOTE: User records may not synchronize if the time clocks on the devices are not in sync. We recommend that you use the same NTP server for each node participating in user record synchronization to keep times accurately adjusted.
>
> The user record synchronization feature will not start automatically after importing a system configuration that has this feature enabled. The workaround is to disable user record synchronization and then enable user record synchronization from the user interface after the configuration import.

Configuring the User Record Synchronization Authentication Server

To set up the authentication server you must define it's logical name:

1. Select **Authentication > Auth Servers**.

2. Click the name of the authentication server you want assign a LAS name.

By assigning the authentication server a LAS name, all users that authenticate using the authentication server are associated with this LAS. In this instance, we are referring to the client nodes, not the user record synchronization server nodes.

3. Select the **User Record Synchronization** check box.

4. Enter a logical name to identify this server.

   This allows you to share user record data across authentication servers on different Connect Secure devices. By assigning a LAS name to an authentication server, you are implicitly assigning it to all users that authenticate with that auth server. The combination of the user's login name and their LAS name uniquely identifies the user's user record across all user record synchronization servers.

5. Click **Save Changes**.

Configuring the User Record Synchronization Server

To set up the user record synchronization server you must define it's peer nodes (optional) and the clients that can access this server.

1. Select **System > Configuration > User Record Synchronization > This Server**.

2. Enter the peer server's node name and IP address, then click **Add**. To specify more than one peer server, enter each server's node name and IP address individually and click **Add**. There is no limit on the number of peer servers you can add.

   Data is replicated from the primary or backup server to its peer servers. If the primary is not available, user data is sent to the backup. User data is then replicated to the peer servers.

3. For each client you want synchronized with this server, enter the client's name and IP address and click **Add**.

Once added, peer servers will have a colored icon next to their name indicating their connection status. Node status is provided to client nodes and LAS mapping servers as well.

Color	Description
Green	Connected
Yellow	Connecting
Gray	Not connected

Configuring the User Record Synchronization Client

To set up the client, you select the primary and backup server you want this client to synchronize with:

1. Select **System > Configuration > User Record Synchronization > This Client**.

2. Select the LAS name you want to synchronize and enter the primary IP of the user record server that will server the user records. If you prefer to synchronize with any available server, select **Any LAS**.

3. Enter the primary and optionally a backup server's IP address and then click **Add**.

   Even if you select Any LAS, you must enter a primary server IP address.

Once added, the primary and backup servers have a colored icon next to their name indicating their connection status.

Configuring the User Record Synchronization Database

With the Database tab, you can delete inactive records from the client cache, retrieve statistics about the database, export and import the data and remove user data from the server's database.

To configure the database:

1. Select **System > Configuration > User Record Synchronization > Database**.

2. Select **Auto-delete inactive synchronized user records from the Cache** to remove inactive user records from the cache. This option does not remove user records from the user record database.

   When this option is selected, the system performs a check every 15 minutes and deletes user records that meet all of the following criteria:

   - There are no active user sessions associated with the user record.

   - The user record does not have any custom settings or the latest version of the user record has been synchronized with the user record database.

   - The authentication server associated with the user record database does not have type "local". For example, the "System Local" auth server that is part of the default configuration has a "local" type, so any user records associated with that auth server will not be auto-deleted. However, user records associated with external authentication servers like Radius or LDAP may be deleted, depending on the two prior criteria.

3. Select **Auto-delete user records from the local synchronization database that have been idle for X days** to permanently remove user records from the database located on the server. Enter the number of days user records must be inactive before being deleted.

   In this instance, "inactive" means that no client as pulled the user record or pushed any modifications to the user record in X days.

4. Click **Retrieve Statistics** to display the number of records in the database. You cannot edit or view records in the database.

5. Under Export, you export user records to a file. The user records can be exported from the user record database, or from the cache. The exported file can be used to pre-populate the user record database on another node.

   - Enter the LAS name of the user records you want to export. If you leave this field blank, all user records are exported. If you enter a LAS name, only user records with the entered LAS name are exported.

   - To encrypt the exported data, select the **Encrypt the exported data with password** check box and enter the password.

   - Click **Export** to export the user records from the specified source (cache or database). You will be prompted where to save the file.

6. Under Import, you import user records into the synchronization database. The user records can be imported from a file or from the cache. Use the Import operation to pre-populate the user record database with user records exported from another node, or with user records from the cache.

   - Click **Browse** to locate the exported file and enter the password if the exported file was encrypted with a password.

   - Select the **Override Logical Auth Servers in imported user records with** check box to replace the LAS name in each imported user record with the LAS name entered.

> For example, you change the LAS name, use this option to update the user records with the new name.

- Click **Import**.

7. Under Delete, specify which user records to permanently remove from the user record database. The options you select apply only to the user record database associated with this server.

- Select **User record with login name and Logical Auth Server** to remove a specific record. The login name and LAS name together uniquely identify a user record. Select this option to remove that record (if it exists).

- Select **User records with Logical Auth Server** to delete all user records with the specified LAS name.

- Select **All user records** to permanently remove user records from the database on this node.

- Click **Delete**.

Enabling User Record Synchronization

The first step in enabling user record synchronizing is to define the node name and the shared secret used to authenticate between the clients and the servers:

1. Select **System > Configuration > User Record Synchronization > General**. See Figure 162.

2. Select the **Enable User Record Synchronization** check box.

3. Enter a unique node name. This name is used when associating a client with a server and is different from the logical name assigned to a server. This node name is also not the same as the cluster node name.

4. Enter the shared secret and confirm it.

   The shared secret is the password used to authenticate the client with its servers and the primary server with its peer servers. Use the same shared secret for all clients and servers participating in user record synchronization.

5. Select whether this node is client only or if this node acts as both a client and server.

6. Click **Save Changes**.

   > *i* NOTE: If you need to make any changes in this window at a later time, you must clear the Enable User Record Synchronization check box and click **Save Changes**. Make your edits, select the Enable User Record Synchronization check box and save your changes.
   >
   > Once you enter a name and shared secret, you cannot clear these fields.

Figure165: User Record Synchronization General Settings Configuration Page

Scheduling User Record Synchronization Backup

You can configure periodic backups of the user record database. User record synchronization backup can be enabled only on a user record synchronization server.

To backup the user record database:

1. Ensure the system is set up as a user record synchronization server. See System > Configuration > User Record Synchronization.

2. Select **Maintenance > Archiving > Archiving Servers**.

3. Select the **Archive User Record Synchronization Database** check box.

4. Specify an archive schedule. Through the options, schedule archives on any combination of weekdays including weekends.

   i NOTE: If you schedule an archival operation to occur during the hour that your system switches to Daylight Savings Time (DST) the operation may not occur as scheduled. For example, if your system is set to change to DST at 1:00 a.m. and you have scheduled an archival operation to occur at anytime between 1:01 a.m. and 1:59 a.m., the operation is not accomplished, because at 1:00 a.m. the system clock is moved forward to 2:00 a.m. and the system never reaches your archival time for that date.

5. Define a specific time when you want the system to archive data or elect to archive data every hour, which produces twenty-four files with unique timestamps.

   i NOTE: We recommend you schedule an archival operation during hours when traffic is light in order to minimize its impact to your users. The automatic

archiving process compresses files and, if the system is busy, can degrade performance for users. Also, a cluster node may appear unresponsive if the system is busy with traffic and performing archiving simultaneously.

6. Provide a password if you want to encrypt system configuration or user account archives with a password (optional).

7. Click **Save Changes**.

Related Documentation

- Network and Host Administration Overview

Using IKEv2 Security

This topic describes how to implement IKEv2 security. It includes the following information:

- IKEv2 Support Overview
- Configuring IKEv2 Ports
- IKEv2 Configuration Overview
- Enabling the IKEv2 Access Feature
- Defining the IKEv2 Role Mapping Rule

IKEv2 Support Overview

This topic gives an overview of support for IKEv2 security. It includes the following information:

- Understanding IKEv2
- Extensible Authentication Protocol
- Client Certificate-Based Authentication
- Client Requirements
- Supported Features

Understanding IKEv2

IKE or IKEv2 (Internet Key Exchange) is the protocol used to set up a security association in the IPsec protocol suite. Microsoft Windows 7 fully supports the IKEv2 standard through Microsoft's Agile VPN functionality and can operate with a VPN gateway using these protocols. Information on IKE and IKEv2 is widely available on the Internet. It is not the intent of this guide to describe details about IKE and IKEv2.

The system supports IKEv2, enabling interoperability with clients or devices, such as smartphones, that have a standards-based IPSec VPN client.

IKEv2 clients count toward the total number of sessions. Thus, the total number of sessions = number of IKEv2 sessions + number of NCP sessions.

The system supports the following methods for authenticating IKEv2 clients:

- Client certificate-based authentication
- Authentication using EAP methods

NOTE: IKEv2 uses port 500 exclusively. Do not configure port 500 in your VPN Tunneling profiles.

Extensible Authentication Protocol

EAP (Extensible Authentication Protocol) is an authentication framework frequently used in wireless communication. It provides functions and negotiation of authentication methods called EAP methods. Connect Secure supports the following EAP methods:

- EAP-MSCHAP-V2 (Microsoft Challenge-Handshake Authentication Protocol version 2)– a mutual authentication method that supports password-based user or computer authentication. During the EAP-MS-CHAP v2 authentication process, both the client and the authentication server must prove that they have knowledge of the user's password for authentication to succeed. Mutual authentication is provided by including an authenticator packet returned to the client after a successful server authentication. In Connect Secure, the local authentication server and the Active Directory server support EAP-MSCHAP-V2.

- EAP-MD5-Challenge – described in RFC 2284, enables an authentication server to authenticate a connection request by verifying an MD5 hash of a user's password. The server sends the client a random challenge value, and the client proves its identity by hashing the challenge and its password with MD5. EAP-MD5-Challenge is typically used on trusted networks where risk of packet sniffing or active attack are fairly low. Because of significant security vulnerabilities, EAP-MD5-Challenge is not usually used on public networks or wireless networks, because third parties can capture packets and apply dictionary attacks to identify password hashes. Because EAP-MD5-Challenge does not provide server authentication, it is vulnerable to spoofing (a third party advertising itself as an access point).

   Only the local authentication server is supported with EAP-MD5-Challenge.

IKEv2 provides a tunnel mechanism for EAP authentication; it does not perform authentication itself. Instead it proxies EAP messages from a client to the EAP server and back.

Client Certificate-Based Authentication

The system supports IKEv2 authentication using client certificates. Note that only certificate authentication server on Connect Secure supports client certification authentication of IKEv2 clients. When using client certificates for authentication, it is not necessary to configure the Realm/Protocol Set Mapping section under System > Configuration > IKEv2.

Client Requirements

Your IKEv2 client should support the following requirements to work with Connect Secure:

- Ability to establish IPsec Security Associations in Tunnel mode (RFC 4301).

- Ability to utilize the AES 128-bit encryption function (RFC 3602).

- Ability to utilize the SHA-1 hashing function (RFC 2404).

- Ability to utilize Diffie-Hellman Perfect Forward Secrecy in "Group 2" mode (RFC 2409).

- Ability to utilize IPSec Dead Peer Detection (RFC 3706).

- Ability to utilize the MD5 hashing function (RFC 1321).

- Ability to handle Internal Address on a Remote Network utilizing CFG_REQUEST-CFG_REPLY exchange.

Optional but recommended requirements include:

- Ability to adjust the Maximum Segment Size of TCP packets entering the VPN tunnel (RFC 4459).

- Ability to reset the "Don't Fragment" flag on packets (RFC 791).

- Ability to fragment IP packets prior to encryption (RFC 4459).

In addition, your client must support certificate authentication and ESP/SHA1.

Supported Features

The following features are unavailable to the end user since you are using a third-party client that are neither controlled nor configured by Pulse Secure.

- Host Checker

- Cache Cleaner

- Idle timeout notifications

- Upload Logs

- Route monitoring feature of split tunnel

- Windows interactive user logon options

- Session startup scripts

- NCP tunnel mode

- DNS search order

- Proxy server settings

Table 117 outlines the behavior of the Network Connect client and the IKEv2 client for certain split tunnel options.

Table119: Split Tunnel Operations with IKEv2 and Network Connect Clients

Option	IKEv2 Client	Network Connect Client
Disable split tunnel mode	Resource—through tunnel Internet—through tunnel local subnet(client)—through physical adapter	Resource—through tunnel Internet—through tunnel local subnet(client)—through tunnel
Enable split tunnel mode	Resource—through tunnel Internet—through tunnel but fails because the resource is not in split tunnel configuration. local subnet(client)—through physical adapt	Internet—through physical adapter local subnet(client)—through physical adapter
Allow local access subnet	Resource—through tunnel Internet—through tunnel local subnet(client)— through physical adapter (same as disable split tunnel mode)	Internet & other traffic—through tunnel local subnet(client)—through physical adapter
Enable split tunnel mode with route monitor (NC proprietary)	Resource—through tunnel Internet—through tunnel but fails because the resource is not in split tunnel configuration. local subnet(client)— through physical adapter Note: route table delete is not monitored.	Resource—through tunnel Internet—through physical adapter local subnet(client)—through physical adapter Note: route table delete is monitored

Enable ST with Allow local access subnet	Resource—through tunnel	Resource—through tunnel
	Internet —through tunnel but fails because the resource is not in split tunnel configuration.	Internet—through physical adapter
	local subnet(client)— through physical adapter	local subnet(client)—through physical adapter

Please note the following:

- IKEv2 does not support automatic cluster failover. After cluster failover, IKEv2 users must reconnect.

- For IKEv2 with client certification authentication to work with Windows 7 IKEv2 client, the certificate imported in to the system must have the enhanced key usage (EKU) value set to **serverAuth(1.3.6.1.5.5.7.3.1)**

Configuring IKEv2 Ports

To configure the IKEv2 ports and EAP protocol:

1. Select **System > Configuration > IKEv2** to display the configuration page.

2. Enter the DPD timeout value in seconds. Valid values are 400-3600.

   DPD is a form of keepalive. When a tunnel is established but idle, one or both sides may send a "hello" message and the other replies with an acknowledgement. If no response is received, this continues until the DPD time value has elapsed. If there still isn't any traffic or acknowledgement, the peer is determined to be dead and the tunnel is closed.

3. Under Port/Realm Mapping, select the port and the realm to use that port.

   To add additional port/realm mapping sets, click **Add**.

   To delete a port/realm mapping set, select the check box next to the set to remove and click **Delete**.

4. Under Realm / Protocol Set Mapping, select the realm and the EAP protocol set to use for that realm.

   To add additional realm/protocol mapping sets, click **Add**.

   To delete a realm/protocol mapping set, select the check box next to the set to remove and click **Delete**.

5. Click Save Changes.

   ⓘ NOTE: Changing IKEv2 configuration (**System > Configuration > IKEv2**) disconnects connections from IKEv2 clients, VPN Tunneling and Pulse. VPN Tunneling and Pulse will reconnect automatically.

IKEv2 Configuration Overview

IKEv2 EAP supports the following authentication server types:

- Local authentication

- Active Directory

If you are using IKEv2 EAP authentication on a local authentication server, you must select the **Password stored as clear text** check box in the Auth Server page of the admin console. Note that you cannot edit an existing local authentication server instance to select this option. If you require IKEv2 EAP authentication on a local authentication server, you must create a new local authentication server instance.

> ℹ️ NOTE: IKEv2 EAP does not work with any preexisting local authentication servers since they do not store passwords in clear text.

To configure support for IKEv2:

1. Configure your client for using IKE. For more information, see your mobile device's documentation.

2. Install client and device certificates.

   - You need a Certificate Authority (CA) that can issue client certificates.

   - On the client side, install this client certificate along with the CA certificate.

   - On the Connect Secure server side, install the CA certificate under Configuration/Certificates/Trusted Client CAs.

   - On the client side, install the Connect Secure certificate corresponding to the port to which the client connects, found under Configuration/Certificates/Device Certificates.

3. Define an IKEv2 rule under the Users > User Realms > User > Role Mapping page of the admin console.

4. Select the IKEv2 access feature under the Users > User Roles > User > General > Overview page of the admin console.

5. Enable Network Connect for the Role and configure an NC Connection Profile (IP pool) to use for that Role.

When a client uses IKEv2 to connect to the host, the Agent Type column of the Active Users page displays IKEv2.

Enabling the IKEv2 Access Feature

Roles specify the session properties, including enabled access features, for users who are mapped to the role.

To enable the IKEv2 access feature:

1. Select **User > User Roles >** *Role Name* **> General > Overview**.

2. Under Access Features, check the **VPN Tunneling** check box.

3. Click **Save Changes**.

Defining the IKEv2 Role Mapping Rule

Role mapping rules are conditions a user must meet for the system to map the user to one or more user roles.

> ℹ️ NOTE: The procedure described in this topic is required only if you want to create a separate role mapping rule specific for IKEv2 users. If you use regular username, group or custom expression-based role mapping rules (typically used for general access to a device), the following procedure is not required.

1. Select **User > User Realms > User > Role Mapping**.

2. Click **New Rule**.

3. Select Custom Expressions as the type of condition on which to base the rule.

4. Click **Update** to display the Expressions list.

5. Click the Expressions button to display the Expressions tab of the server catalog.

6. Create a rule: **userAgent = "IKEv2"**.

7. Click **Add Expression** and then **Close**.

8. Select the rule you just created from the Available Expressions list and click **Add** to move it to the Selected Expressions list.

9. Specify the roles to assign to the authenticated user by adding roles to the Selected Roles list.

10. (optional) Check the Stop processing rules when this rule matches check box if you want the system to stop evaluating role mapping rules when the user meets the conditions specified for this role.

11. Click **Save Changes**.

Related Documentation

- Network and Host Administration Overview

Using the Traffic Segregation Feature

This topic describes the traffic segregation feature that is available on the Connect Secure virtual appliance (service provider edition). It includes the following information:

- Traffic Segregation Feature Overview
- Configuring AAA Traffic Over the Management Port
- Configuring AAA Traffic Through Both the Internal and Management Ports

Traffic Segregation Feature Overview

Service providers often need a way to cleanly segregate the system-generated network traffic across interfaces (such as Internal, Management and VLAN ports), to differentiate AAA traffic from others.

Two typical service provider deployment models are:

- Authentication server of the customer is hosted by the customer
- Authentication server for the customer is hosted and managed by the service provider

In both models, the service provider's authentication server is always hosted in the service provider's network and is reachable either through the internal or management port. In the first model, the customer's authentication servers are reachable through the internal port of the virtual appliance. In the second model, the customer's authentication server must be routed either through the internal or management port, depending on where the service provider has hosted the customer's authentication server.

A Traffic Segregation menu is available only on virtual appliances to allow system providers to configure the interface and traffic. The "Default Network" is used as the primary logical network for the service provider and customer configuration. If traffic segregation across different logical networks is needed, the "Administrative Network" can be used.

You can differentiate AAA traffic from other traffic and route it through the management port. This is useful when the only the authentication servers are hosted on the network reachable through the management port and all other resources uses a different port. This option is available on both the Default Network and the Administrative Network.

The configurations to do on the virtual appliance depend on the logical network setup around the virtual appliance and the type of service provider deployment model:

- If both the service provider's and customer's authentication server are reachable through the same interface, the entire configuration for the service provider and customer is done under the Default Network. It is not necessary to enable the Administrative Network.

- If the service provider's and customer's authentication servers are located on two different networks, the Administrative Network must be created. Table 120 shows where the administrator configures the options in the system GUI.

Table120: Configuring Traffic Segregation Options

Task	For Service Provider	For Customer
AAA configuration for administrators	Under Administrative Network	Delegated admin under Default Network (only if needed)
AAA configuration for users	Not required	Under Default Network
System configuration	Under Default Network	Not required

Configuring AAA Traffic Over the Management Port

This topic describes how to configure all AAA traffic through the management port. Note that if the management port is enabled, system management traffic (like SNMP) continues to flow through the management port. All other traffic flows through the internal port. When all authentication servers are hosted in the same network, you need only to configure the Default Network.

To configure AAA traffic over the management port:

1. Select **System > Traffic Segregation**.

2. Click **Default Network**.

3. Click the **Send AAA Traffic via Management Port** check box.

4. Click **Save Changes**.

Configuring AAA Traffic Through Both the Internal and Management Ports

When authentication servers are hosted in two different networks, enable the Administrative Network to allow the servers to be reachable through two different virtual appliance interfaces. Service provider authentication server settings are defined under the Administrative Network. User authentication servers are configured in the Default Network using the internal port.

To enable the Administrative Network and set the Management Port as the default:

1. Select **System > Traffic Segregation**.

2. Check the **Enable Administrative Network** check box.

3. Click **Save Changes**.

   The Administrative Network appears in the Traffic Segregation table.

4. Click **Administrative Network**.

5. Under Interface: Available Interface, select **Management Port** and click **Add**.

6. Under Selected Interfaces, select **Management Port** and click **Set default interface**.

7. Click **Save Changes**.

To configure service provider AAA settings on the Administrative Network:

1. Select **Administrative Network** from the menu in the upper left corner and click **Go**.

   A limited set of options are available on the Administrative Network.

2. Select **Authentication > Auth Servers** and configure your service provider AAA configurations as needed.

By default, the Default Network uses the internal port so no additional configuration is needed for the Default Network.

To log in to the virtual appliance through the administrative network, either:

- Use the internal or external port (https://*virtual-appliance*/VA/adminnet/admin)

- Assign a virtual port to the administrative network and log in through that virtual port (https://*virtual-appliance-VIP*/admin)

  NOTE: If you enable AAA traffic over the management port, the next time you log in to the virtual appliance over the management port the administrator sign-in page appears with "Welcome to Pulse Connect Secure" and not "Sign in via Administrative Network to Pulse Connect Secure."

Related Documentation

- *Service Provider Virtual Appliance Deployment Guide*

- Network and Host Administration Overview

- Using the Management Port

Using the Serial Port

This topic describes use of the serial port and serial port console. It includes the following information:

- Connecting to the Serial Port Console

- Using the Serial Console to Roll Back to a Previous OS Version

- Using the Serial Console to Reset the System to the Factory Image

Connecting to the Serial Port Console

In cases where the admin console is unavailable, you can perform network and host configuration tasks and troubleshooting using the serial port console.

To connect to the serial console:

1. Plug a null modem crossover cable from a console terminal or laptop into the device serial port. This cable is provided in the product box. Do not use a straight serial cable.

2. Configure a terminal emulation utility, such as HyperTerminal, with the following serial connection parameters:

   - 9600 bits per second

   - 8-bit No Parity (8N1)

   - 1 Stop Bit

   - No flow control

3. Press Enter until the serial console is displayed.

Figure 163 shows the serial console menu.

Table 121 describes the serial console menu options.

Figure166: Serial Console Menu
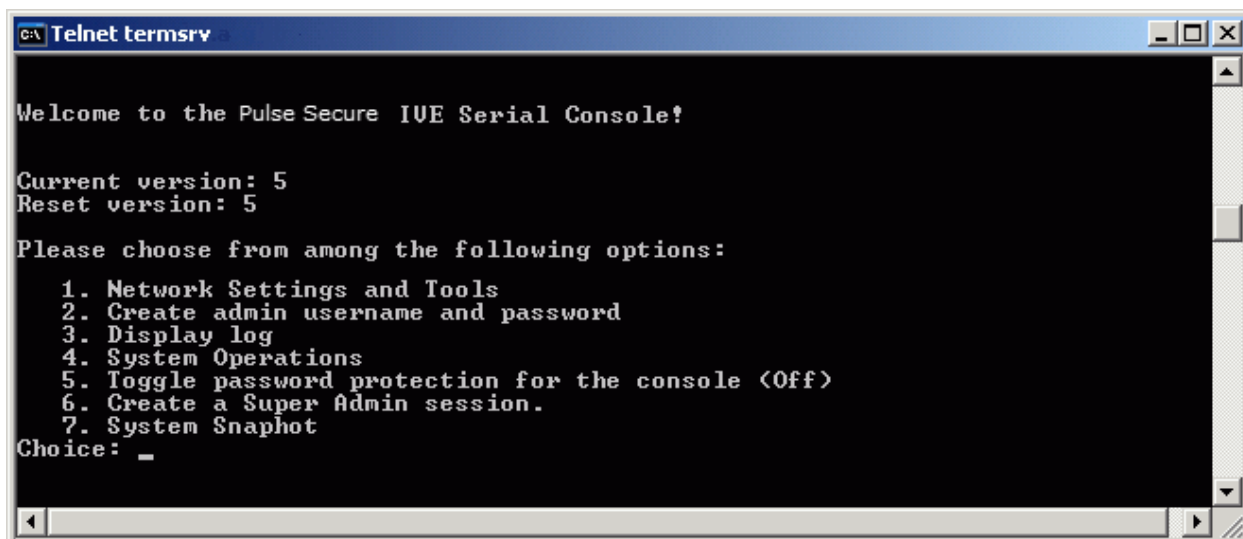
Table121: Serial Console Menu Options

Options	Description
1. Network Settings and Tools	Enables you to change standard network settings; print a routing table; print or clear an ARP cache; run the ping and traceroute commands, remove static routes, and add an ARP entry.
2. Create admin username and password	Enables you to create a new super administrator account.
3. Display log	Enables you to display system configuration, user access logs, or administrator access logs through the serial console. Note that must enter **q** to return to serial console options after viewing the logs.
4. System Operations	Enables you to reboot, shut down, restart, roll back, or factory reset the system without using the admin console.
5. Toggle password protection for the console	Enables you to password protect the serial console. When you toggle this option to "on," only super administrators are allowed access.
6. Create a Super Admin session	Enables you to create a recovery session to the admin console, even if you have configured the system to block access to all administrators. When you select this option, the system generates a temporary token that is valid for 3 minutes. Enter the following URL into a browser window: `https://<fully-qualified-domain-name>/dana-na/auth/recover.cgi` Then, enter the temporary token when prompted to sign in to the admin console. When you select this option, the system blocks any additional administrators from signing in to the admin console until you sign in to the specified URL and initiate a session using your token. The appliance blocks additional sign-in attempts so that you can fix any configuration problems that the system may have encountered without conflicting with another session.
7. System Snapshot	Enables you to take a system snapshot without using the admin console. When you select this option, the system takes the snapshot immediately. You can then send the snapshot file, by way of SCP, to a remote system. The system prompts you for the destination server port, user ID, password, and the destination path to the remote directory. If you choose not to send the snapshot file to a remote system, the system saves the file locally. The next time you log in to the admin console, the System Snapshot tab contains a link to the snapshot file.

Using the Serial Console to Roll Back to a Previous OS Version

You can use the admin console to roll back the configuration to a previous state. If the rollback option is not available in the admin console, you can use the procedure described in this section to perform the system rollback.

If you have not yet performed an OS service package upgrade, there is no previous state to roll back to, and the rollback option is not available. If you have performed an OS service package upgrade, any system and user configuration data created after the upgrade is lost unless you export the most current configuration files before rolling back the system and then import them afterwards.

To roll back to the previous OS service package:

1. Connect to the serial console.

2. In a browser window, sign in to the admin console.

3. Select **Maintenance > System > Platform**.

4. Click **Reboot Now** and then return to the console utility window. The window displays a message that the system is restarting.

5. After several moments, you are prompted to use the Tab key to select options. Press Tab, and when prompted for the configuration to load, type **rollback** and then press Enter.

After you click **Reboot Now,** the rollback status is output to the screen, and when complete, you are prompted to press Return (Enter) to modify system settings, which returns you to the initial setup options. When you are finished entering data, simply close the serial console window.

If you wait more than 5 seconds to enter your choice, the current system configuration is automatically loaded and you must go back to the admin console and click **Reboot Now** to start the process again. If you have already performed a system rollback, the rollback option is not available again until you upgrade the OS service package again.

Using the Serial Console to Reset the System to the Factory Image

In rare cases, you might need to reset the system to its original factory settings. Before performing this advanced system recovery option, contact Pulse Customer Support (http://www.juniper.net/support/). If possible, export the most current system and user configuration data before performing a factory reset.

To perform a factory reset:

1. Connect to the serial console.

2. In a browser window, sign in to the admin console.

3. Select **Maintenance > System > Platform**.

4. Click **Reboot** and then go back to the console utility window. The window displays a message that the system is restarting.

5. After several moments, you are prompted to use the Tab key to select options. Press Tab , and when prompted for the configuration to load, type **factory-reset** and then press Enter.

   If you wait more than 5 seconds to enter your choice, the current system configuration is automatically loaded, and you must go back to the admin console and click **Reboot Now** to start the process again.

6. When you are prompted to confirm performing a factory reset, type **proceed** and then press Enter.

The system begins the process of resetting the machine to its original settings and outputs several screens of data. After several minutes, you are prompted to use the Tab key to select configuration choices.

7. When prompted to press the Tab key, do one of the following:

- Wait for the default selection (current) to start automatically.

- Press **Tab**, type **current**, and then press **Enter**.

You are then prompted to enter the initial configuration settings. For details on how to proceed, see the Installation Guide provided in the product packaging or on the Pulse Secure Support site.

After you complete the initialization process, you can upgrade to the latest OS service package and import saved system and user configuration files to return to the last good working state of your system.

You might receive errors from the system during the initial setup or on a factory reset. Before the system starts services, it monitors the network port for a maximum of 120 seconds. The system checks the link status and sends ARP requests to the default gateway. If there is a problem, after 5 seconds, the system displays a message on the serial console that starts with **NIC:......** If the link recovers within 120 seconds, the startup process continues. If the link does not recover, the following message is displayed:

```
Internal NIC: ...............[Down code=0x1]
```

Two codes can appear:

- **0x1** means that the interface link status reported by the NIC remains off (for example, a disconnected cable or a cable is in the wrong port).

- **0x2** means that the gateway is unreachable. The system boots but is not reachable from IP addresses bound to that network port.

Related Documentation

- Network and Host Administration Overview

- Using the Admin Console Troubleshooting Tools

CHAPTER 33 Certificate Security Administration

Understanding Digital Certificate Security

Policy Secure and Connect Secure use Public Key Infrastructure (PKI) to secure the data sent to clients over the Internet. PKI is a security method that uses public and private keys to encrypt and decrypt information. These keys are enabled and stored through digital certificates. A digital certificate is an encrypted electronic file issued by a certificate authority (CA) that establishes credentials for client/server transactions.

In public key cryptography, a public/private key pair is used to encrypt and decrypt data. Data encrypted with a public key, which the owner makes available to the public, can be decrypted with the corresponding private key only, which the owner keeps secret and protected. For example, if User1 wants to send User2 an encrypted message, User1 can encrypt it with User2's public key and send it. User2 then decrypts the message with the private key. The reverse process is also useful: encrypting data with a private key and decrypting it with the corresponding public key. This process is known as creating a digital signature. For example, if User1 wants to present their identity as the sender of a message, they can encrypt the message with her private key and send the message to User2. User2 then decrypts the message with User1's public key, thus verifying that User1 is indeed the sender.

Policy Secure and Connect Secure systems use the following types of digital certificates to establish credentials and secure session transactions:

- Device certificates—A device certificate helps to secure network traffic to and from the Pulse Secure client service using elements such as company name, a copy of your company's public key, the digital signature of the CA that issued the certificate, a serial number, and expiration date. In addition, Policy Secure uses a device certificate for communications with the ScreenOS Enforcer and the Pulse Policy Secure.

- Trusted client CAs—A trusted client CA is a client-side certificate issued by a CA. You can use trusted client CAs in the access management framework realm and role configurations to require certificates or certificates with specific attributes. For example, you may specify that users must present a valid client-side certificate with the OU attribute set to "yourcompany.com" to sign into the Users authentication realm.

- Trusted server CAs—A trusted server CA is the certificate of a Web server that you trust. You can install a trusted server CA to validate the credentials of the web sites that users access through the Pulse Secure client service.

- Code-signing certificates—Connect Secure only. A code-signing certificate (also called an applet certificate) is a type of server-side certificate that re-signs Java applets that are intermediated by Connect Secure. You can use the self-signed code-signing certificate that comes pre-loaded, or you can install your own code-signing certificate.

- Client auth certificates—Connect Secure only. In this context, the client auth certificate is used when backend SSL servers require Connect Secure to present a client certificate.

  NOTE:

  - The system can verify certificates that use SHA2 as the message digest.

  - DSA certificates are not supported.

Related Documentation

- Using Device Certificates
- Using Trusted Client CAs
- Using Trusted Server CAs
- Understanding ECC Certificates
- *Using Certificate-Based Security with Infranet Enforcer Deployments*

Using Device Certificates

This topic describes how to use device certificates. It includes the following information:

- Understanding Device Certificates
- Understanding Self-Signed Certificates
- Importing a Device Certificate and Private Key
- Creating a Certificate Signing Request
- Importing a Signed Certificate Created from a CSR
- Understanding Intermediate Certificates
- Importing Intermediate CA Certificates
- Importing a Renewed Certificate That Uses the Existing Private Key
- Downloading a Device Certificate
- Using Device Certificates with Virtual Ports

Understanding Device Certificates

A device certificate helps to secure network traffic to and from the Pulse Secure client service using elements such as your company name, a copy of your company's public key, the digital signature of the Certificate Authority (CA) that issued the certificate, a serial number, and an expiration date. The system also uses device certificates for secure communications with the Infranet Enforcer.

When receiving encrypted data from the system, the client's browser first verifies whether the device certificate is valid and whether the user trusts the CA that issued the certificate. If the user has not already indicated that they trust the certificate issuer, the Web browser prompts the user to accept or install the certificate.

The system supports X.509 device certificates in DER and PEM encode formats (file extensions include .cer, .crt, .der, and .pem) as well as PKCS #12 (file extensions include .pfx and .p12). The system also supports the following features:

- Intermediate device CA certificates—Within a certificate hierarchy, one or more intermediate certificates are branched off a single root certificate.

- Multiple device certificates—When using multiple device certificates, each certificate handles validation for a separate hostname or fully qualified domain name (FQDN) and can be issued by a different CA.
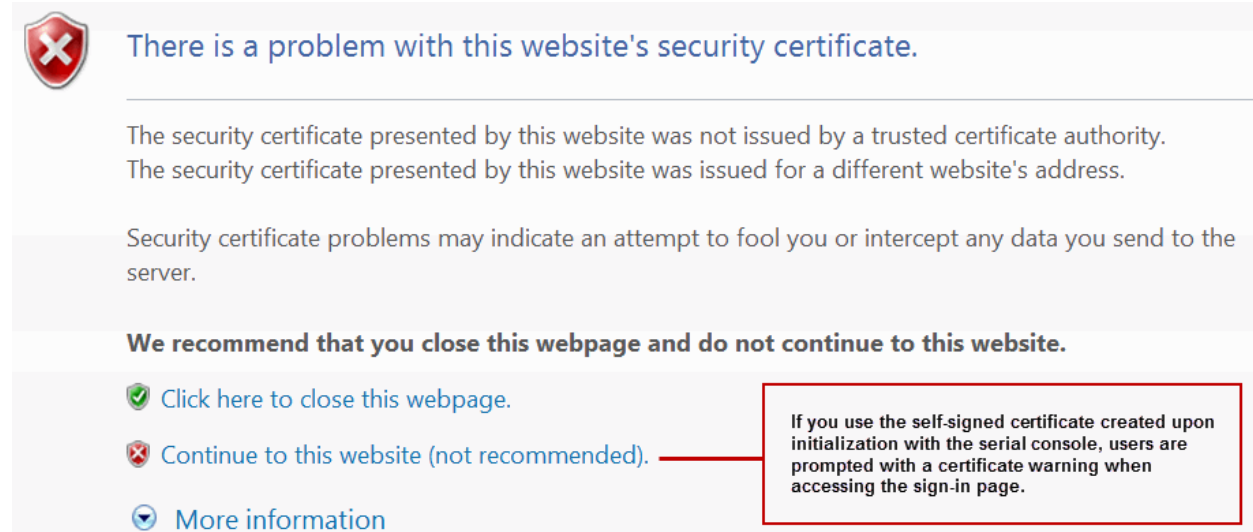
  i NOTE: Beginning with Connect Secure system software release 7.2, you can assign device certificates to the Connect Secure VLAN interfaces.

Understanding Self-Signed Certificates

When you initialize the system with the serial console, the system creates a self-signed certificate that enables you to immediately begin setting up the system. Users are prompted with a security alert each time they sign in because the certificate is not issued by a trusted CA. Figure 164 shows the security alert.

Figure167: Security Alert When the Device Certificate Is Not Issued by a Trusted CA

Before promoting the system to production use, we recommend you replace the self-signed certificate with a certificate issued by a trusted CA.

i NOTE: In Policy Secure deployments with ScreenOS Enforcers, you must use a CA-signed device certificate. If you use a self-signed certificate, the ScreenOS Enforcer does not allow a connection. Import a CA-signed device certificate into the Policy Secure , and then import the certificate of the CA that signed the device certificate into the ScreenOS Enforcer.

Importing a Device Certificate and Private Key

The system uses certificates to verify itself to other network devices. A digital certificate is an electronic means of verifying your identity through a trusted third party, known as a Certificate Authority (CA). Your company might use its own enterprise CA server, or it might use a reputable third-party CA.

To import an enterprise root server certificate and private key:

1. Select **System > Configuration > Certificates > Device Certificates**.

2. Click **Import Certificate & Key** to display the configuration page.

3. Use one of the following options to complete the import procedure:

   - **If certificate file includes private key**—When the certificate and key are contained in one file.

   - **If certificate and private key are separate files**—When the certificate and key are in separate files.

   - **Import via System Configuration file**—When the certificate and key are contained in a system configuration file. With this option, the system imports all of the certificates specified (including private keys and pending CSRs, but not the corresponding port mappings).

   In the appropriate form, browse to the certificate and key files. If the file is encrypted, enter the password key.

4. Click **Import**.

   NOTE: The **Import Certificate and Key** button is disabled on FIPS hardware platforms because importing private keys is not allowed. On a FIPS hardware platform, you must create a CSR and then import a signed certificate from the CA.

Creating a Certificate Signing Request

If your company does not own a digital certificate for its Web servers, you can create a certificate signing request (CSR) and then send the request to a CA for processing. When you create a CSR, a private key is created locally that corresponds to the CSR. If you delete the CSR at any point, this file is also deleted, prohibiting you from installing a signed certificate generated from the CSR.

To create a certificate signing request:

1. Select **System > Configuration > Certificates > Device Certificates**.

2. Click **New CSR** to display the configuration page.

3. Complete the required information and click **Create CSR**.

4. Follow the onscreen instructions, which explain what information to send to the CA and how to send it.

   When you submit a CSR to a CA authority, you might be asked to specify either the type of Web server on which the certificate was created or the type of Web server the certificate is for. Select **apache** (if more than one option with apache is available, select **any**). If you are prompted for the certificate format to download, select the standard format.

   Do not send more than one CSR to a CA at one time. Doing so can result in duplicate charges.

   NOTE: To view details of any pending requests that you previously submitted, click the **Certificate Signing Request Details** link.

Importing a Signed Certificate Created from a CSR

When you receive the signed certificate from the CA, import it.

To import a signed device certificate created from a CSR:

1. Select **System > Configuration > Certificates > Device Certificates**.

2. Under Certificate Signing Requests, click the **Pending CSR** link that corresponds to the signed certificate.

3. Under Import signed certificate, browse and select the certificate file you received from the CA, and then click **Import**.

Understanding Intermediate Certificates

Within a certificate hierarchy, one or more intermediate certificates are branched off a single root certificate. The root certificate is issued by a root CA and is self-signed. Each intermediate certificate is issued by the certificate preceding it in the chain.

To use chained certificates in your deployment, you must ensure that the server (Policy Secure or Connect Secure) and client (Web browser) together contain the entire certificate chain. For example, you can secure traffic using a chain that stems from a VeriSign root certificate. If your users' browsers come preloaded with VeriSign root certificates, you need to install only the lower-level certificates in the chain. When your users sign in, the system presents any required certificates within the chain to the browser to secure the transaction. The system creates the proper links in the chain using the root certificate's IssuerDN. If the system and browser together do not contain the entire chain, the user's browser does not recognize or trust the device certificate because it is issued by another certificate instead of by a trusted CA.

You can upload one or more intermediate CAs in a PEM file. The entire chain must be sent to the client in descending order, starting with the root certificate.

Within a certificate hierarchy, one or more intermediate certificates are branched off a single root certificate. The root certificate is issued by a root CA and is self-signed. Each intermediate certificate is issued by the certificate preceding it in the chain.

To use chained certificates in your deployment, you must install the appropriate client-side certificates in each user's Web browser and then upload the corresponding CA certificates to Pulse Secure gateway's Intermediate CA store. Use one of the following methods to upload the certificate chain:

- Import the entire certificate chain in one file. The file must contain the root certificate and any subcertificates whose parents are in the file or already imported. You can include certificates in any order in the import file.

- Import the certificates one at a time in descending order. You must install the root certificate first, and then install the remaining chained certificates in descending order.

If you follow one of these methods, the system automatically chains the certificates together in the correct order and displays them hierarchically in the admin console.

> NOTE: If you install multiple certificates in a user's Web browser, the browser prompts the user to choose which certificate to use when signing in.

Importing Intermediate CA Certificates

To import an intermediate CA certificate:

1. Select **System > Configuration > Certificates > Device Certificates**.

2. Click the **Intermediate Device CAs** link to display the management page.

3. Click **Import CA certificate**.

4. Browse to the certificate file, select it, and click **Import Certificate** to complete the import operation.

Importing a Renewed Certificate That Uses the Existing Private Key

You can renew a device certificate in two ways:

- Submit a new CSR to a CA—This process is more secure because the CA generates a new certificate and private key and retires the older private key. To use this renewal method, you must first create a CSR through the admin console.

- Request renewal based on the CSR previously submitted to the CA—This process is less secure, because the CA generates a certificate that uses the existing private key.

When you order a renewed certificate, you must either resubmit your original CSR or ensure that the CA has a record of the CSR that you submitted for your current certificate.

To import a renewed device certificate that uses the existing private key:

1. Follow your CA's instructions for renewing a certificate that you previously purchased through them. Be sure to specify the same information you used in the original CSR. Your CA uses this information to create a new certificate that corresponds to the existing key.

   > **ⓘ** NOTE: Even though you specify the same information used in the original CSR, your root CA might have different serial numbers and keys from the original. You might need to support both new client and old client certificates during the transition period, which also requires that you maintain two root CA certificates (your existing certificate and the renewed certificate), at least temporarily

2. Select **System > Configuration > Certificates > Device Certificates**.

3. Click the link that corresponds to the certificate you want to renew.

4. Click **Renew Certificate** to display the page.

5. In the Renew the Certificate form, browse to the renewed certificate file, enter the password for the certificate key, and click **Import**.

Downloading a Device Certificate

You download the device certificate to your local host so that you can import it into other network devices as needed.

To download a device certificate:

1. Select **System > Configuration > Certificates > Device Certificates**.

2. Click the link of the device certificate you want to download to display the configuration page.

3. Click the **Download** link.

4. Save the file to the desired location.

Using Device Certificates with Virtual Ports

Virtual ports can be used to create multiple fully qualified domain names for user sign-in.

When a user tries to sign in using the IP address defined in a virtual port, the system uses the certificate associated with the virtual port to initiate the SSL transaction and for NetScreen Address Change Notification (NACN) communications with the Infranet Enforcer.

You must associate the signed certificate with the port that is connected to the Infranet Enforcer. You can use the same port and certificate for OAC or Pulse. Or, you can import other signed certificates and associate them with ports connected to OAC.

You can implement digital certificate security with virtual ports in either of the following ways:

- Associate all hostnames with a single certificate—With this method, you use a single wildcard certificate to validate the identity of all system hostnames, regardless of which hostname is used to sign into. A wildcard certificate includes a variable element in the domain name, making it possible for users who sign in from multiple hosts to map to the "same" domain. For example, if you create a wildcard certificate for *.yourcompany.com, the system uses the same certificate to validate its identity to users who sign in to employees.yourcompany.com as it does to users who sign into partners.yourcompany.com.

- Associate each hostname with its own certificate—With this method, you associate different hostnames with different certificates. Create a virtual port for each hostname. A virtual port activates an IP alias on a physical port. For example, you can create two virtual ports on a single appliance, mapping the first virtual port to the IP address 10.10.10.1 (sales.yourcompany.com) and the second virtual port to the IP address 10.10.10.2 (partners.yourcompany.com). Then you can associate each of these virtual ports with its own certificate, ensuring that users authenticate through different certificates.

To associate certificates with virtual ports:

1. Create the virtual ports.

2. Import the device certificates.

3. Associate the device certificates with the virtual ports:

   1. Select **System > Configuration > Certificates > Device Certificates**.

   2. Click the link of the device certificate you want to configure to display the configuration page.

   3. Use the controls in the "Present certificate on these ports" section to associate ports with the certificate.

   > *i* NOTE: You can assign only one device certificate to the Management Port. If you assign a certificate other than the default device certificate to the Management Port, the default device certificate is automatically deselected as the default. If you do not select a device certificate for the Management Port, the system uses the default device certificate that is presented on the Internal port. You cannot assign certificates to Management Port VIPs.

Related Documentation

- Understanding Digital Certificate Security

- Understanding ECC Certificates

Using Trusted Client CAs

This topic describes how to use trusted client Certificate Authorities (CAs). It includes the following information:

- Understanding Trusted Client CAs

- Trusted Client CA Implementation Notes

- Understanding CRLs

- Understanding OCSP

Understanding Trusted Client CAs

A trusted client CA is a CA that you deem trusted by adding it the trusted client CA store. The system trusts any certificate issued by that CA. To use client CA certificates, you must install and enable the proper certificates. Additionally, you must install the corresponding client-side certificates in your users' Web browsers, or you must use the MMC snap-in in your users' computer accounts (machine certificate). When validating a client-side CA certificate, the system verifies that the certificate is not expired or corrupt and that the certificate is signed by a CA that the system has been configured to recognize. If the CA certificate is chained, the system also follows the chain of issuers until it reaches the root CA, validating each issuer in turn. The system supports X.509 CA certificates in DER and PEM encode formats.

When you install a client-side certificate, you must determine whether to use the certificate to identify individual users or individual machines. To use the certificate to identify individual users, you must install the certificate in each user's individual certificate store. Then you must enable authentication using a certificate server, or you must enable authorization using realm, role, and/or resource policy settings. To use the certificate to identify individual machines, you must install the certificate in each computer's certificate store. Then you must configure a Host Checker policy that checks for the machine certificate and authorizes access to realms, roles, or resource policies based on the certificate's validity.

The system supports using the following additional features with CA certificates:

- Certificate servers—A certificate server is a type of local authentication server that allows you to authenticate users based solely on their certificate attributes rather than authenticating them against a standard authentication server (such as LDAP or RADIUS), and it requires specific certificates or certificate attributes.

- Certificate hierarchies—Within a certificate hierarchy, one or more subordinate certificates (called intermediate certificates) are branched off a root certificate to create a certificate chain. Each intermediate certificate (also called a chained certificate) handles requests for a part of the root CA domain. For example, you can create a root certificate that handles all requests to the yourcompany.com domain and then branch off intermediate certificates that handle requests to partners.yourcompany.com and employees.yourcompany.com. When you install a chained certificate, the system confirms that the chain is valid and allows users to authenticate using the leaf certificate (that is, the lowest certificate in the chain).

- Certificate revocation lists—Certificate revocation is a mechanism by which a CA invalidates a certificate before its expiration date. The CA publishes a certificate revocation list (CRL) which is a list of revoked certificates. Within CRLs, each entry contains the serial number of the revoked certificate, the date that the certificate was revoked, and the reason the certificate was revoked. The CA can invalidate a certificate for various reasons such as when the employee to whom the certificate is issued leaves the company, the certificate's private key is compromised, or the client-side certificate is lost or stolen. When the CA revokes a certificate, the system can appropriately deny access to users who present a revoked certificate.

Trusted Client CA Implementation Notes

Uploading a trusted client CA certificate does not enable client-side SSL authentication or authorization. To do so, you must use a certificate server, or enable certificate restrictions at the realm, role, or resource policy level, or create a Host Checker policy that verifies a machine certificate.

With client-side certificates, we strongly recommend that you advise users to close their Web browsers after signing out. If they do not, other users might be able to use their open browser sessions to access certificate-protected resources without reauthentication. After loading a client-side certificate, Internet Explorer caches the certificate's credentials and private key. The browser keeps this information cached until the user closes the browser (or, in some cases, until the user reboots the workstation). For details, see http://support.microsoft.com/?kbid=290345.) To remind users to close their browsers, you can modify the sign out message on the Sign-in Pages tab.

> NOTE: Certificate authentication does not work on Internet Explorer 8 and 9 if SSL 2.0 is enabled with other SSL and TLS versions. For details, see http://support.microsoft.com/kb/2851628.

Understanding CRLs

A certificate revocation list (CRL) is a mechanism for canceling a client-side certificate. As the name implies, a CRL is a list of revoked certificates published by a CA or a delegated CRL issuer. The system supports base CRLs, which include all of the company's revoked certificates in a single, unified list.

The system determines the correct CRL to use by checking the client's certificate. (When it issues a certificate, the CA includes CRL information for the certificate in the certificate itself.) To ensure that it receives the most up-to-date CRL information, the system periodically contacts a CRL distribution point to get an updated list of CRLs. A CRL distribution point (CDP) is a location on an LDAP directory server or Web server where a CA publishes CRLs. The system downloads CRL information from the CDP at the interval specified in the CRL, at the interval that you specify during CRL configuration, and when you manually download the CRL. The system also supports CRL partitioning. CRL partitioning enables you to verify portions of very large CRLs without spending the time and bandwidth necessary to access and validate a very large CRL or collection of large CRLs. CRL partitioning is only enabled when you employ the Specify the CDP(s) in the client certificates method (described below). In this case, the system validates the user by verifying only the CRL specified in the client certificate.

Although CAs include CRL information in client-side certificates, they do not always include CDP information as well. A CA can use any of the following methods to notify the system of a certificate's CDP location:

- Specify the CDP(s) in the CA certificate—When the CA issues a CA certificate, it might include an attribute specifying the location of the CDPs that the system should contact. If more than one CDP is specified, the system chooses the first one listed in the certificate and then fails over to subsequent CDPs, if necessary.

- Specify the CDP(s) in the client certificates—When the CA issues a client-side certificate, it might include an attribute specifying the location of the CDPs that the system must contact. If more than one CDP is specified, it chooses the first one listed in the certificate and then fails over to subsequent CDPs, if necessary. When the system employs CRL partitioning and the client certificate specifies only one CRL, it performs verification using only that CRL.

> **NOTE:** If you choose this method, the user receives an error on the first sign-in attempt because no CRL information is available. Once the system recognizes the client's certificate and extracts the CRL location, it can start downloading the CRL and subsequently validate the user's certificate. In order to successfully sign in, the user must try to reconnect after a few seconds.

- Require the administrator to manually enter the CDP location—If the CA does not include the CDP location in the client or CA certificates, you must manually specify how to download the entire CRL object. You can specify a primary and backup CDP. (Manually entering the CDP location provides the greatest flexibility because you do not need to reissue certificates if you change the CDP location.)

The system compares the user's certificate against the appropriate CRL during authentication. If it determines that the user's certificate is valid, the system caches the certificate attributes and applies them, if necessary, during role and resource policy checks. If it determines that the user's certificate is invalid, if it cannot contact the appropriate CRL, or if the CRL is expired, it denies the user access.

> **NOTE:**
> - The system supports only CRLs that are in a PEM or DER format and that are signed by the CA for which the revocations apply.
> - The system only saves the first CRL in a PEM file.

Understanding OCSP

The Online Certification Status Protocol (OCSP) is a service that enables you to verify client certificates. When OCSP is enabled, the system becomes a client of an OCSP responder and forwards validation requests for users based on client certificates. The OCSP responder maintains a store of CA-published certificate revocation lists (CRLs) and maintains an up-to-date list of valid and invalid client certificates. After the OCSP responder receives a validation request, it validates the status of the certificate using its own authentication database, or it calls upon the OCSP responder that originally issued the certificate to validate the request. After formulating a response, the OCSP responder returns the signed response, and the original certificate is either approved or rejected.

Importing a Trusted Client CA Certificate

If you require users to provide a client-side certificate to sign in, you must upload the corresponding CA certificate. You can upload CA certificates manually, or you can configure the system to upload CA certificates automatically. The system uses the uploaded certificate to verify that the browser-submitted certificate is valid. In addition, you can specify whether or not to automatically import CA certificates for validation, and you can specify a CRL or OCSP retrieval method to use to automatically import CA certificates.

To import a trusted client CA certificate:

1. Select **System > Configuration > Certificates > Trusted Client CAs** to display the configuration page.

2. Click **Import CA Certificate** to display the configuration page.

3. Browse to the certificate file, select it, and click **Import Certificate** to complete the import operation.

Renewing a Certificate

To renew a certificate:

1. Select **System > Configuration > Certificates > Trusted Client CAs**.

2. Click the link for the certificate you want to renew.

3. Click **Renew Certificate** to display the import certificate page.

4. Browse to the certificate file, select it, and click **Import Certificate** to complete the import operation.

Configuring Auto-Importing of Client Certificates

To enable auto-importing:

1. Select **System > Configuration > Certificates > Trusted Client CAs**.

2. Click the **Auto-Import Options** button to display the options.

3. Complete the configuration described in Table 122.

4. Save your changes.

Table122: Auto-Import Options Settings

Settings	Guidelines
Auto-import trusted CAs	Select this option to enable auto-import and display its configuration settings.
Client Certificate Status Checking	Select a method to validate the trusted client certificate: • **None**—Do not validate. • **Use OCSP**—Use the OCSP method, validating the client certificate in real-time, as needed. After you select this option, you can specify options for OCSP. • **Use CRLs**—Use CRLs to validate the client certificate. After you select this option, you can specify options for OCSP. • **Use OCSP with CRL fallback**—Use the OCSP validation method when possible, but attempt to validate client certificates using CRLs if the OCSP method fails (for example, if the link to the OCSP responder fails). After you select this option, you can specify options for OCSP. • **Inherit from root CA**—Use the method configured for the device certificate.
CDP(s)/OCSP responder	Select the location of the responder value: • **None**—Do not use the responder. • **From client certificate**—Use the responder value configured in the client certificate. • **From trusted CA certificate**—Use the responder value configured in the trusted CA certificate that has been uploaded to the system.
Verify imported CA certificates	Select this option to verify that this trusted client CA is valid. Enabling this will check the CRL of this certificate's issuer, and repeat up the chain until reaching the root trusted client CA.

Configuring Options for Trusted Client CA Certificates

To configure options for the trusted client CA certificate:

1. Select **System > Configuration > Certificates > Trusted Client CAs**.

2. Click the certificate you want to configure.

3. Complete the configuration described in Table 123.

Table123: Trusted Client CA Settings

Settings	Guidelines
Certificate	Use the expander buttons to display the following details: • Issued To—Name and attributes of the entity to whom the certificate is issued. • Issued By—Name and attributes of the entity that issued the certificate. Note that the value of this field must match either the Issued To field (for root certificates) or the Issued To field of the next highest certificate in the chain (for intermediate certificates). • Valid Dates—Time range for which the certificate is valid. • Details—Various certificate details, including its version, serial number, signature algorithm, CRL distribution points, public key algorithm type, and public key.
Client Certificate Status Checking	Select a method to validate the trusted client certificate: • **None**—Do not validate. • **Use OCSP**—Use the OCSP method, validating the client certificate in real-time, as needed. After you have selected this option and saved the configuration, you can specify options for OCSP. • **Use CRLs**—Use CRLs to validate the client certificate. After you have selected this option and saved the configuration, you can specify options for CRL. • **Use OCSP with CRL fallback**—Use the OCSP validation method when possible, but attempt to validate client certificates using CRLs if the OCSP method fails (for example, if the link to the OCSP responder fails). After you have selected this option and saved the configuration, can specify options for OCSP and CRL.
Verify Trusted Client CA	Select this option to verify that this trusted client CA is valid. Enabling this will check the CRL of this certificate's issuer, and repeat up the chain until reaching the root trusted client CA.
Trusted for Client Authentication	Clear this check box to exclude the CA from being trusted for client certificate authentication. You might want to do this if this CA was added for another trusting purpose, such as SAML signature verification or machine certificate validation.
Participate in Client Certificate Negotiation	This option is available only on Connect Secure. Select this option to include the CA participation in client certificate selection for authentication. NOTE: In client certificate authentication or restriction, the device sends a list of all trusted client CAs configured in the trusted client CA store with this flag enabled to the user's browser for user certificate selection. The browser prompts the client certificates whose issuer CA and/or root CA is in that list. This option allows you to control which client certificate(s) are prompted for selection. Clearing this option for all certificates in a CA chain results in those certificates not being prompted.

4. Save your changes.

5. If you have enabled CRL Checking, click **CRL Checking Options**.

6. If you have enabled OCSP options:

    1. Click **OCSP Options**.

    2. Complete the configuration described in Table 124.

Table124: OCSP Options Settings

Settings	Guidelines
Use	Select the type of OCSP responder to validate trusted client CAs: • **None**—The system does not use OCSP to verify the status of certificates issued by this CA. • **Responder(s) specified in the CA certificate**—The system uses OCSP responders specified in the imported client CA to perform verification. When you select this option, the system displays a list of OCSP responders specified in the imported CA (if any) and the last time they were used. • **Responder(s) specified in the client certificates**—The system uses responders specified during client authentication to perform verification. When you select this option, the system displays a list of known OCSP responders (if any) and the last time they were used. • **Manually configured responders**—The system uses primary and secondary OCSP responders at the addresses you specify.
Device Certificate to sign the request	Select the appropriate device certificate or leave the default (unsigned).
Signature Hash Algorithm	Select **SHA-1** or **SHA-2**.
Use Nonce	A nonce is random data the system includes in an OCSP request and the OCSP responder returns in the OCSP response. The system compares the nonce in the request and response to ensure that the response is generated by the OCSP responder. If the two do not match, the system disregards the response and sends a new request. Nonces are a common way of preventing replay attacks.

3. Save the configuration.

4. After you have added an OCSP responder to the list, you can click its link to display the page.

5. Complete the configuration described in Table 125.

Table125: Responder Signer Certificate Settings

Settings	Guidelines
Responder Signer Certificate	Browse to the network path or local directory location of a Responder Signer Certificate. This is the certificate the OCSP responder uses to sign the response. You must specify the Responder Signer Certificate if the signer certificate is not included in the response.
Trust Responder Certificate	Select this option to allow an OCSP responder certificate that matches the responder signer certificate.
Revocation Checking	Select this option to ensure that the certificate has not recently been revoked. This option has implications only if you specified the Use OCSP with CRL fallback option.
Allow clock discrepancy	Use this option to account for possible mismatches in timestamps between the system clock and the OCSP responder clock. If the mismatch is significant, the system disregards the response from the OCSP responder as out of date or expired.

6. Save the configuration.

Configuring a Proxy Server for CRL Downloads and OCSP Status Checks

You can configure the system to send CRL download requests and OCSP status checks to the proxy server and collect the response. You might want to do this if you deploy proxy server to control access to the Internet.

The following types of CRL downloads can use the proxy server:

- CRL distribution points (CDPs) specified in the trusted client CAs

- CDPs specified in client certificates

- Manually configured CDPs

Similarly, the system can send OCSP requests to the OCSP responder through the proxy server. The OCSP responses are also received through the proxy server. This feature is useful when you deploy a large number of Pulse access systems and the OCSP responders are located outside the network.

To configure a proxy server:

1. Select **System > Configuration > Certificates > Trusted Client CAs**.

2. Click **Proxy Settings** to display the page.

3. Complete the configuration described in Table 126.

4. Save the configuration.

Table126: Proxy Settings

Settings	Guidelines
Use Proxy Server for HTTP-based CRL download	Select to enable the CRL operations to use a proxy server.
	NOTE: You can configure a proxy server for web-based URLs, not LDAP URLs.
Use Proxy Server for HTTP-based OCSP status checking	Select to enable the OCSP operations to use a proxy server.
Host Address	Specify either an IP address or a fully qualified domain name.
Port	Enter the proxy server port number if it is different from the default value of 80.
Username/password	If your proxy server required authentication, enter a username and password to log in to the proxy server.

Related Documentation

- Understanding Digital Certificate Security

Using Trusted Server CAs

This topic describes trusted server certificate authorities (CAs). It includes the following information:

- Understanding Trusted Server CAs

- Uploading Trusted Server CA Certificates

- Restoring the Prepopulated Group of Trusted Server CA Certificates

- Renewing a Trusted Server CA Certificate

- Deleting a Trusted Server CA Certificate

Understanding Trusted Server CAs

All of the trusted root CAs for the Web certificates installed in Internet Explorer are preinstalled. You might need to install a trusted server CA for additional Web servers in the following situations:

- If you are using third-party integrity measurement verifiers (IMVs) that are installed on a remote server, you must upload the trusted root certificate of the CA that signed the remote server's server certificate.

- If you are using virus signature version monitoring with your own staging site for storing the current virus signatures list, you must upload the trusted root certificate of the CA that signed the staging server certificate.

You can install the trusted root CA certificate on the endpoint in any of the following ways:

- Use a CA certificate that is chained to a root certificate that is already installed on the endpoint, such as VeriSign.

- Upload the CA certificate and any intermediate CA certificates to the Pulse Secure client system. During client installation, the system automatically installs the trusted root device CA certificates on the endpoint. When prompted during installation, the user must allow the installation of the CA certificate(s).

- Prompt users to import the CA certificates on the endpoint using Internet Explorer or other Microsoft Windows tools. In other words, you can use common methods organizations use to distribute root certificates.

   NOTE: You cannot use CRL revocation checks for trusted server CA certificates.

Uploading Trusted Server CA Certificates

You can use the Trusted Server CAs page to upload the trusted root certificate of the CA that signed the Pulse Secure client service device certificate. If you upload a certificate chain, you must install the certificates one at a time in descending order starting with the root certificate (DER or PEM files), or you must upload a single file that contains the entire certificate chain (PEM files only). The system supports X.509 CA certificates in PEM (Base 64) and DER (binary) encode formats.

To upload CA certificates:

1. Select **System > Configuration > Certificates > Trusted Server CAs** to display the page.

2. Click **Import Trusted Server CA** to display the page.

3. Browse to the certificate file, select it, and click **Import Certificate** to complete the import operation.

Restoring the Prepopulated Group of Trusted Server CA Certificates

The System > Configuration > Certificates > Trusted Server CAs page is prepopulated with some of the trusted root CAs for the Web certificates installed in Internet Explorer and Windows. You can use the delete functionality on this page to delete CAs and the reset functionality to restore the list to the set that was installed during the upgrade. The reset operation clears all manually imported certificates.

To restore the prepopulated group of trusted CA certificates:

1. Select **System > Configuration > Certificates > Trusted Server CAs**.

2. Click **Reset Trusted Server CAs**.

3. Confirm that you want to restore the set of trusted server CAs that was installed when you upgraded.

   Connect Secure restores the group of prepopulated trusted server CAs that were installed upon upgrade. This operation clears all manually imported certificates.

Renewing a Trusted Server CA Certificate

If a trusted CA renews its certificate, you must upload the renewed CA certificate.

To import a renewed CA certificate:

1. Select **System > Configuration > Certificates > Trusted Server CAs**.

2. Click the link that corresponds to the certificate that you want to renew to display the page.

3. Click **Renew Certificate**.

4. Browse to the certificate file, select it, and click **Import Certificate** to complete the import operation.

Deleting a Trusted Server CA Certificate

You can delete any trusted server CA certificate, including preinstalled certificates.

To delete a trusted server CA certificate:

1. Select **System > Configuration > Certificates > Trusted Server CAs**.

2. Select the check box for the certificate you want to delete.

3. Click **Delete**, and then confirm that you want to delete the certificate.

Related Documentation

- Understanding Digital Certificate Security

Using Code-Signing CAs

This topic describes how to use code-signing Certificate Authorities (CAs). It includes the following information:

- Understanding Code-Signing CAs
- Additional Considerations for Oracle JVM Users
- Importing a Code-Signing CA Certificate
- Using Code-Signing Certificates for Java Applets

Understanding Code-Signing CAs

In a basic setup, the only required certificates are a device certificate and a code-signing certificate. Connect Secure can use a single code-signing certificate to resign all Java applets and a single device certificate to intermediate all other PKI-based interactions. If the basic certificates do not meet your needs, however, you may install multiple device and applet certificates on Connect Secure or use trusted CA certificates to validate users.

When Connect Secure intermediates a signed Java applet, it re-signs the applet with a self-signed certificate by default. This certificate is issued by a nonstandard trusted root CA. As a result, if a user requests a potentially high-risk applet (such as an applet that accesses network servers), the user's Web browser alerts him that the root is untrusted.

If you import a code-signing certificate, Connect Secure uses the imported certificate to re-sign applets instead of the default self-signed certificate. As a result, if a user requests a potentially high-risk applet, the user's Web browser displays an informational message instead of a warning. The message informs the user that the applet is signed by a trusted authority.

The system supports the following types of code-signing certificates:

- Microsoft Authenticode Certificate—The system uses this certificate to sign applets that run on either Microsoft JVM or Oracle JVM. Note that we only support Microsoft Authenticode Certificates issued by Verisign. You may purchase Microsoft Authenticode Certificates at the following location:

  http://www.verisign.com/products-services/security-services/code-signing/index.html

- JavaSoft Certificate—The system uses this certificate to sign applets that run on Oracle JVM. Note that we only support JavaSoft Certificates issued by Verisign and Thawte.

When deciding which code-signing certificate to import, consider the following browser dependencies:

- Internet Explorer—Internet Explorer running on new computers shipped with Windows pre-installed typically runs the Oracle JVM, which means that Connect Secure needs to re-sign applets using the JavaSoft certificate.

  Internet Explorer running on an older version of Windows that has been upgraded may run the Microsoft JVM, which means that Connect Secure needs to re-sign applets using the Authenticode certificate.

- Netscape, Firefox, and Safari—These browsers only support the Oracle JVM, which means that Connect Secure needs to re-sign applets using the JavaSoft certificate.

  NOTE: If you create IVS systems, you can also import code-signing certificates for each IVS. You must navigate to each IVS system, using the IVS system drop-down menu in the admin console header, then import the code-signing certificate for each IVS on the System > Configuration > Certificates > Code-signing Certificates page.

Additional Considerations for Oracle JVM Users

By default, the Java Plug-in caches an applet along with the code-signing certificate presented when a user accesses the applet. This behavior means that even after importing a code-signing certificate to Connect Secure, the browser continues to present applets with the original certificate. To ensure that JVM users are not prompted with an untrusted certificate for applets accessed prior to importing a code-signing certificate, users need to flush the Java Plug-in cache. Alternatively, users can disable the cache, but this option may impact performance since the applet needs to be fetched each time the user accesses it.

The Java Plug-in maintains its own list of trusted Web server certificates that is different from the browser's list of trusted certificates. When a user accesses an applet, the JVM makes its own connection (in addition to the browser) to the Web server on which the applet resides. The user is then presented with the option to accept the Web server certificate in addition to the code-signing certificate. In these cases, the user needs to select the Always Trust button for the Web server certificate. Due to a built-in timeout in the Java Plug-in, if the user waits too long to select this button for the Web server certificate, the applet does not load.

Importing a Code-Signing CA Certificate

To import a code-signing certificate:

1. Select **System > Configuration > Certificates > Code-Signing Certificates** to display the configuration page.

2. Click **Import Certificates** to display the configuration page.

3. Complete the configuration described in Table 127.

Table127: Import Certificates Configuration Guidelines

Settings	Guidelines
Microsoft Authenticode or Multipurpose Certificate for Internet Explorer (Microsoft JVM)	
Certificate File	Browse to the network path or local directory location of your certificate key file.
Private Key File	Browse to the network path or local directory location of your private key file.
Password Key	Enter the password key.
Javasoft Certificate for Internet Explorer & Netscape (Sun JVM)	
Keystore File	Browse to the network path or local directory location of the keystore file.
Password key	Enter the password key.

4. Click **Import** to complete the import operation.

5. When you have successfully imported a certificate, the system displays the Sign Pulse Secure Web Controls With dialog box. Specify the signing option:

   - **Default Pulse Secure Certificate**—Select this option to sign all ActiveX and Java applets originating from Connect Secure using the default Pulse Secure certificate. If you have previously selected an imported code-signing certificate and are reverting back to this option, after you click Save, a process icon appears indicating that the system is processing the request and re-signing all of the relevant code. This process can take several minutes to complete.

   - **Imported Certificate**—Select this option to sign all ActiveX and Java applets using the certificate or certificates imported in the previous step. When you click Save, a process icon appears indicating that the system is processing the request and signing all of the relevant code. This process can take several minutes to complete.

6. Use settings in the following tabs to specify which resources are signed or re-signed by the applet certificate:

   - Users > Resource Policies > Web > Java > Code Signing

   - Users > Resource Policies > Web > Java > Applets

Using Code-Signing Certificates for Java Applets

To use code-signing certificates with Java applets:

1. Install the Java code-signing certificates. Use the **System > Configuration > Certificates > Code-Signing Certificates** page.

2. Use any of the following methods:

   - Create code-signing policies specifying which applets to re-sign. Use the **Users > Resource Policies > Web > Java > Code Signing** page or the **Users > Resource Profiles > Web Application Resource Profiles > Profile** page. The policies must specify the hostnames from which the applets originate.

   - Upload your own Java applets to Connect Secure and configure the system to sign or re-sign the applets.

Related Documentation

- Creating a Custom Web Application Resource Profile

- Understanding Digital Certificate Security

Using Client Auth Certificates

This topic describes how to use client auth certificates. It includes the following information:

- Understanding Client Auth Certificates
- Importing a Client Auth Certificate
- Renewing a Client Auth Certificate
- Configuring Two-Way SSL Authentication

Understanding Client Auth Certificates

In certain corporate environments, servers on the LAN are protected with two-way SSL authentication. These servers require the client to authenticate by presenting a valid certificate.

In the remote access scenario, Connect Secure is a client of these servers. You can configure Connect Secure to present client authentication certificates to servers whenever it communicates over SSL. Note that Connect Secure will present client certificates only when the SSL handshake requires it.

> **ⓘ** NOTE: This feature authenticates Connect Secure (as a client) to back-end servers. It also authenticates end users or end-user machines to servers on the corporate LAN.

The SSL protocol provides for mutual authentication of server and client at the time of session initiation. The client part of the authentication is optional. For enhanced security, some deployments may require that the client also authenticate itself with a certificate. Normally, when setting up an SSL connection with a server on behalf of the end user, Connect Secure does not present any certificate to the server. It needs to be explicitly configured to present such certificate. This section explains how such configuration may be performed.

The basic idea is to upload a certificate, private key pair to the Pulse Secure access management framework, and configure a mapping between this pair and a server resource. Subsequently, when an end user attempts to establish a connection with that server, Policy Secure and Connect Secure presents the associated certificate to the server. If no certificate is associated with the server in Policy Secure and Connect Secure's certificate store, then it is assumed that the server does not demand client certificate.

If, during the SSL handshake, the back-end server requests a client certificate but Connect Secure doesn't send a certificate, the end user sees an "access denied" error message. Similarly, if the back-end server rejects the Connect Secure certificate, the end user sees an "access denied" error message. If a certificate is configured, is successfully retrieved and no error is encountered during handshake, the user is granted access to the server.

> **ⓘ** NOTE: The Pulse Secure access management framework allows client authentication certificates to be uploaded to the device in two ways: generate a CSR and upload the signed certificate returned by the CA, or directly import the certificate if one is available. The first option is not available on FIPS devices.

Importing a Client Auth Certificate

The Pulse Secure access management framework allows certificates that include the private key and for instances where the private key is in a separate file from the certificate. In addition, if your certificates have been exported into a system configuration file, you can import the system configuration file to upload the certificates.

To import the client auth certificates files:

1. Select **System > Configuration > Certificates > Client Auth Certificates**.

2. Click **Import Certificate & Key** to display the configuration page.

3. Complete the configuration described in Table 128.

4. Click **Import**.

Table128: Import Certificate and Key Settings

Settings	Guidelines
If certificate file includes private key	
Certificate File	Browse to the network path or local directory location of your private key file.
Password Key	Enter the password key.
If certificate and private file are separate keys	
Certificate File	Browse to the network path or local directory location of your certificate key file.
Private Key File	Browse to the network path or local directory location of your private key file.
Password Key	Enter the password key.
Import via System Configuration file	
System Configuration File	Browse to the network path or local directory location of the system configuration file.
Password	Enter the password.

Renewing a Client Auth Certificate

To renew a certificate:

1. Select **System > Configuration > Certificates > Client Auth Certificates**.

2. Click the link that corresponds to the certificate you want to renew.

3. Click **Renew Certificate** to display the configuration page.

4. In the Renew the Certificate form, browse to the renewed certificate file, enter the password for the certificate key, and click **Import**.

Configuring Two-Way SSL Authentication

To configure two-way SSL authentication:

1. Import the certificates used for two-way SSL handshake in the System > Configuration > Certificates > Client Auth Certificates window.

2. Define the back-end resource and assign a certificate to be presented when accessing it using the Users > Resource Policies > Web > Client Authentication window.

Related Documentation

- Mapping Resource Policies to the Certificate

- Understanding Digital Certificate Security

- Configuring Syslog

Mapping Resource Policies to the Certificate

Once the certificates have been uploaded, you can map resources to the certificates and the roles to which they apply.

1. Select **Users > Resource Policies > Web > Client Authentication**.

2. If you do not see the Client Authentication menu item, select Users > Resource Policies > Web.

   1. Click the **Customize** button in the upper right corner of the console.

   2. In the Customize View dialog box, select **Client Authentication**.

   3. Click **OK**.

   4. Click the **Client Authentication** tab.

5. Click **New Policy**.

6. On the New Policy page:

   - Enter a name to label this source interface policy.

   - Enter an optional description.

7. In the Resources section, specify the back-end servers to which this policy applies. Valid values/formats are: hostnames, IP addresses, IP Address:Port and Hostname:Port.

   If you specify * as the resource, one certificate is used for all resources requesting a back-end certificate authentication. This certificate becomes the default certificate. Defining a default certificate is not required.

8. In the Roles section, select one of the following options:

   - **Policy applies to ALL roles**—To apply this policy to all users.

   - **Policy applies to SELECTED roles**—To apply this policy only to users who are mapped to roles in the Selected roles list. Make sure to add roles to this list from the Available roles list.

   - **Policy applies to all roles OTHER THAN those selected below**—To apply this policy to all users except for those who map to the roles in the Selected roles list. Make sure to add roles to this list from the Available roles list.

9. In the Action section, select one of the following options:

   - **Use the Client Authentication Certificate Below**—Select this option to associate this policy with a client authentication certificate. Select the certificate to use from the Certificate menu.

     If the Certificates menu is blank, no certificates have been uploaded to the System > Configuration > Certificates > Client Auth Certificates window.

   - **Do not use Client Authentication**—If this option is selected, the system does not perform client authentication for the configured resource.

   - **Use Detailed Rules**—Select this option to specify one or more detailed rules for this policy.

10. Click **Save Changes**.

Related Documentation

- Using Client Auth Certificates

- Writing a Detailed Rule for Resource Policies

Mapping a Client Authentication Auto-Policy

A client authentication auto-policy option is available on the Users > Resource Profiles > Web page. If the back-end server requires two-way SSL authentication, this auto-policy lets you configure a certificate to be presented during the SSL handshake.

1. Select **Users > Resource Profiles > Web**.

2. Follow the process as a regular resource profile for defining the name and type.

3. Select the **Autopolicy: Client Authentication** check box.

4. In the Resource field, specify the back-end server. Valid formats/values are: hostnames, IP addresses, IP Address:Port, and HostName: Port.

   If you specify * as the resource, one certificate is used for all resources requesting a back-end certificate authentication. This certificate becomes the default certificate. Defining a default certificate is not required.

5. Click **Save Changes**.

Related Documentation

- Resource Profiles

CHAPTER 34 Elliptic Curve Cryptography

Understanding ECC Certificates

Public-key cryptography is a cryptographic system that requires a secret key and a public key that are mathematically linked with each other. One key encrypts the plain text while the other decrypts the cipher text. RSA is the most widely used public-key algorithm.

Elliptic Curve Cryptography (ECC) were introduced as an alternative to RSA in public key cryptography. One advantage of ECC over RSA is key size versus strength. For example, a security strength of 80 bits can be achieved through an ECC key size of 160 bits, whereas RSA requires a key size of 1024. With a 112-bit strength, the ECC key size is 224 bits and the RSA key size is 2048 bits.

The most popular signature scheme that uses elliptic curves is called the Elliptic Curve Digital Signature Algorithm (ECDSA). The most popular key agreement scheme is called Elliptic Curve Diffie-Hellman (ECDH). An ECDH exchange is a variant of the Diffie-Hellman (DH) protocol and is an integral part of the Suite B cryptography standards proposed by the National Security Agency (NSA) for protecting both classified and unclassified information.

About Suite B

The Advanced Encryption Standard (AES) is a specification for the encryption of electronic data established by the U.S. National Institute of Standards and Technology (NIST) in 2001. Because a single encryption algorithm cannot satisfy all of the needs of the national security community, NSA created a larger set of cryptographic algorithms, called Suite B, which can be used along with AES in systems used by national security users. In addition to AES, Suite B includes cryptographic algorithms for hashing, digital signatures, and key exchanges.

Per RFC 6460, to be Suite B TLS 1.2 compliant the server and client should negotiate with the following ciphers:

- TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256
- TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384

RFC 6460 also lists a transitional Suite B profile for TLS 1.0 and TLS 1.1. Clients and servers that do not yet support Suite B TLS 1.2 should negotiate with the following ciphers:

- TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA
- TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA

There is no special configuration to ensure that Connect Secure and Policy Secure negotiates Suite B ciphers. However, the following general steps should be performed to enable Suite B compliance:

- An ECC certificate signed by an ECC Root CA is associated with a network port.

  - A P-256 CSR is signed by either a P-256 or P-384 Root CA.

  - A P-384 CSR is be signed by a P-384 Root CA.

- Manually enable only AES128 and/or AES256 custom ciphers.

  NOTE: Connect Secure and Policy Secure cannot be configured to allow only Suite B ciphers.

Using ECC Certificates

ECC certificates are currently supported only on the MAG Series Pulse Secure Gateways and virtual appliance platforms. As with RSA certificates, ECC certificates are associated with a network port. You can create multiple virtual ports on the server with each port supporting a specific certificate. For example, external virtual port 1 can use a 1024-bit RSA while external virtual port 2 uses ECC P-256 and external virtual port 3 uses ECC P-384. Only clients that support ECC cipher suites can connect to the web server on that network port.

When an Elliptic Curve Cryptography (ECC) certificate is associated with a network port, only clients that support ECC cipher suites can connect to the Web server on that network port.

Except for the key and certificate generation process, the use of ECC certificates is basically the same as using RSA certificates.

Related Documentation

- Example: Assigning an ECC P-256 Certificate to an External Virtual Port and Giving Preference to Suite B Ciphers

Example: Assigning an ECC P-256 Certificate to an External Virtual Port and Giving Preference to Suite B Ciphers

This example outlines the general steps for creating an external port and assigning an ECC P-256 certificate. The steps are generally the same as assigning an RSA certificate to a port.

- Configuring the External Port
- (optional) Configuring the Virtual Ports
- Creating the Certificate Signing Request for a New Certificate
- Importing the Signed Certificate Created from a CSR
- Presenting the Certificate on the Network
- Setting the Security Options

Configuring the External Port

The external port handles all requests from users signed into the server from outside the customer LAN, for example, from the Internet.

To configure the external port:

1. In the admin GUI, choose **System > Network > External Port > Settings**.

2. Modify the settings as needed. In this example, only IPv4 is enabled. See Figure 165.

Figure168: Configuring the External Port for IPv4

3. Click **Save Changes**.

(optional) Configuring the Virtual Ports

A virtual port is an IP alias bound to a physical port. It shares all of the network settings, except IP address, with the associated physical port. You can use virtual ports for different purposes, depending on the physical port or the VLAN on which you base the virtual port. In this example, we are configuring the virtual port on the external port to support external sign-ins. This is an optional step that shows one way of allowing multiple certificates on the device.

To configure the external virtual port:

1. In the admin GUI, choose **System > Network > External Port > Virtual Ports**.

2. Click **New Port**.

   In this example, the port is named p_ecdsa256 and accepts only IPv4 addresses. See Figure 166.

Figure169: Creating the Virtual Port on the External Port

3. Click **Save Changes**.

Creating the Certificate Signing Request for a New Certificate

A certificate signing request (CSR) is a message sent from an applicant to a certificate authority (CA) to apply for a digital identity certificate. You create a CSR through the admin console. When you create a CSR, a private key is created locally that corresponds to the CSR. If you delete the CSR at any point, the private key is deleted too, prohibiting you from installing a signed certificate generated from the CSR.

In this example, a CSR for an ECC P-256 certificate is requested.

To create a CSR for a new certificate:

1. In the admin console, choose **System > Configuration > Certificates > Device Certificates**.

2. Click **New CSR**.

3. Enter the required requestor information. In this example, the common name is ecc-p256.juniper.net.

4. Click **ECC** and select **P-256** from the ECC Curve menu. See Figure 167.

Figure170: Creating an ECC P-256 Certificate Signing Request

5. Click **Create CSR**. A CSR is successfully created, as shown in Figure 168.

Figure171: CSR Successfully Created
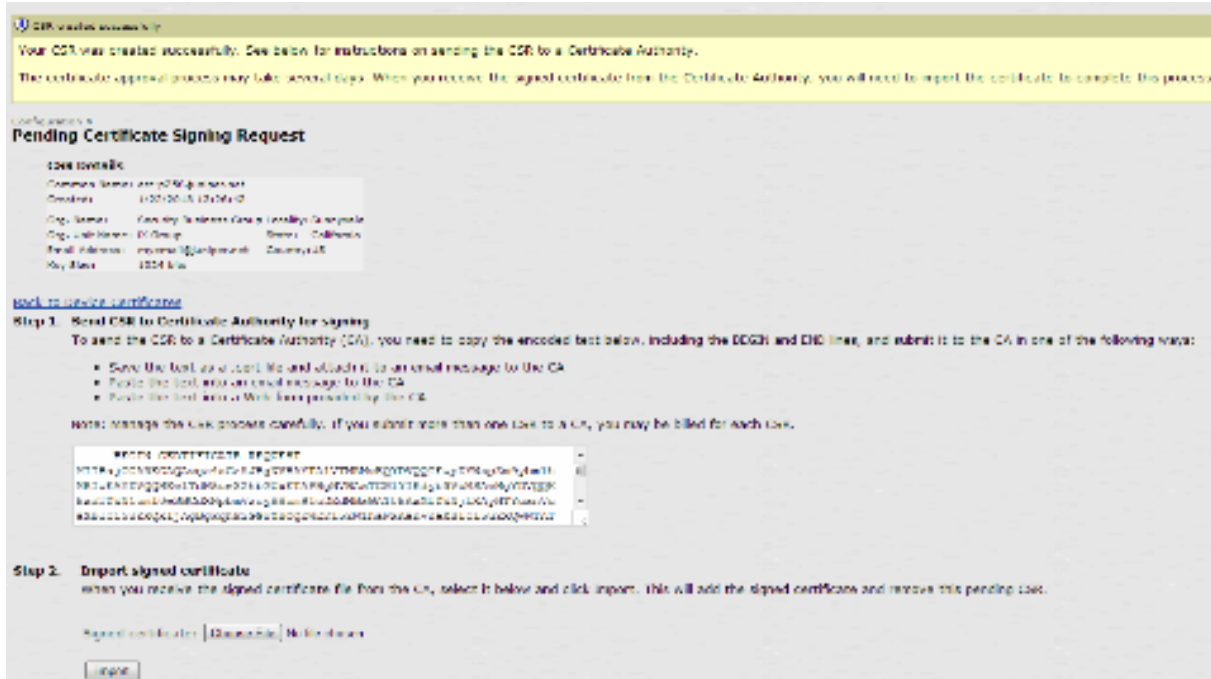
6. The CSR is encoded and can be copied or saved to a file. The ECC certificate should be signed by an ECC CA for Suite B compliance. Follow your CA's process for sending a CSR.

7. Click the **Back to Device Certificates** link. Until you import the signed certificate from your CA, your CSR is listed as Pending. See Figure 169.

Figure172: Pending CSR
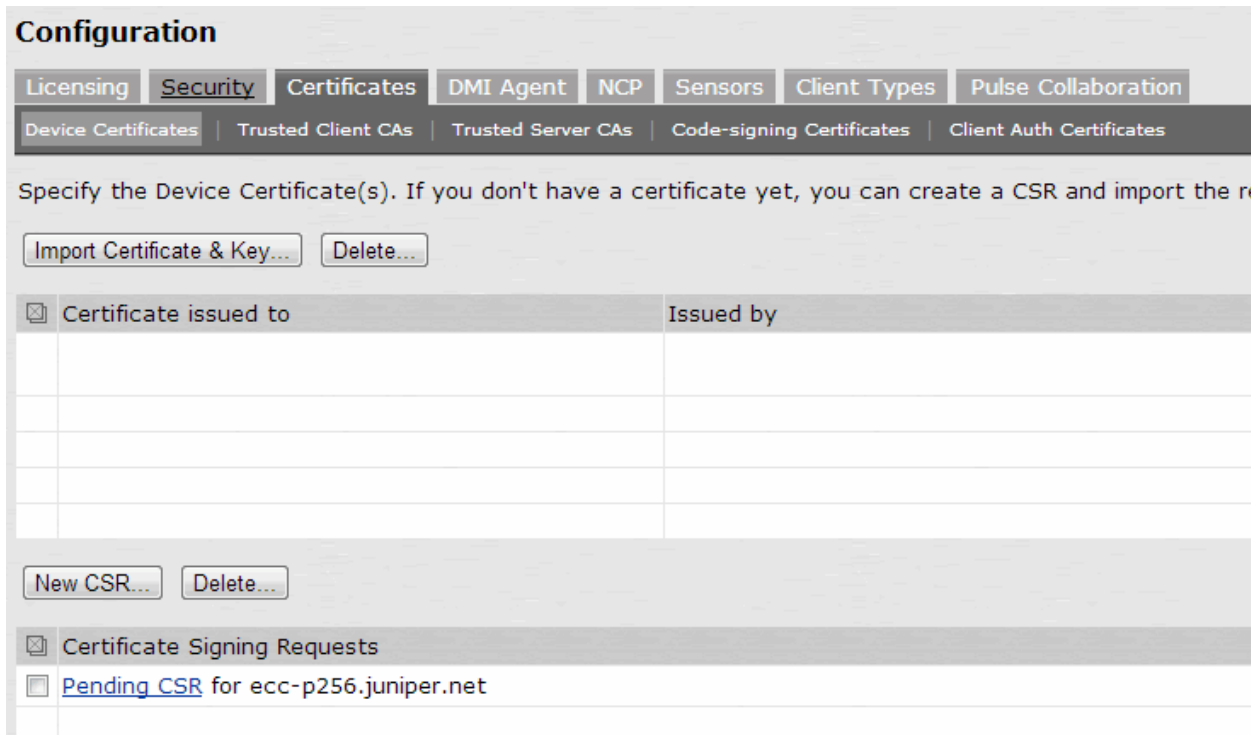
Importing the Signed Certificate Created from a CSR

Once your CA has sent your signed certificate, you must import that into the pending CSR.

To import a signed device certificate created from a CSR:

1. In the admin console, choose **System > Configuration > Certificates > Device Certificates**.

2. Under Certificate Signing Requests, click the **Pending CSR** link that corresponds to the signed certificate. See Figure 169.

3. Under Import signed certificate, browse to the certificate file you received from the CA and then click **Import**. See Figure 168.

Presenting the Certificate on the Network

You can present a certificate many ways, depending on your configuration. For example, you can present the certificate to one or more virtual ports or on an internal or external port. In this example, the ECC P-256 certificate is presented on the external virtual port p1.

To present a certificate on an external virtual port:

1. In the admin console, select **System > Configuration > Certificates > Device Certificates**.

2. Click the certificate name you want to assign to a port. In this example, click **ecc-p256.juniper.net**.

3. Under External Ports, select **p_ecdsa256** and click **Add**. See Figure 170.

Figure173: Associating the ECC P-256 with the External Virtual Port p_ecdsa256

4. Click **Save Changes**.

Setting the Security Options

To specify the cipher suites for the incoming connection to the Web server, use the SSL Options page and select the **Custom SSL Cipher Selection** option. This step is required in our example to give Suite B cipher suites preference. If you do not want to give Suite B cipher suites preference, you do not have to perform this step.

To set the security options:

1. In the admin console, select **System > Configuration > Security > SSL Options**.

2. Under Allowed Encryption Strength choose **Custom SSL Cipher Selection**. See Figure 171.

Figure174: Setting Custom SSL Cipher Selections

Allowed Encryption Strength
Strong ciphers (rated by the number of bits in the cipher) improve the security of SSL encryption, but some browsers ma
Secure Access Service will give preference to the cipher with the fastest data transfer rate, regardless of its relative enc
Please see the Setting Security Options section in the Admin guide for more details.

- ○ Accept only 168-bit and greater (maximize security)
- ○ Accept only 128-bit and greater (security and browser compatibility)
- ○ Accept 40-bit and greater (maximize browser compatibility)
- ● Custom SSL Cipher Selection

Note: Custom cipher selection disables the **Encryption Strength option**.

Cipher Suites	High (>=168-bit)	Medium (>=128-bit and <168-bit)	Low (>=40-bit and <128-bit)	Description
AES/3DES	☑			cipher suites using triple DES and AES (Maximizes security)
AES		☑		cipher suites using AES (Maximizes performance)
DES			☐	cipher suites using DES (not triple DES)
RC4		☐		cipher suites using RC4

■ Recommended for FIPS deployment

Encryption Strength option
Normally, the allowed encryption strength is enforced after an SSL session is established, so that a user that connects us
option below will prevent a browser with a weak cipher from establishing a connection. Changing this option will cause the

☑ Do not allow connections from browsers that only accept weaker ciphers

SSL Handshake Timeout option
By default, the SSL handshake has a timeout of 60 seconds. Use the text box below to set a different value.

Timeout: [60] seconds 10-600 seconds

SSL Legacy Renegotiation Support option
When this option is enabled, renegotiation with clients and servers, which dont support the new TLS Renegotiation Info e
and servers will not be allowed. Changing this option will cause the web service to restart.

☑ Enable support for SSL legacy renegotiation

3. Choose the recommended options (in green) if they are not already selected.

4. If you are using client certificate authentication (Connect Secure only):

- Select **Enable client certificate on the external port** under ActiveSync Client Certificate Configuration.

- Move p_ecdsa256 to the Selected Virtual Ports column.

5. Click **Save Changes**.

A list of the custom ciphers to be used on the device's port is displayed in the order the web server will select them. Note that Suite B ciphers are listed on top. See Figure 172. End

users who now log in to external virtual port p_ecdsa256 must have at least one of the listed ciphers installed on their browser or else they cannot log in to the server.

Figure175: Confirming Custom Ciphers

6. Click **Change Allowed Encryption Strength**.

Related Documentation

- Using Device Certificates

- Verifying the Certificate on the Client

- Understanding ECC Certificates

Verifying the Certificate on the Client

End users can check which certificate their browser is using to connect to the server. In the following example, the end user connects to server port p3, which uses an ECC curve P-256 certificate. See Figure 173.

Figure176: Connecting to a Port Using an ECC Curve P-256 Certificate
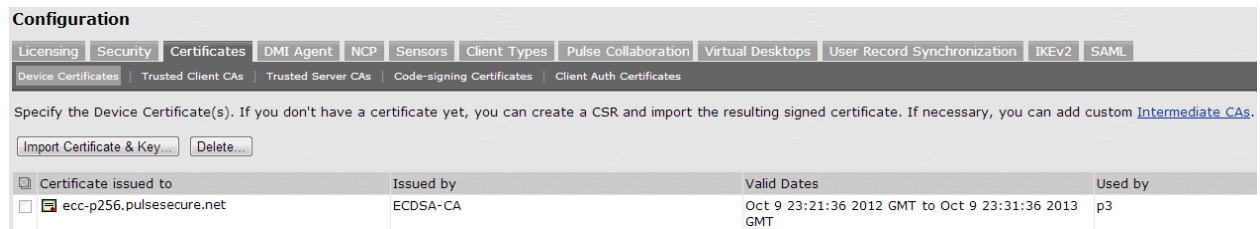
To view the certificate from an Internet Explorer 8 browser:

1. Open an Internet Explorer 8 browser and point to the server to which you want to connect.

2. Click the lock icon located at the end of the address bar and then click the **View Certificate** link. See Figure 174.

Figure177: Viewing the Connection Certificate Information

3. Click the **Details** tab and scroll down until you see the Public key field. In this example, the public key value is ECC (256 Bits) which matches the server port p3 certificate shown in Figure 173. See Figure 175.

Figure178: Certificate Public Key
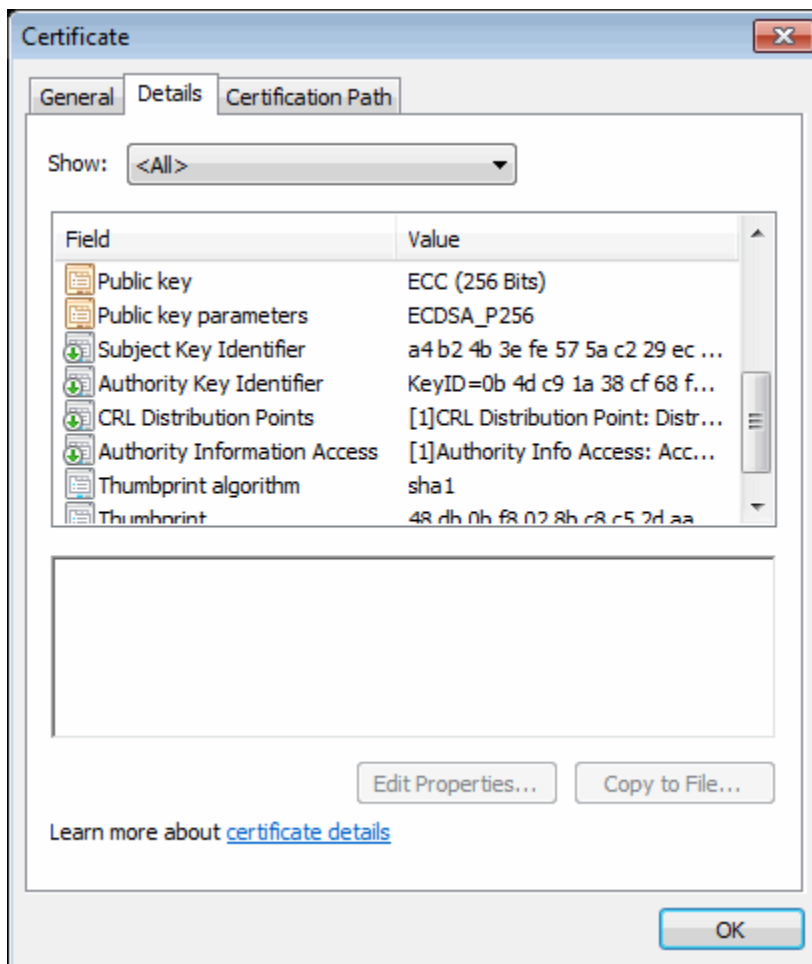
Related Documentation

- Understanding ECC Certificates

- Using TCP Dump to View Cipher Information

Using TCP Dump to View Cipher Information

You can use the TCP Dump tool to view which cipher each client uses to connect to the server. TCP Dump is a packet analyzer that intercepts (sniffs) and displays TCP/IP and other packets transmitted or received between the server and clients.

NOTE: To permit debugging, it is recommended that the ECC certificate be replaced by an RSA certificate so that an RSA cipher suite gets selected and then the application data can be decoded.

To capture packet headers:

1. Select **Maintenance > Troubleshooting > Tools > TCP Dump**.

2. Select the interface, internal or external or both, you wish to sniff and then the VLAN port.

3. Click **Start Sniffing**.

   The next time a user points a browser window to the server or logs in to the server, handshake information is obtained.

4. Click **Stop Sniffing** when done.

To view the packet headers:

1. Select **Maintenance > Troubleshooting > Tools > TCP Dump**.

2. Under Dump file, select **SSLDump** from the file menu and the certificate to use. See Figure 176.

Figure179: Viewing the TCP Dump Output
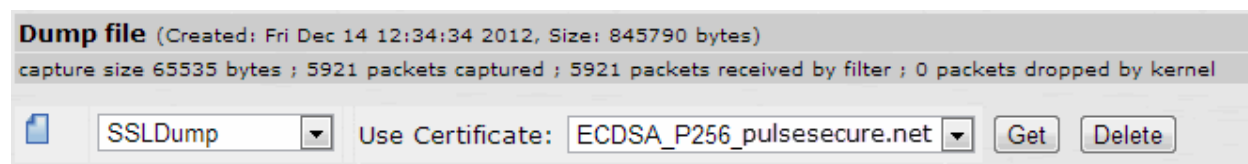
The certificate names in the TCP Dump window are the same as the "Certificate issued to" names in the Device Certificates window. Select the certificate corresponding to the port you wish to view packet information. See Figure 177.

Figure180: Issued to Certificate on the Device Certificates Pages
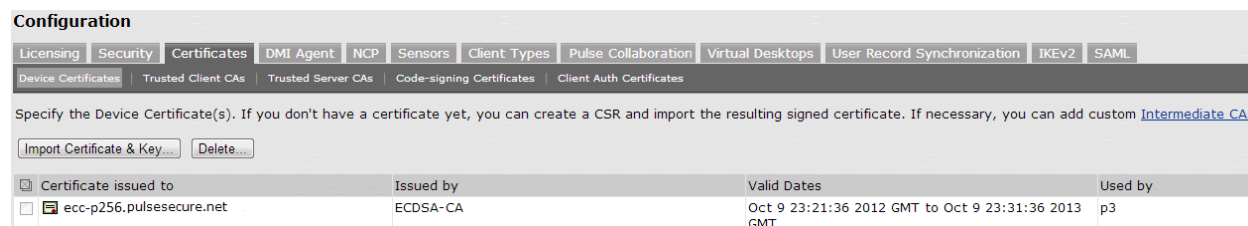
3. Click **Get**.

Portions of a TCP dump output follow.

The client starts a handshake with the server:

```
1 1  0.0007 (0.0007)  C>S  Handshake
```

The client then lists its supported cipher suites:

```
cipher suites
        TLS_ECDH_ECDSA_WITH_AES_256_GCM_SHA384
        TLS_ECDH_ECDSA_WITH_AES_256_SHA384
        TLS_ECDH_ECDSA_WITH_AES_256_SHA
        TLS_ECDH_ECDSA_WITH_DES_CBC3_SHA
    ...
```

The server acknowledges the handshake:

```
1 2  0.0010 (0.0003)  S>C  Handshake
```

The server compares the cipher suites on the client with the ones on the server and picks the cipher suite that is preferred by the server based on SSL options:

```
cipherSuite        TLS_ECDH_ECDSA_WITH_AES_256_GCM_SHA384
```

Example TCP Dump Output

```
New TCP connection #1: 10.64.8.3(46200) <-> 10.64.90.21(443)
1 1  0.0007 (0.0007)  C>S  Handshake
      ClientHello
        Version 3.3
        cipher suites
        TLS_ECDH_ECDSA_WITH_AES_256_GCM_SHA384
        TLS_ECDH_ECDSA_WITH_AES_256_SHA384
        TLS_ECDH_ECDSA_WITH_AES_256_SHA
        TLS_ECDH_ECDSA_WITH_DES_CBC3_SHA
        TLS_ECDH_ECDSA_WITH_AES_128_GCM_SHA384
        TLS_ECDH_ECDSA_WITH_AES_128_SHA256
        TLS_ECDH_ECDSA_WITH_AES_128_SHA
        TLS_ECDH_ECDSA_WITH_RC4_SHA
        Unknown value 0xc001
        TLS_EMPTY_RENEGOTIATION_INFO_SCSV
        compression methods
                NULL
        ClientHello Extensions [113]=
          00 6f 00 0b 00 04 03 00 01 02 00 0a 00 34 00 32
          00 0e 00 0d 00 19 00 0b 00 0c 00 18 00 09 00 0a
          00 16 00 17 00 08 00 06 00 07 00 14 00 15 00 04
          00 05 00 12 00 13 00 01 00 02 00 03 00 0f 00 10
          00 11 00 23 00 00 00 0d 00 22 00 20 06 01 06 02
          06 03 05 01 05 02 05 03 04 01 04 02 04 03 03 01
          03 02 03 03 02 01 02 02 02 03 01 01 00 0f 00 01
          01
1 2  0.0010 (0.0003)  S>C  Handshake
      ServerHello
        Version 3.3
        session_id[32]=
          a3 07 40 6e 73 12 c2 4d f3 7d b9 77 f8 97 e1 94
          fc 1b 51 6a 66 3c 99 d6 c7 7d 0e fa 29 2e d0 c4
        cipherSuite        TLS_ECDH_ECDSA_WITH_AES_256_GCM_SHA384
        compressionMethod                      NULL
        ServerHello Extensions [20]=
          00 12 ff 01 00 01 00 00 0b 00 04 03 00 01 02 00
          0f 00 01 01
1 3  0.0010 (0.0000)  S>C  Handshake
      Certificate
1 4  0.0010 (0.0000)  S>C  Handshake
      ServerHelloDone
1 5  0.1413 (0.1403)  C>S  Handshake
      ClientKeyExchange
1 6  0.1413 (0.0000)  C>S  ChangeCipherSpec
1 7  0.1413 (0.0000)  C>S  Handshake
1 8  0.1464 (0.0051)  S>C  ChangeCipherSpec
1 9  0.1464 (0.0000)  S>C  Handshake
```

```
1 10 9.2389 (9.0924)  C>S  application_data
1 11 9.5828 (0.3438)  C>S  application_data
1 12 9.5833 (0.0004)  S>C  application_data
1    9.5833 (0.0000)  S>C  TCP FIN
1 13 9.5999 (0.0166)  C>S  Alert
1    9.5999 (0.0000)  C>S  TCP FIN
```

Related Documentation

- Understanding ECC Certificates

- Example: Assigning an ECC P-256 Certificate to an External Virtual Port and Giving Preference to Suite B Ciphers

CHAPTER 35 Configuration File Administration

Configuration File Administration Overview

The system supports multiple administrator utilities related to configuration file management. Table 129 describes the purpose of the different utilities.

Table129: Utilities for Configuration File Administration

Utility	Recommended Usage
Archiving	Schedule periodic backups to a remote backup server. You should schedule archiving for both the system configuration binary file (system.cfg) and the user configuration binary file (user.cfg). If necessary, you can import an archived configuration using the configuration binary file import/export feature.
Local backup and restore	Create backups on the local system as a precaution when making significant configuration changes. With this utility, you can quickly restore to a previous configuration.
Binary configuration file import/export	Export binary configuration files to a local host (an alternative to the remote archiving server and archiving process that runs as a scheduled job). You might do this if you do not use or do not have access to an archiving server, or if you want to make use of a configuration that has not yet been archived. You can export the binary system configuration file (system.cfg) and the binary user configuration file (user.cfg).

You can use the binary file import/export feature to clone a configuration that you want to deploy more broadly, such as deploying a backup device or to a group of devices. You can use "selective import" options to exclude unique network identifiers (such as IP address) that would cause problems if the configuration were to be wholly imported and activated. |

| XML configuration file import/export | Import or export the configuration for only the features and settings you select. This enables you to take a more granular approach to mass configuration management than the binary file import/export feature. For example, you might want to populate an authentication server configuration across a large number of nodes. You can export just that configuration element, and when you import it in the other nodes, you do not overwrite the large number of configuration elements that you would if you had imported the user.cfg file.

You might also find the XML file import/export feature useful when managing a single node. For example, you might want to add many new users to the local authentication server, which can be faster editing the XML than using the user interface. Or you might want to make global changes to the configuration object naming conventions or descriptions as part of a "housekeeping" initiative. This, too, might be accomplished faster editing the XML than clicking through the user interface. |
|---|---|
| Push configuration | Push a partial configuration from the running configuration on the source system to the running configuration on one or more target systems. This is the best option to instill common configuration elements if the devices that already deployed and currently online. |

Related Documentation

- Configuring Archiving for System Logs, Configuration Files, and Snapshots

- Using the Configuration Backup and Restore Feature

- Using the Import/Export Feature for Binary System Configuration Files

- Using the Import/Export Feature for Binary User Configuration Files

- Using the Import/Export Feature for XML Configuration Files

- Using the Push Configuration Feature

Configuring Archiving for System Logs, Configuration Files, and Snapshots

You can schedule periodic archiving for system logs, system configuration files, and system snapshots. Periodic archiving occurs only at the scheduled time. "Unscheduled" archiving does not occur automatically. For example, if a log file exceeds the maximum file size, the archiving process does not automatically back up the file prior to the scheduled time to prevent data loss.

If the archive process fails, archiving restarts at the next scheduled time. The system does not continue to retry the process if it fails, and log files are not deleted.

We recommend that you schedule an archive operation during hours when traffic is light to minimize its impact on users. The automatic archiving process compresses files and, if the system is busy, can degrade performance for users. Also, a cluster node might appear unresponsive if the system is busy with traffic and performing archiving simultaneously.

> NOTE: If you schedule an archive operation to occur during the hour that your system switches to daylight savings time (DST), the operation might not occur as scheduled. For example, if your system is set to change to DST at 1:00 AM, and you scheduled an archive job to occur at anytime between 1:01 AM and 1:59 AM, then the operation does not take place because at 1:00 AM the system clock is moved forward to 2:00 AM, and the system never reaches the archive time for that date.

To configure log archiving:

1. Select **Maintenance > Archiving > Archiving Servers** to display the configuration page.

   Figure 178 shows the archiving configuration page for Policy Secure.

   Figure 179 shows the configuration page for Connect Secure.

2. Complete the configuration as described in Table 130.

3. Save the configuration.

Figure181: Archiving Configuration Page – Policy Secure

Figure182: Archiving Configuration Page – Connect Secure

Table130: Archiving Configuration Guidelines

Settings	Guidelines
Archive Settings	
Archive Server	Specify the fully qualified domain name or IP address of the server to which to send the archive files.
Destination Directory	Specify the destination directory. Follow these recommendations: • For UNIX systems, you can specify an absolute or relative path. We recommend you specify a full path. • For Windows systems, specify a path that is relative to the ftproot directory. We recommend you specify a full path. Do not include a drive specification for the destination directory, such as: juniper/log.

Username	Specify a username that has privileges to log into the server and write to the destination directory.
Password	Specify the corresponding password.
Method	Select **SCP** or **FTP**. SCP is the default method. SCP is a file transfer utility similar to FTP. SCP encrypts all data during transfer. When the data reaches its destination, it is rendered in its original format. SCP is included in most SSH distributions and is available on all major operating system platforms.
Archive Schedule	
Archive events log	Schedule archiving for the Events log. The archive file has the following format: **JuniperEventsLog-[clustername|standalone]-[nodename|hostname]- [IVSname|root]-[date]-[time]** For example, an archive file for a cluster named Gen has a filename similar to the following: **JuniperEventsLog-Gen-node1-Root-20090109-1545.gz**. The archiving schedule configuration includes the following options: • **Use this filter**—Select a log format filter. • Day of week—Select the days of the week on which to run the archiving job. • **Every hour** or a **Specified Time**. The Every hour option runs a job every hour on the hour for the selected days. The specified time option runs a job once on the selected days. • **Clear log after archiving**. Select this option to clear the local log file after the archiving job is successfully completed. If an archive job fails, the log files are not deleted. • Password—(Optional) Specify a password to secure and encrypt system configuration or user account archives.
Archive user access log	Schedule archiving for the User Access log. The archive file has the following format: **JuniperAccessLog-[clustername|standalone]- [nodename|hostname]-[IVSname|root]-[date]-[time]** The archiving schedule configuration includes the same options as those described for the Events log.
Archive admin access log	Schedule archiving for the Admin Access log. The archive file has the following format: **JuniperAdminLog-[clustername|standalone]- [nodename|hostname]-[IVSname|root]-[date]-[time]** The archiving schedule configuration includes the same options as those described for the Events log.
Archive sensors log	Schedule archiving for the Sensors log. The archive file has the following format: **JuniperSensorsLog-[clustername|standalone]- [nodename|hostname]-[IVSname|root]-[date]-[time]** The archiving schedule configuration includes the same options as those described for the Events log.
Archive client-side log uploads	Schedule archiving for client-side log uploads. This option is available only on Connect Secure. The archiving schedule configuration includes the same options as those described for the Events log, except for log filter format, which is not applicable to the client-side logs.
Archive system configuration	Schedule archiving for the system configuration binary file (system.cfg). The archive file has the following format: **JuniperConf-[clustername|standalone]- [nodename|hostname]-[IVSname|root]-[date]-[time]** The archiving schedule configuration includes the same day, time, and password-protection options as those described for the Events log.

Archive user accounts	Schedule archiving for user account configuration binary file (user.cfg). The archive file has the following format: **JuniperUserAccounts-[clustername|standalone]- [nodename|hostname]-[IVSname|root]-[date]-[time]** The archiving schedule configuration includes the same day, time, and password-protection options as those described for the Events log.
Archive XML configuration	Schedule archiving for the XML configuration files. The archiving schedule configuration includes the same day and time options as those described for the Events log.
Archive debug log	Enable archiving for collected debug logs. You cannot specify a day and time for archiving debug logs. If you select this option, debug logs are archived periodically and cleared if the **Clear log after archiving** option is selected.
Archive periodic snapshots	Enable archiving for snapshots. You cannot specify a day and time for archiving periodic snapshots. If you select this option, snapshots are archived periodically.

Related Documentation

- *Logging Overview*

- Logging Overview

- Configuration File Administration Overview

Archiving Pulse Collaboration Meetings

Connect Secure enables you to archive Pulse Collaboration instances. You can:

- Set up a recurring archival process.

- Perform a one-time archive.

- Archive the deleted meetings into an XML file for later download or deletion. One file is created for each archive run.

- Define the number of days a meeting instance remains on the system before archiving (instances older than x days are archived).

- Define which node in a cluster performs the archive.

The archival process removes completed standalone meetings, completed recurring meeting instances and completed MyMeeting instances. For recurring meeting with end dates already passed, the recurring meetings and their parent meeting are removed. The parent meeting, however, is not archived since the parent meeting information is already captured in the recurring instances. The archival process does not remove meetings in progress or scheduled (future) meetings.

> NOTE: By default, archiving Pulse Collaboration is turned off. Also by default, MyMeetings instances older than 90 days are removed. If the Pulse Collaboration archive feature is turned off, the automatic deletion of MyMeetings is not saved into the archive file.

In a cluster configuration, only one node performs the archival task and only the files stored on that node are archived. You must log in to the archive node using the node IP instead of the virtual IP to download or delete the archived files.

For IVS, you must configure the settings on each IVS you want to archive meetings.

Shown below is an example snippet of an XML file created by the Pulse Collaboration archival process:

```
 <meetings>
  <meeting>
  <id>20993310</id>
  <creator><![CDATA[gary (Users)]]></creator>
  <name><![CDATA[Support Meeting (20993310)]]></name>
  <agenda><![CDATA[]]></agenda>
  <teleconference_info><![CDATA[]]></teleconference_info>
  <date><![CDATA[4:11 PM May 15, 2007 (GMT-08:00) Pacific Time (US &
  Canada); Tijuana]]></date>
  <duration>1 hour</duration>

  <meeting_type>support</meeting_type>
  <invitees>
   <invitee><![CDATA[gary (System Local) Host]]></invitee>
   ...
  </invitees>
  <attendees>
   <attendee>
    <name><![CDATA[gary]]></name>
    <join_time>04:11 PM</join_time>
    <duration>50 minutes </duration>
   </attendee>
   ...
  </attendees>
  ...
  </meeting>
  ...
 </meetings>
```

To archive meetings:

1. In the admin console, choose **Maintenance > Archiving > Pulse Collaboration**.

   Figure 180 shows the archiving page.

2. To schedule a recurring archival process, select the **Perform automatic clean up every** option and specify how often the archiving process should run.

3. In the Delete meetings older than field, enter how old (in days) meetings must be before being archived. Meetings older than this number are archived and removed from the system.

4. To archive meetings in a cluster configuration, select the **Archive meeting records on node** option and then select the node that performs the archive.

5. Click **Clean Up Now** to perform the archive process immediately. Meetings older than the specified age are archived and removed from the system.

6. Click **Save Changes** to save your edits.

Figure183: Pulse Collaboration Archiving Page
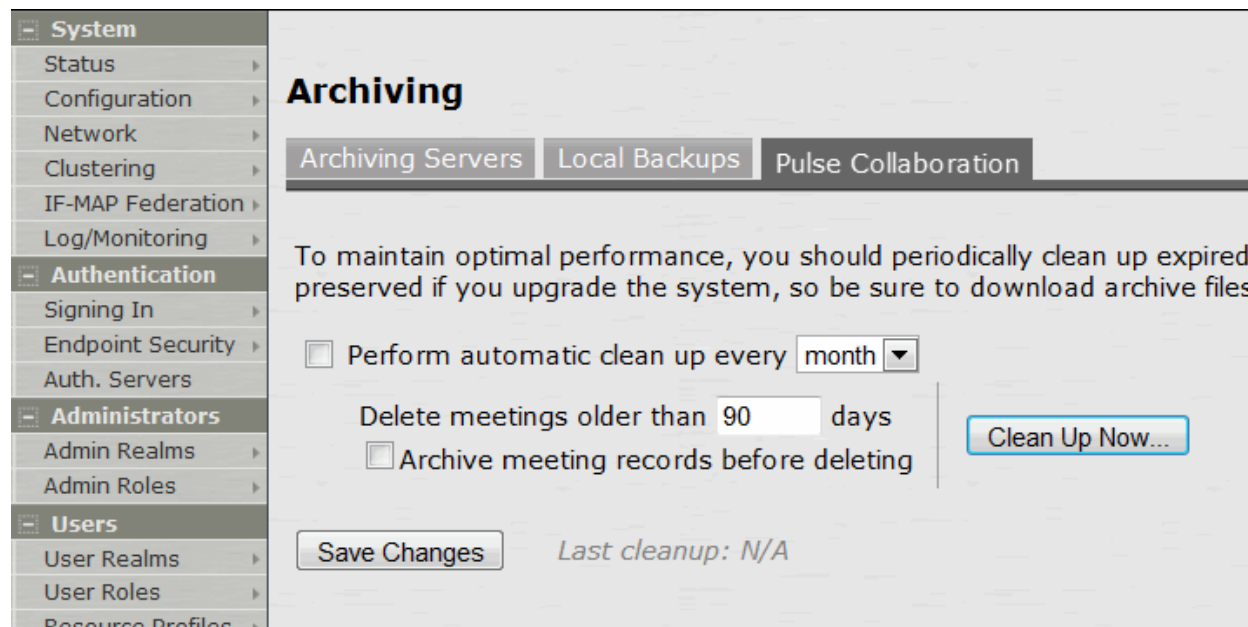
Once the archive process completes, the archive files are listed in the Pulse Collaboration archive table.

To view or download an archive file, click its name.

To delete an archive file, select the check box next to its name and click **Delete**.

Related Documentation

- Pulse Collaboration Overview

Using the Configuration Backup and Restore Feature

You can save up to five system configuration backups and five user account backups on the local server. If you exceed this limit, the system overwrites the oldest backup with the new backup. If you do not want to overwrite the oldest backup, select and delete another backup instead, before you save the most current one.

To manage configuration file backups:

1. Select **Maintenance > Archiving > Local Backups** to display the configuration page.

   Figure 181 shows the archiving configuration page for Policy Secure.

   Figure 182 shows the archiving configuration page for Connect Secure.

2. Use the controls to backup or restore the configuration as described in Table 131.

3. Save the configuration.

Figure184: Local Backups Management Page – Policy Secure
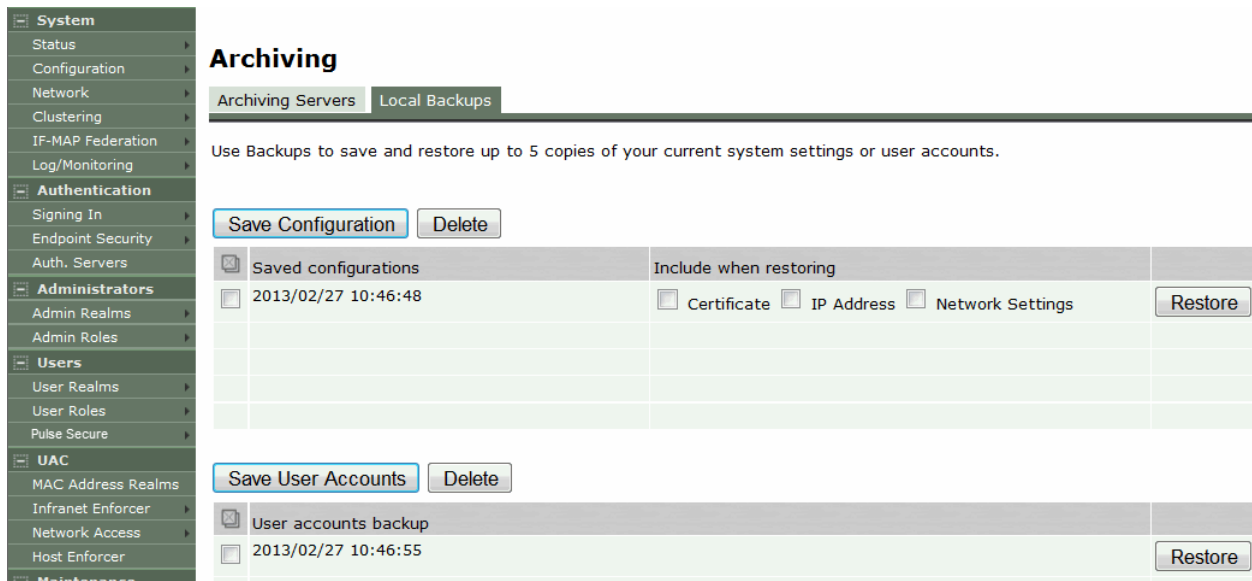
Figure185: Local Backups Management Page – Connect Secure
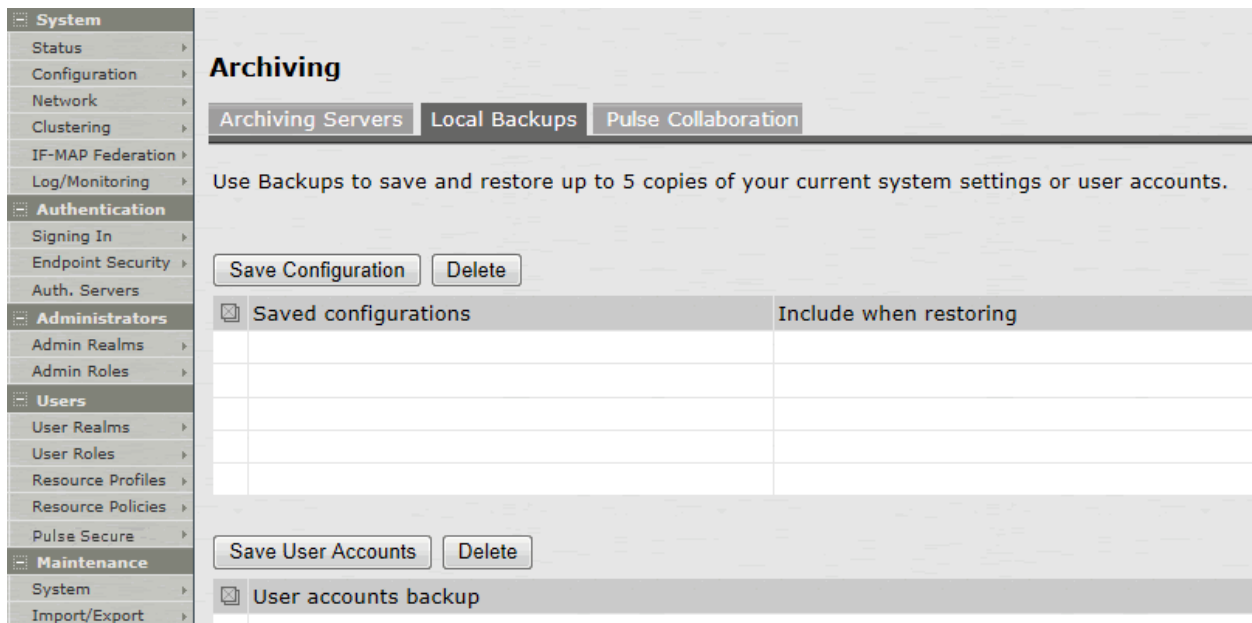
Table131: Local Backups Management Guidelines

Controls	Guidelines
System Configuration	
Save Configuration	Create a backup of the running configuration.
Delete	Select a row in the table and click **Delete** to delete the backup.
Restore	Select a row in the table and components in the "Include when restoring" column and click **Restore** to replace the running configuration with the archived configuration.
User Configuration	

Save Configuration	Create a backup of the running configuration.
Delete	Select a row in the table and click **Delete** to delete the backup.
Restore	Select a row in the table and click **Restore** to replace the running configuration with the archived configuration.

Related Documentation

- Configuration File Administration Overview

Using the Import/Export Feature for Binary System Configuration Files

This topic describes the import/export feature for binary system configuration files. It includes the following information:

- Binary System Configuration File Overview
- Exporting a Binary System Configuration File
- Importing a Binary System Configuration File

Binary System Configuration File Overview

The Pulse Secure access management framework enables you to import and export the system and network settings using binary system configuration files. When importing a system configuration file, you can exclude the device certificate and the server's IP address or network settings from the imported information. For example, to set up multiple Policy Secure or Connect Secure systems behind a load balancer, import everything except for the IP address. To set up the system as a backup server, import everything except for the digital certificate and the network settings.

The binary system configuration file includes the following settings:

- Network settings
- Certificates. The system imports only device certificates, not the chains that correspond to the device certificates or trusted client CAs.
- Cluster configuration
- Licenses. When you import a configuration file that contains licenses, the system gives precedence to any existing licenses. Licenses are imported only if no licenses are currently installed.
- SNMP settings
- Sensor configuration. Sensor configurations are included in the system configuration file while sensor event policies are included in the user configuration file. To import or export all sensor-related settings, import or export both the system and user configuration files. The user configuration file, not the system configuration file, includes resource profiles, resource policies, and the local user database. To perform a complete backup, export both the system and user configuration files.
- Client-side logs. To import or export client-side logs, import or export both the system and user configuration files.
- Mail proxy. Connect Secure only.
- Web proxy servers. Connect Secure only. To export all web proxy related information, both the system and user configuration files are needed.
- Web caching options. Connect Secure only.
- Rewriter filters. Connect Secure only.

NOTE:

- Import of system and user configuration across different hardware platforms is not supported. For example, exporting a system configuration from an SA6000 device and importing it to an SA4500 device is not supported and may result in unexpected behavior.

- You can import a FIPS configuration file into a non-FIPS device and vice versa provided that you do not include the certificate and security world in the import process.

Exporting a Binary System Configuration File

To export a binary system configuration file:

1. Select **Maintenance > Import/Export > Import/Export Configuration** to display the configuration page.

   Figure 183 shows the configuration page for Policy Secure.

   Figure 184 shows the configuration page for Connect Secure.

2. Complete the configuration and export operation as described in Table 132.

Figure186: Export Binary System Configuration File Configuration Page – Policy Secure

Figure187: Export Binary System Configuration File Configuration Page – Connect Secure

Table132: Export Binary System Configuration File Configuration and Action Guidelines

Settings	Guidelines
Password for configuration file	Specify a password to encrypt and secure the configuration file.
Confirm password	Specify the password.
Save Config As	Display a dialog box to save the file to your local host.

Importing a Binary System Configuration File

To import a binary system configuration file:

1. Select **Maintenance > Import/Export > Import/Export Configuration** to display the configuration page.

   Figure 185 shows the configuration page for Policy Secure. The configuration page for Connect Secure is similar.

2. Complete the configuration and import operation as described in Table 133.

Figure188: Import Binary System Configuration File Configuration Page

Table133: Import Binary System Configuration File Configuration and Action Guidelines

Settings	Guidelines
Options	
Import Device Certificate(s)?	Overwrite the existing device certificate(s) with the ones in the imported configuration file. NOTE: When importing a device certificate in to a FIPS device, note that you must choose a certificate that uses a FIPS-compliant private key. To ensure FIPS-compliance, select a certificate and corresponding security world private keys were generated on a FIPS device.
Other Import Options	
Import everything (except Device Certificate(s))	Import all settings except the device certificate.
Import everything but the IP address	Do not overwrite the existing configuration for network interface IP addresses, netmask, default gateway, virtual interfaces, ARP tables, and route tables. Use this option only if the exported configuration file is from a standalone node. TIP: To set up multiple nodes in a cluster behind a load balancer, import everything except the IP address.

Import everything except network settings and licenses	Do not allow the imported configuration to change the existing configuration for settings found in the Network Settings and Licensing sections. With this option, network configurations, licenses, cluster configurations, certificates, defined SNMP settings and syslog configurations are not imported. Always use this option if configuration file was exported from a node that is part of a cluster. TIP: To set up a backup node, import everything except network settings and digital certificates.
Import only Device Certificate(s)	Import the device certificate(s) only. You also must select the **Import Device Certificate(s)** check box.
Config file	Use the browse button to locate and select the file from your local host.
Password	Specify the password (if applicable).
Import Config	Import the configuration file.

Using the Import/Export Feature for Binary User Configuration Files

This topic describes the import/export feature for user configuration binary files. It includes the following information:

- Binary User Configuration File Overview
- Exporting a Binary User Configuration File
- Importing a Binary User Configuration File

Binary User Configuration File Overview

In general, if a menu item falls under the Authentication, Administration, or Users menu, the item is included in the user configuration file (user.cfg). The exception is Sensors event policies, which are under System, but which are exported in the user configuration file. In particular, the user configuration file includes the following settings:

- Sign-in settings (includes sign-in policies, sign-in pages, all authentication servers, authentication protocol sets, Pulse, and OAC settings)
- Authentication realms (including admin realms, user realms, and MAC authentication realms)
- Roles
- Network access. Policy Secure only.
- Infranet Enforcers. Policy Secure only.
- Host Enforcer. Policy Secure only.
- Resource profiles. Connect Secure only.
- Resource policies
- Sensor event policies
- User accounts
- Client-side logs. To export or import client-side logs, export or import both the system and user configuration files.

Exporting a Binary User Configuration File

To export a binary user configuration file:

1. Select **Maintenance > Import/Export > Import/Export Users** to display the configuration page.

   Figure 186 shows the configuration page for Policy Secure.

   Figure 187 shows the configuration page for Connect Secure.

2. Complete the configuration and export operation as described in Table 134.

Figure189: Binary Export User Configuration File Configuration Page – Policy Secure

Figure190: Binary Export User Configuration File Configuration Page – Connect Secure
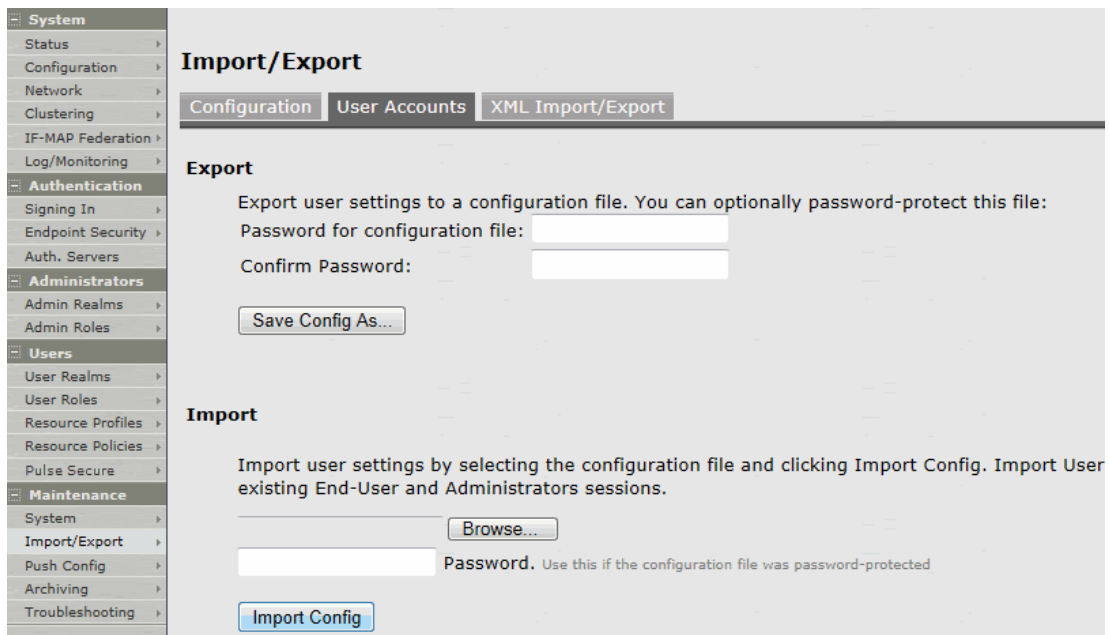
Table134: Binary Export User Configuration File Configuration and Action Guidelines

Settings	Guidelines
Password for configuration file	(Optional) Specify a password to encrypt and secure the configuration file.

Confirm password	Specify the password.
Save Config As	Display a dialog box to save the file to your local host.

Importing a Binary User Configuration File

To import a binary user configuration file:

1. Select **Maintenance > Import/Export > Import/Export Users** to display the configuration page.

   Figure 188 shows the configuration page for Policy Secure. The configuration page for Connect Secure is similar.

2. Complete the configuration and import operation as described in Table 135.

Figure191: Import User Configuration Binary File Configuration Page
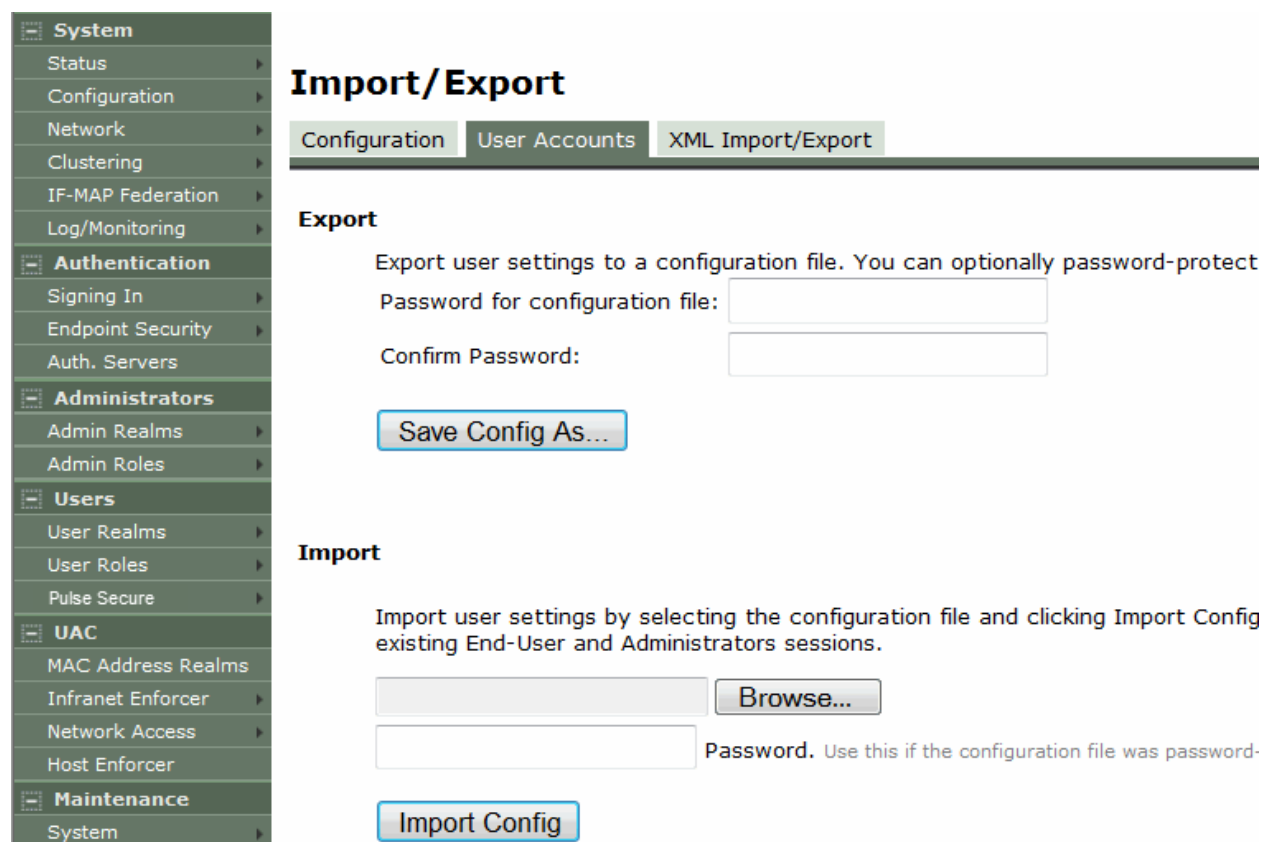
Table135: Import Binary User Configuration File Configuration and Action Guidelines

Settings	Guidelines
Browse	Locate and select the file from your local host.
Password	Specify the password (if applicable).
Import Config	Import the configuration file.

Importing and Exporting IVS Configuration Settings

Exporting IVS configuration saves the following settings in an encrypted file:

- IVS Profiles

- IVS System

- IVS Authentication

- IVS Administrators

- IVS Users

- IVS Resource Policies

- IVS Maintenance

To export IVS settings:

1. In the admin console, choose **Maintenance > Import/Export > Import/Export IVS**.

2. Under Export, enter a password if you'd like to password-protect the configuration file.

3. Click **Save Config As** to save the file.

To import IVS settings:

1. In the admin console, choose **Maintenance > Import/Export > Import/Export IVS**.

2. Browse to the configuration file, which is named ivs.cfg by default.

3. Enter the password you specified for the file. If you did not specify a password before exporting the file, then leave this field blank.

4. Select the **Import IVS Profile Network Settings** option to import references to VLAN ports and virtual ports listed in the imported IVS profiles.

5. Click **Import Config**.

Related Documentation

- Configuration File Administration Overview

Using the Import/Export Feature for XML Configuration Files

This topic describes the import/export feature for XML configuration files. It includes the following information:

- XML Configuration File Overview
- Guidelines and Limitations
- Exporting an XML Configuration File
- Importing an XML Configuration File

XML Configuration File Overview

The system maintains its configuration in a structured XML file. This enables the system to support an alternative to the complete configurations that are exported and imported with the configuration binary files. You can use the export/import configuration XML pages to export and import selected configuration elements.

You might find the feature useful when performing the following tasks:

- Adding to the configurations of peer nodes, for example, adding a large number of users.

- Modifying multiple instances of a single setting, for example, an authentication server name.

- Deleting settings, for example, deleting authentication servers that are no longer used.

- Creating a configuration template to use for setting up new nodes.

- Tracking configuration changes by comparing differences on periodic exports.

Guidelines and Limitations

Table 136 summarizes the guidelines and limitations for using the XML import/export feature.

Table136: XML Import/Export Guidelines and Limitations

Category	Guidelines and Limitations
General	The following guidelines and limitations apply:
	• You can import and export configuration files only between systems running the same software version.
	• You might find it useful to use a text editor to modify configuration elements that ought to be distinguished, such as configuration object names and descriptions. Never modify the names of the NIC identifiers. The system relies on knowing that each appliance has two interface cards, known as NIC0 and NIC1.
	• Immediately after importing an Active Directory authentication server configuration, you must edit the configuration to change the Computer Object name. Unexpected problems might arise if two systems join an Active Directory domain using the same Computer Object name.
Licenses	The following rules apply to exported and imported licenses:
	• You cannot edit the license data that is exported. It is encrypted.
	• An XML import of licenses is valid only if the system does not currently have a license installed. If a license is installed already, any imported licenses are dropped. If you still intend to import a license, you must perform a factory reset before you perform the import operation.
	• If you import a license after deleting a temporary license, the imported license is dropped because you might still be able to reactivate the deleted license. The import operation preserves any licensing data.
Clusters	The following guidelines apply to importing a configuration file for nodes that belong to a cluster:
	• When you perform an import operation on a cluster, all of the cluster nodes must be enabled and running. If you attempt to import a configuration into a cluster in which a node is not running, the import operation might hang or your import results might be unpredictable.
	• The XML configuration that you import must contain the same set of nodes as the original cluster. The signature used to synchronize the cluster when the nodes are reenabled is derived from the IP addresses of the cluster nodes. Therefore, the remaining nodes cannot rejoin the cluster if the imported configuration yields a different signature.
	• When import occurs, the imported configuration file overwrites the node-specific cluster configuration network settings of the remaining nodes. If you change the node-specific network settings, make sure you do not make the remaining nodes unreachable.
	• After you have exported the file, do not modify settings that could render the primary node unreachable, such as changes to network settings.
	• After you have exported the file, do not modify the XML to change the node name, IP address, or IP netmask.
	• After you have exported the file, do not modify virtual port settings or add new virtual port settings.

Exporting an XML Configuration File

To export an XML configuration file:

1. Select **Maintenance > Import/Export > Export XML** to display the configuration page.

   Figure 189 shows the configuration page for Policy Secure.

   Figure 190 shows the configuration page for Connect Secure.

2. Complete the configuration and export operation as described in Table 137.

Figure192: Export XML File Configuration Page – Policy Secure

Figure193: Export XML File Configuration Page – Connect Secure

Table137: Export XML File Configuration and Action Guidelines

Settings	Guidelines
Schema Files	
Schema files	Download the XML schema definition (.xsd) files that describe the XML.
Select Settings and Export	
Expand All	Expand the display of all settings groups.
Select All	Select all settings for all groups.
Export	Export the selected configuration data to an XML file.
Settings	
System	Expand this group and select settings found under the System menu.
	NOTE: Do not select the DMI Agent unless Pulse Secure Technical Support instructs you to do so.
Sign-in	Expand this group and select settings found under the Sign-in menu.
Endpoint Security	Expand this group and select settings found under the Endpoint Security menu.
	NOTE: ESAP and JEDI packages are encrypted when exported.
Authentication Realms	Expand this group and select authentication realm settings, including user realms and MAC address authentication realms.
Roles	Expand this group and select settings found under the Roles menu.
Resource Profiles	Connect Secure only.
	Expand this group and select settings found under the Resource Profiles menu.
Resource Policies	Expand this group and select settings resource policies settings.
Pulse Secure Client	Expand this group and select settings found under the Pulse Secure client menu.
Local User Accounts	Expand this group and select local authentication server settings.
Infranet Enforcer	Policy Secure only.
	Expand this group and select settings found under the Infranet Enforcer menu.
Network Access	Policy Secure only.
Maintenance	Expand this group and select settings found under the Maintenance menu.
Export Settings?	
Export	Export the selected configuration data to an XML file.

Importing an XML Configuration File

To import an XML configuration file:

1. Select **Maintenance > Import/Export > Import XML** to display the configuration page.

   shows the configuration page for Policy Secure. The configuration page for Connect Secure is similar.

2. Complete the configuration and import operation as described in .

Figure194: Import XML File Configuration Page
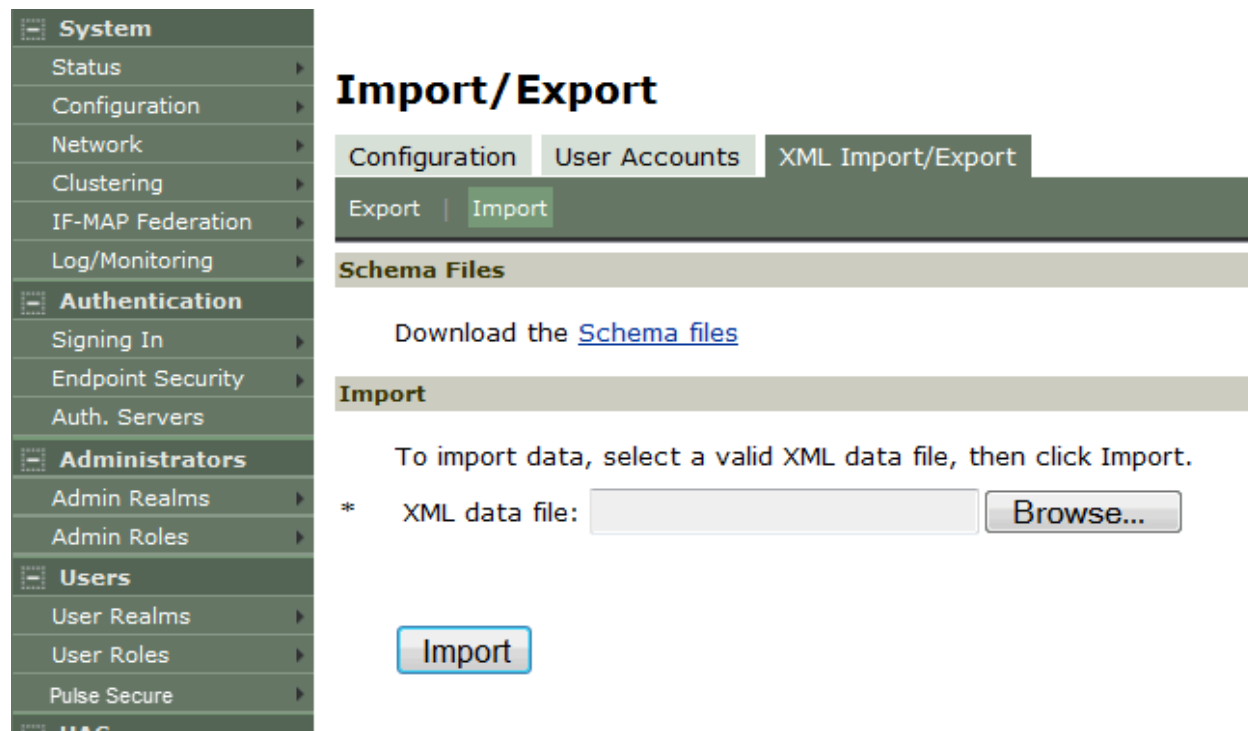
Table138: Import XML File Configuration and Action Guidelines

Settings	Guidelines
Schema Files	
Schema files	Download the XML schema definition (.xsd) files that describe the XML.
Import	
XML data file	Locate and select the XML file.
Import	Import the file. The Import XML Results page is displayed. This page contains information about the imported network settings, roles, resource policies, and other settings. If there are errors in the XML, the import operation stops and rolls back the configuration to the previous state. Error messages are displayed on the Import XML Results page.

Related Documentation

- Example: Using the Configuration XML File Import/Export Feature to Add Multiple Users

- Guidelines for Modifying Configuration XML Files

- Configuration File Administration Overview

Example: Using the Configuration XML File Import/Export Feature to Add Multiple Users

This example shows how to use the configuration XML file import/export feature. The example is illustrative. There are additional ways to use export files.

Assume you have just added a new device to the network, and you want to add your 2,000 users to the system. Instead of adding them one at a time in the admin console, you want to perform a mass import You can export the user accounts, extract the relevant XML that defines

users, replicate each element as needed, and then import them. In this situation, your configuration should include the option to force the users to change their passwords the first time they log in to the system.

In this procedure, you only see examples for User 1, User 2, and User 2000. All other users are included in your import file. You set the passwords to numbered instances of the word password, such as password1, password2, and so on. All users in this example are assigned to the same auth server, although you can specify any combination of auth servers that are valid on your system.

To add multiple new users:

1. Select **Maintenance > Import/Export > Export XML**.

2. Follow the instructions to export local user accounts.

3. Save the exported file as users.xml.

4. Open the users.xml file.

5. Copy and paste the User container element repeatedly until you have added the necessary number of users. Although the example shows only three new users, you might add hundreds of new users to the file.

6. Update the appropriate data in each User container element as shown in "Example: Updating the User container".

7. Save the users.xml file.

8. Select **Maintenance > Import/Export > XML Import/Export > Import**.

9. Click **Browse** to locate and select your users.xml file.

10. Click **Import**.

Example: Updating the User container

```
<configuration xmlns="http://xml.juniper.net/ive-sa/6.2R1"
xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance">
 <authentication>
  <auth-servers>
   <auth-server>
    <local>
     <users>
       <user>
       <username>user1</username>
       <fullname>User1</fullname>
       <password-cleartext>password1

       </password-cleartext>
       <one-time-use>false</one-time-use>
       <enabled>true</enabled>
       <change-password-at-signin>true

       </change-password-at-signin>
      </user>
       <user>
       <username>user2</username>
       <fullname>User2</fullname>
       <password-cleartext>password2

       </password-cleartext>
       <one-time-use>false</one-time-use>
       <enabled>true</enabled>
       <change-password-at-signin>true
```

```
            </change-password-at-signin>
          </user>
          <name>System Local</name>
        </auth-server>
      </auth-servers>
    </authentication>
  </configuration>
```

Related Documentation

- Guidelines for Modifying Configuration XML Files

- Using the Import/Export Feature for XML Configuration Files

Guidelines for Modifying Configuration XML Files

This topic provides guidelines for modifying an exported configuration file. It includes the following information:

- Preparing to Modify a Configuration XML File

- Understanding the XML Export File

- Comparing Configuration Settings and Values Shown in the User Interface with the Ones in the XML File

- Understanding Referential Integrity Constraints

- Using Operation Attributes

Preparing to Modify a Configuration XML File

The following practices might be useful when you export and import XML configuration elements:

- Define your goals for a particular task:

  - What object or objects do you need to add, update, or delete?

  - Do you need to complete all modifications in one operation, or can you modify the configuration in separate operations?

  - Is your process a one-time operation, or do you need to perform the same operation multiple times?

  - Are you updating an active system or are you using one configuration as a template for configuring systems that have not yet been brought online?

- Document the intended changes to the configuration objects:

  - Make a list of objects to be added, updated, or deleted.

  - For objects to add or update, list specific attribute data.

  - List pages or tabs from the admin console that correspond to the objects and attributes you intend to change.

- Make a binary system snapshot or a binary configuration backup immediately before you perform the import.

- Make a plan to verify that the completed configuration meets your goals.

  - View the Admin Access log to make sure the export and import operations succeeded.

  - Perform a random check of the modified items. Make sure items were added, updated, or deleted as you expected.

- Make sure you are able to view configuration details in both the admin console and XML file while you work on the modifications, typically in the following sequence.

  1. Use the admin console to correlate the configuration data with the data in the XML file.

  2. Use the XML file to locate and modify the configuration data.

  3. Use the admin console to verify the successful import.

- Use an XML editor. The exported XML files have a standard structure. Once you become familiar with the structure, you can navigate the files easily. The files can become large, so you might find it more efficient to use a commercial or open source XML editor. XML editors often separate the editable data from the structural display. This separation reduces or eliminates the risk of accidentally modifying an XML element rather than its data.

Understanding the XML Export File

When you export a configuration file, the system saves the configuration as an XML file. The data in the exported file is based on the selections you make when you configure the export operation. The file contains all of the required XML processing instructions and namespace declarations, which must be included exactly as defined.

Table 139 provides some basic information and guidelines to help you understand the structured XML used in the export file.

Table139: Structured XML Files: Basic Information and Guidelines

Topic	Guideline
XML schema definition (.xsd) file	The export is based on an XML schema. The schema is a separate file that defines the metadata, and that serves as a model or a template for the exported file. Use the XML schema file to: • Identify the structure and sequence of configuration objects. • Identify optional and required elements, allowable values, default values, and other attributes of the configuration objects. You can download the XML schema definition (.xsd) file in either of the following ways: • From the XML Import/Export pages by clicking a link. • From the URL where the files are stored on the system (you do not need to sign in). To access the .xsd file, access the following URL: https://<*IP-or-hostname*>/dana-na/xml/config.xsd
Elements	An element is a discrete XML unit that defines an object or part of an object. The element typically consists of a pair of tags that may or may not surround string data. Tags are surrounded by angle brackets (< >).
Namespaces	Namespaces allow you to use the same words or labels in your code from different contexts or XML vocabularies. Prefixing elements with namespace qualifiers allows the XML file to include references to different objects that originate in different XML vocabularies and that share the same name. If you do not prefix elements with namespace qualifiers, the namespace is the default XML namespace, and you refer to element type names in that namespace without a prefix. A namespace declaration looks like the following example: <configuration xmlns="http://xml.juniper.net/ive-sa/6.2R1" xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"> When you see namespace identifiers in your XML files, you do not need to be concerned about them, as long as you do not delete or modify the identifiers.
Element Sequence	You should avoid changing the sequence of elements in the XML file, whenever possible. Although the schema does not enforce sequence in all cases, you gain no benefit from changing the order of elements from the order in which they appear in the exported file, and, in some cases, you might invalidate the XML structure by changing element sequence.

Every XML tag fits into one of the following XML tag types:

- Start tag—Defines the beginning of an element. The start tag consists of an open angle bracket (<), a name, zero or more attributes, and a close angle bracket (>). Every start tag must be followed by an end tag at some point in the document.

- End tag—Defines the end of an element. The end tag consists of an open angle bracket and a forward slash (</), followed by the same name defined in its corresponding start tag, and ends with a close angle bracket (>).

- Empty tag—The empty tag is denoted in two forms. If a tag pair has no data between them, the tag pair is considered an empty tag. Officially, according to the XML specification, an empty tag looks something like this:

      <<empty tag example/>>

In this form, the empty tag consists of an open angle bracket (<), followed by an element name, a slash and a close angle bracket (/>). When you see an empty tag in your configuration files, it signifies an element that the schema requires to be included in the XML file, but whose data is optional.

Start tags can contain attributes, and tag pairs (elements) can contain additional elements. The following example shows an XML file for the Users object. In this example, you see only the Administrator configuration settings.

```
<configuration xmlns="http://xml.juniper.net/ive-sa/6.2R1"
xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance">
    <authentication>
     <auth-servers>
      <auth-server>
       <local>
        <users>
         <user>
        <username>admin</username>
        <fullname>Platform Administrator</fullname>
        <password-encrypted>3u+U</password-encrypted>
        <one-time-use>false</one-time-use>
        <enabled>true</enabled>
        <change-password-at-signin>false</change-password-at-signin>
         </user>
        </users>
       </local>
        <name>Administrators</name>
      </auth-server>
```

You make changes to the string data that is displayed between start and end tags. For example, using the preceding example, you can add to or change the following elements:

- <username>**admin**</username>

- <fullname>**Platform Administrator**</fullname>

- <password-cleartext>**password**</password-cleartext>

- <change-password-at-signin>**false**</change-password-at-signin>

- <name>**Administrators**</name>

NOTE: The preceding sample displays the password element's data as encrypted data. You can modify the password if you change the element to password-cleartext. If you modify the password, the password value is visible until it is imported back into the system. Once imported, the system encrypts the password.

If you enter passwords for new users in cleartext format, the passwords are visible in the file, therefore, you might consider setting the Change Password at Next Login option to true.

> **NOTE:**
>
> - Because passwords are encrypted by default, they are fully portable from one system to another.
>
> - Use the password-cleartext element and enter a text password when changing passwords through the XML file.
>
> If you change a user for a given authentication server or an authentication server for a given user, you are creating a different user, not updating an existing user or authentication server. User and authentication server together logically define a unique user.

Comparing Configuration Settings and Values Shown in the User Interface with the Ones in the XML File

The elements in the XML file are closely related to the objects and their options as you see them in the admin console. The element names in the XML instance file correlate closely with the displayed object and option names.

For example, select **Users > User Roles > [Role] > General > Session Options**. The admin console renders the possible values for a roaming session as an option button group, consisting of the values:

- **Enabled**

- **Limit to subnet**

- **Disabled**

The following snippet from the exported configuration file shows the session options for the Users role. On the bolded line, the roaming session option is disabled:

```
<session-options><SessionOptions>
 <MaxTimeout>60</MaxTimeout>
 <RoamingNetmask />
 <Roaming>disabled</Roaming>
 <PersistentSession>false</PersistentSession>
</SessionOptions>
```

In the schema file, you can locate the allowable values for the roaming session option:

```
<Attribute roaming:START>
<xsd:element name="roaming" minOccurs="0">
...
  <xsd:enumeration value="enabled">
 ...
  <xsd:enumeration value="limit-to-subnet">
 ...
  <xsd:enumeration value="disabled">
...
</xsd:element>
<Attribute roaming:END>
```

To change the value for the roaming session from **Disabled** to **Limit to subnet**, replace **disabled** with **limit-to-subnet**.

This example shows that the admin console often provides all of the allowable values, displayed either in an option button group, as check boxes, as list boxes, or as other types of user interface components. The XML file displays only the current state of your configuration.

The schema file displays all of the actual values for the configuration options that are supported.

Understanding Referential Integrity Constraints

The system configuration objects are part of a data model that is enforced through the use of referential integrity constraints. You cannot change these constraints, but you should understand them before you attempt to delete objects that maintain dependencies to other objects.

If you violate the referential integrity constraints, your import operation fails.

In Figure 192 the boxes represent object types and the arrows represent dependent relationships between the object types. Arrows point from dependent objects to objects.

Figure195: Object Referential Integrity Constraints
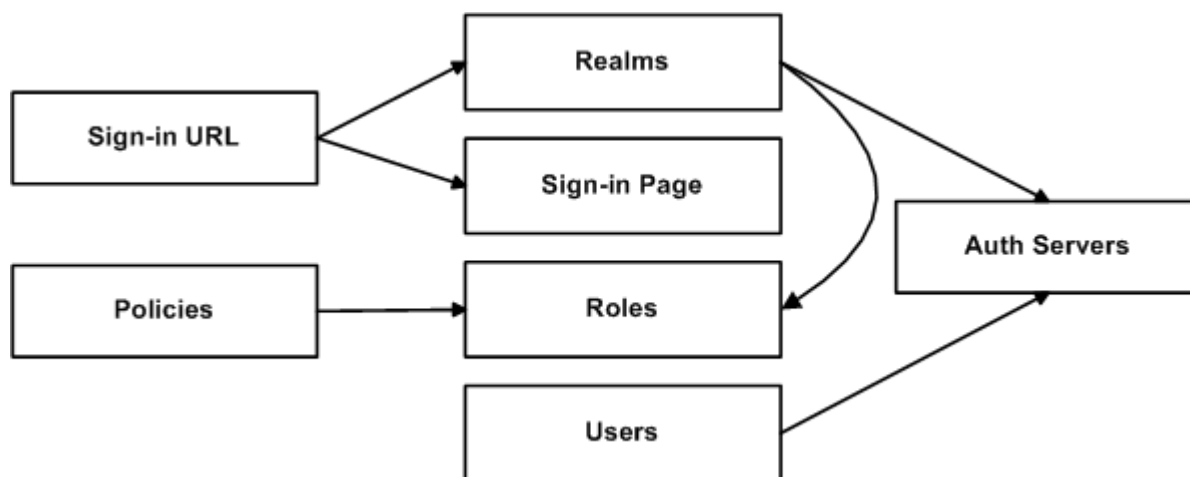
The system does not allow you to delete an object on which another object depends. Conversely, when you add an object, you must add any other objects on which that object depends.

Sign-in URLs depend upon realms and sign-in pages. Realms depend upon both authentication servers and roles. Policies depend upon roles. Users depend upon authentication servers.

Consider the following scenarios based on the preceding figure:

- If you add sign-in URLs, you must add realms, sign-in pages, roles, and authentication servers. You need to add an authentication server and at least one role to support the realm, and you must add the realm and the sign-in page to support the new sign-in URL.

- If you add a user, you must be able to assign it to an authentication server. If there is no authentication server on the target node yet, you must add one in the XML file.

- If you add a policy, you must be able to assign it to a role. If there is no role on the target system, you must add one in the XML file.

- If you delete an authentication server, you might strand realms and users, therefore, you need to make sure no realms or users depend on the authentication server before you attempt to delete it.

- If you delete a role, you might strand policies and realms. To delete a role, you must first delete any policy that depends upon the role, or reassign associated policies to another role. Also, to delete a role, you must first delete or reassign any realm that depends upon that role.

- If you delete a sign-in page, you might strand one or more sign-in URLs. To delete a sign-in page, you must first delete any associated sign-in URLs or reassign them to other sign-in pages.

Referential integrity checks are performed only during XML import.

Using Operation Attributes

Operation attributes define the positioning or action of XML data within the schema. If you do not specify an operation attribute, the modified data is merged by default.

XML data with an operation attribute has the following format:

```
<object1 xc:operation="operator for object1 and its children unless new operator is defined">
    ….
    <object2>
     …
      <object3 xc:operation="operator for object3">
       …
      </object3>
     …
    </object2>
    …
</object1>
```

The operation attribute is applied to all children objects unless a different operation attribute is defined in children objects.

The following operation attributes are supported:

- Merge—The configuration data identified by the element that contains this attribute is merged with the configuration at the corresponding level in the configuration datastore identified by the target parameter. This is the default behavior.

- Replace—The configuration data identified by the element that contains this attribute replaces any related configuration in the configuration datastore identified by the target parameter. Only the configuration actually present in the configuration parameter is affected.

- Create—The configuration data identified by the element that contains this attribute is added to the configuration if and only if the configuration data does not already exist on the device.

- Delete—The configuration data identified by the element that contains this attribute is deleted in the configuration datastore identified by the target parameter.

- Insert before—Changes the position of a configuration element in an ordered set.

- Insert after—Changes the position of a configuration element in an ordered set.

- Rename—Changes the name of one or more of a configuration object's identifiers.

If you are merging a list of objects to an existing list of objects in the configuration store, the results of the merged list might be unexpected. During a merge operation the order of the objects in the new list is not maintained. If you are importing a list of objects and would like to preserve the order of the new list, you should use the replace operation attribute. You can also use insert before or insert after to ensure that you produce the hierarchy that you intended.

Operation attributes are applied to elements recursively unless new operators are also defined within lower-level elements. There are limitations on the legal operator that can be used in child elements without conflict with the parent operator. Table 140 displays the legal operator relationships between parent and child elements.

Table140: Legal Operator Attribute Relationships

Child > V-Parent	Create	Merge	Replace	Delete	Insert before	Insert after	Rename
None	OK	OK	OK	OK	OK	OK	OK
Create	OK	OK	Error	Error	OK	OK	Error
Merge	OK	OK	OK	OK	OK	OK	OK
Replace	Error	OK	OK	Error	OK	OK	Error
Delete	Error	OK	Error	OK	Error	Error	Error
Insert before	OK	OK	OK	OK	OK	OK	OK
Insert after	OK	OK	OK	OK	OK	OK	OK
Rename	OK	OK	OK	OK	OK	OK	OK

The following examples demonstrate the import operation:

Example 1: Set the MTU to 1500 on an interface named "Ethernet0/0" in the running configuration.

```
<interface>
    <name>Ethernet0/0</name>
    <mtu>1500</mtu>
</interface>
```

Example 2: Add an interface named "Ethernet0/0" to the running configuration, replacing any previous interface with that name.

```
<interface xc:operation="replace">
  <name>Ethernet0/0</name>
  <mtu>1500</mtu>
  <address>
   <name>192.0.2.4</name>
   <prefix-length>24</prefix-length>
   </address>
</interface>
```

> **NOTE:**
> The default import modes have the following equivalent attributes on the root object of the configuration tree:
>
> - Standard Import is always a merge operation.
> - Quick Import is a create operation.
> - Full Import is a replace operation.

Related Documentation

- Using the Import/Export Feature for XML Configuration Files
- Example: Using the Configuration XML File Import/Export Feature to Add Multiple Users

Using the Push Configuration Feature

This topic describes the push configuration feature. It includes the following information:

- Push Configuration Overview
- Guidelines and Limitations

Push Configuration Overview

The push configuration feature supports simple configuration management across an enterprise without requiring you to deploy the systems as a cluster. You push a partial configuration from the running configuration on the source system to the running configuration on one or more target systems.

It is not desirable to push the some groups of settings to a running configuration, so the following groups of settings are not supported:

- Network configurations

- Licenses

- Cluster configurations

- Certificates

- SNMP settings

- Syslog server settings

- Push configuration targets

Guidelines and Limitations

Table 141 summarizes the guidelines and limitations for using the push configuration feature.

Table141: Push Configuration Guidelines and Limitations

Category	Guidelines and Limitations
General	The following guidelines and limitations apply: - You can push a configuration to systems running the same software version (same build number) or higher software version. - The source device pushes data over the management port (if configured) or the internal port. The target device can receive data over the internal or external port or management port. - You can push to a single target or to multiple targets. For example, if you install several new systems, you can push a common configuration to set their initial configuration. - When a configuration push job begins on a target, no warning is displayed, and the administrators are automatically logged out to avoid potential conflicts. - For selected configuration push, if the configuration to be pushed contains one or more JAM packages, we recommend you to push the JAM packages first considering one JAM package per push and then push the remaining configurations. - When the job has completed on a target, the target device restarts its services. Brief interrupts might occur while the service restarts. You must push to targets when they are idle or when you can accommodate brief interruptions. - Immediately after pushing an Active Directory authentication server configuration, you must edit the configuration to change the Computer Object name. Unexpected problems might arise if two systems join an Active Directory domain using the same Computer Object name. - You must delete the failed push jobs before performing a new push. - When performing an entire configuration push, with changes to settings such as security settings, the target web server might get restarted during configuration import. As a result, the source might experience a lost connection. You can resume the job from the source at a later point to see the result of the import.

Licenses	The push configuration job does not push licenses or licensing settings.
Clusters	• You can push a configuration to target that is a member of a cluster, as long as the target is not a member of the same cluster as the source.
	• You can push a configuration to multiple targets, as long as targets are not part of the same cluster.
	• You must not perform the clustering operations such as adding a cluster, deleting a cluster, and so on when performing a push configuration. If such events occur then unsuccessful jobs will be aborted and the back up files will be deleted.
	• You must not use VIPs during push configuration. Instead you must use the internal IP or the management IP of one of the nodes to create the target.
	• You must delete the backed up configuration on the target node(s) as soon as possible to free up the disk space.

Configuring Targets

To configure push configuration targets:

1. Select **Maintenance > Push Config > Targets** to display the target list and source options configuration page.

   Figure 193 shows the configuration page for Policy Secure.

   Figure 194 shows the configuration page for Connect Secure.

2. Complete the configuration for the source options as described in Table 142.

3. Click **New Target** to display the configuration page for targets.

   Figure 195 shows the configuration page for Policy Secure. The configuration page for Connect Secure is similar.

4. Complete the configuration as described in Table 143.

5. Save the configuration.

Figure196: Push Configuration Target List and Source Device Settings Page – Policy Secure

Figure197: Push Configuration Target List and Source Device Settings Page – Connect Secure

Targets

| Push Configuration | Targets | Results | History |

☑ Allow this Pulse Connect Secure to be a target
·This Pulse Connect Secure will accept configuration pushed from another Pulse Connect Secure . This box must be checked on each target that is configured on this page.

☐ Validate target server certificate
To be configured on source device. This enables source device to validate target device server certificate before pushing configuration.

| New Target... | Delete... | Save Changes |

	Target Name	Target Sign-in URL	Admin Username	Auth. Realm
☐	n156	https://10.204.50.156/admin	push	Admin Users
☐	n16	https://10.204.51.16/admin	push	Admin Users

Table142: Push Configuration Target Source Device Configuration Options

Settings	Guidelines
Allow this system to be a target	Select this option to allow the system to accept configuration pushed from another system. This option must be selected on targets, but does not have to be selected on the source system.
Validate target server certificate	Select this option on the source system if you want the source system to validate the target system server certificate before pushing the configuration.
Save Changes	Click this button if you have changed the source device configuration options described above.
Delete	Select a row in the table and click **Delete** to remove the target from the list. You cannot delete a target if it has push configuration results associated with that target.

Figure198: Push Configuration Targets Configuration Page

Targets >
n156

Name: *	n156
Sign-in URL: *	https://10.204.50.156/admin
Admin Username: *	push
Password: *	••••••••
Auth. Realm: *	Admin Users

Save changes?

| Save Changes |

* indicates required field

Table143: Push Configuration Targets Configuration Guidelines

Settings	Guidelines
Name	Specify a name to identify the target within the system. Target names and target sign-in URLs cannot be edited after they have been saved.
Sign-in URL	Specify the URL for the administrator sign-in page. Sign-in URLs cannot be edited after they have been saved.

Admin Username	Specify an account on the target system that the push configuration job can use to sign-in and make changes to the configuration. The job can make wide-ranging configuration changes, so the user must have full administrative privileges. In other words, the user must belong to the .Administrators role.
Password	Specify the corresponding password.
Auth. Realm	Specify the administrator authentication realm on the target system. The access management framework must be configured so that the job process (run as the username specified above) can sign in without any human interaction. For example, you cannot have dynamic credentials or multiple roles that are not merged, as these both require manual interaction. We recommend that you create an administrator account on each target that can be used exclusively for push configuration. Configure the administrator realm so that the realm policy and role mapping rules do not result in prompts requiring human interaction. For example, the user must be able to log in with static password authentication or two-factor tokens that do not use challenge-response type authentication. For example, certificates, Soft ID, and Defender Authentication are not supported.

Configuring Push Settings

To configure the settings to be pushed:

1. Select **Maintenance > Push Config > Push Configuration** to display the configuration page.

   Figure 196 and Figure 197 shows the configuration page for Connect Secure. The configuration page for Policy Secure is similar.

2. Complete the configuration and push configuration operation as described in Table 144.

Figure199: Push Configuration Selected Settings Page
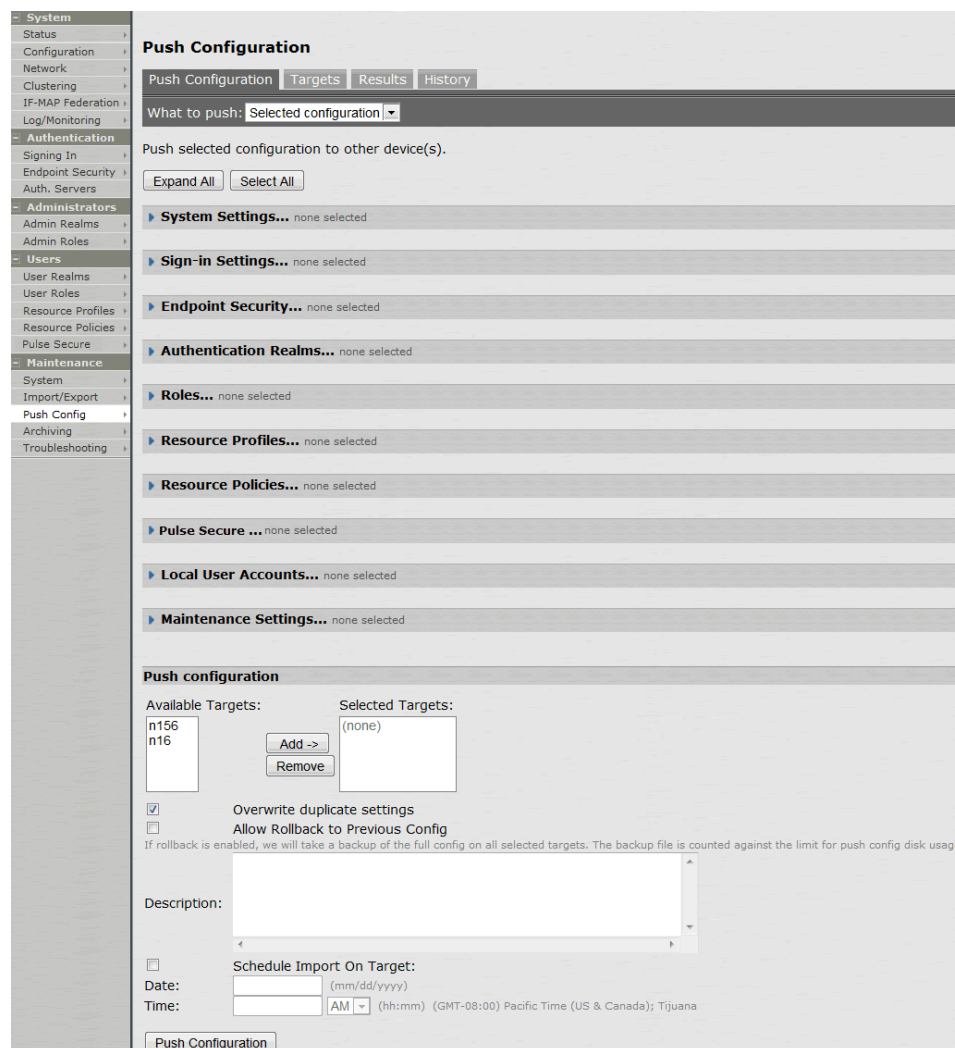
Figure200: Push Configuration Selected Settings Page
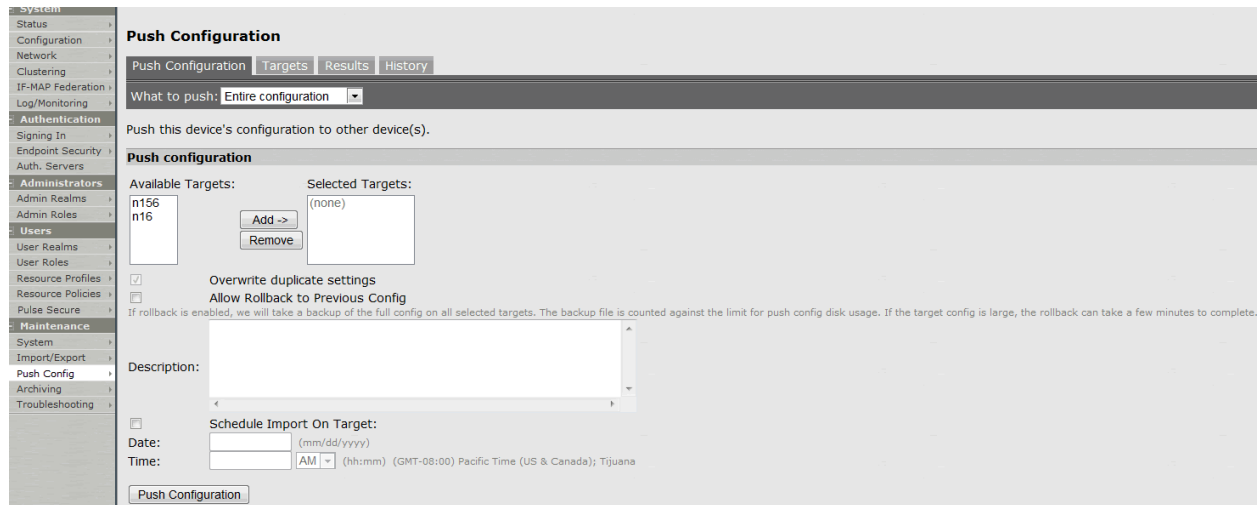
Table144: Push Configuration Selected Settings and Action Guidelines

Settings	Guidelines
Select Settings and Export	
What to push	Select **Selected configuration** or **Entire configuration**.
	If you select **Selected configuration**, the page displays controls to select settings groups.
	If you select **Entire configuration**, all settings from the source system are pushed, except for the following:
	• Network configurations
	• Licenses
	• Cluster configurations
	• Certificates
	• SNMP settings
	• Syslog server settings
	• Push configuration targets
Expand All	Click this button to expand the display of all settings groups.
Select All	Click this button to select all settings for all groups.
Settings	
System	Expand this group and select settings found under the System menu.
	NOTE: You cannot push host-specific network settings to a target. If you want to copy these settings to another system, use the configuration XML file import/export feature.
Sign-in	Expand this group and select settings found under the Sign-in menu.
Endpoint Security	Expand this group and select settings found under the Endpoint Security menu.
	NOTE: ESAP and JEDI packages are encrypted when exported.
Authentication Realms	Expand this group and select authentication realm settings, including user realms and MAC address authentication realms.
Roles	Expand this group and select settings found under the Roles menu.
Resource Profiles	Connect Secure only.
	Expand this group and select settings resource profiles settings.

Resource Policies	Expand this group and select settings resource policies settings.
Pulse Secure Client	Expand this group and select settings found under the Pulse Secure menu.
Local User Accounts	Expand this group and select local authentication server settings.
Infranet Enforcer	Policy Secure only. Expand this group and select settings found under the Infranet Enforcer menu.
Network Access	Policy Secure only. Expand this group and select settings found under the Network Access menu.
Maintenance	Expand this group and select settings found under the Maintenance menu.
Push Configuration	
Available Targets / Selected Targets	Use the **Add** and **Remove** buttons to select the targets.
Overwrite duplicate settings	Select this option to overwrite settings on the target that have the same name as settings being pushed. If you do not select this option, the push configuration job copies only configuration objects that have names different from the configuration objects on the target.
Allow Rollback to previous configuration	Select this option to revert to a previous configuration state, effectively rolling back configuration changes. If you select this option, the local configurations on the target node will be backed up before importing the configurations. You can also undo the push configuration if you want to discard the changes and revert back to the previous state. We recommend you to delete the backed up configuration if the import is successful. NOTE: If the target configuration is large the rollback of configurations can take several minutes to complete.
Description	Enter the description for the job. The job description is limited to 100 characters.
Schedule Import on Target	Select this option to allow a delayed import on the target node. If you select this option the selection applies to all the targets in the job. The import schedule is measured in HH:MM (hours, minutes) format. The schedule is specified according to source's timezone.
Push Configuration	Click this button to push the selected configuration data to the specified targets. You can pause the push for a target during the push process. If errors occur during the push process, the job stops, and the configuration for the target is not imported. However, you can resume the failed push jobs. Error messages are displayed on the Results page. If you have specified multiple targets and a push configuration job to a target fails, the job continues to the next target until specified targets are updated (or fail). The results page displays the status and any problems encountered during the process.

Viewing Configuration Push Results

Purpose The source system saves and displays the push configuration results in the Results tab.

Action To view push configuration job results:

1. Select **Maintenance > Push Config > Results** to display the results page.

   Figure 198 shows the results page for Connect Secure. The page for Policy Secure is similar. The push configuration results page auto refreshes for every 30 seconds.

2. Examine the results to verify success or learn the reasons the push job failed.

   Click the job name to display additional information about the job.

   Select a job and click **Delete** to remove it from the results page.

Figure201: Push Configuration Results Page

Table 145 describes the information displayed on the Results page and the various management tasks you can perform.

Table145: Push Configuration Results

GUI Element	Guidelines
Disk Usage Details	Displays the disk space available for push configuration and the disk space consumed by all the push jobs in the device.
	The disk space consumed by individual push jobs are also mentioned across each push job under the disk usage column. When total disk space consumed reaches the total disk space push jobs may fail and you can see the results column to see the failure message. You need to monitor the disk space consumed by push configuration to avoid push failures related to disk space limits.
Description Column	Displays the type of the push configuration.
Disk Usage Column	Displays the disk space used by the job.
Results Column	Displays the status of the transfer and result of post push action. It also displays the status of the push such as login, export, transfer, backup , import and so on. The status result message shows the type of data that is getting transferred. For a paused or failed target the information on the current state of the job when it is paused, or failure reasons if any is displayed. This column also shows the progress of data transfer using a bar chart. For selected push additional configuration data (additional configuration data refers to configuration that is transferred only if it is modified or not available on the target) includes ESAP package, JAM package, VDI configurations, Terminal services, Host Checker files, Custom sign in pages and notifications, and Applet files. For complete configuration push additional data includes ESAP and JAM packages.
Post Push Action	Displays the options that the user can perform after the push such as roll back and delete backup. It also displays the post push actions such as rollback done, backup deleted, rollback failed, performing rollback, deleting back up and so on.
Resume	Select this option to resume a paused or a failed push.
Undo	Select this option to rollback to previous configuration that was backed up. Note that you can perform this operation only when the push is successful and Allow Rollback to Previous Configuration is selected. This option is available only if the backup is not deleted or undo is not done yet.

Abort	Select this option to cancel an entire push job or push to particular target within a job. An aborted push cannot be resumed.
Pause	Select this option to temporarily pause the push operation to a specified target.
Delete Backup	Select this option to delete the backup configuration on the specified target. Note that this option is available only when the users selects the Allow Rollback to Previous Configuration option during the push job.

Viewing Configuration Push History

Purpose   The source/target system saves and displays up to 5 push history results per target/source in the History tab. When the history table reaches 5 entries, the system removes the oldest result data when the next push configuration job is started.

Action   To view push configuration push history:

1. Select **Maintenance > Push Config > History** to display the history page.

   Figure 199 shows the history page for Connect Secure. The page for Policy Secure is similar.

2. Examine the history to verify success or learn the reasons the push job failed. The history page displays rollback history however the failure reason is not displayed. You can check the failure reason in the details page for each job. It also displays the timestamp history information of successful, failed push jobs, or if a configuration is undone.

   Select the source name and click **Delete History** to remove it from the History page.

Figure202: Push Configuration History Page

Related Documentation

- Configuration File Administration Overview

CHAPTER 36 System Maintenance

Using the System Maintenance Pages

You can use the System > Maintenance pages to perform the following tasks:

- Enable system maintenance options, such as software version monitoring and disk clean-up.

- Upgrade, downgrade, or rollback the system software.

- Download client installer files so that you can distribute them in out-of-band methods to end users.

- Test network connectivity between the system and servers that have been configured to be used with it.

- Display hardware status.

Related Documentation

- Configuring System Maintenance Options
- Upgrading the System Software
- Downloading Client Installer Files
- Restarting, Rebooting, and Shutting Down the System
- Testing Network Connectivity
- Displaying Hardware Status

Configuring System Maintenance Options

You can use the maintenance options page to enable various system maintenance features.

To enable various system maintenance features:

1. Select **Maintenance > System > Options** to display the maintenance options page.

2. Select options as described in Table 146.

3. Save the configuration.

Table146: System Maintenance Options Configuration Guidelines

Options	Guidelines

Automatic version monitoring	If you enable this option, the system reports to Pulse Secure the following data:
	
	• Machine identifier.
	• Information describing your current software, including:
	  ◦ Software build number and build name.
	  ◦ An MD5 hash of your license settings.
	  ◦ An MD5 hash of the internal interface IP address.
	  ◦ If this node is in a cluster, the number of nodes within that cluster.
	  ◦ Current state of the node.
	  ◦ Cluster type (active/active, active/passive).
	  ◦ Total number of unique subnets on the cluster nodes.
	  ◦ Version of Pulse Secure client.
	  ◦ Version of ESAP.
	  ◦ Cluster log synchronization status.
	  ◦ Number of Infranet Enforcers configured (Policy Secure only.)
	  ◦ Names of Radius vendors configured (Policy Secure only.)
	• Total number of concurrent users on the device.
	• Number of Network Connect and Pulse tunnels.
	
	We strongly recommend that you enable this service.

Gzip compression	Connect Secure only. Use gzip compression to reduce the amount of data sent to browsers that support HTTP compression. This can result in faster page downloads for some users.
Kernel Watchdog	Enables the kernel watchdog that automatically restarts the system under kernel deadlock or when kernel runs low on some key resources. NOTE: Enable the kernel watchdog only when instructed by Pulse Secure Technical Support.
File System Auto-clean	Enables the system to automatically clean up the file system when disk utilization reaches 90%. NOTE: The clean-up operation deletes files that might be relevant in debugging—for example, debug logs, core files, and snapshots might be deleted.
Web installation and automatic upgrade of Pulse Secure Clients	After you deploy Pulse Secure client software to endpoints, software updates occur automatically. A Pulse client can receive updates from the server. If you upgrade the Pulse software on your Pulse server, updated software components are pushed to a client the next time it connects. A bound endpoint receives connection set options and connections from its binding server, but it can have its Pulse client software upgraded from any Pulse server that has the automatic upgrade option enabled. During a client software upgrade the client loses connectivity temporarily.
Hardware acceleration	Enables hardware acceleration (if available on the device) to offload cryptographic operations from the main CPU. This can significantly improve performance.
Virtual Terminal console	Enables the virtual terminal on a virtual appliance. Clear this check box to use the serial console. Changing this setting will restart the system.
Java instrumentation caching	Connect Secure only. Caches the Java instrumentation to improve the performance of Java applications.
Show Auto-allow	Connect Secure only. The auto-allow option provides the means to automatically add bookmarks for a given role to an access control policy, for example, Web bookmarks with auto-allow set are added to the Web access control policy. You only use this feature if you also use Resource Policies. We recommend that you use Resource Profiles instead.

End-user Localization	Select one of the following options:
	• Automatic (based on browser settings) • English (U.S.) • Chinese (Simplified) • Chinese (Traditional) • French • German • Japanese • Korean • Spanish
External User Records Management	
Persistent user records limit	Specify the maximum number of user records. This feature is useful when system performance is affected due to a large number of user records. We highly recommend you consult Pulse Secure Technical Support prior to using this feature. Deleting a user record removes all persistent cookies, SSO information, and other resources for that user. It does not remove the user record from the external or internal authentication server. If you delete a user record and that user logs back in to the authentication server, new user records are created. Records are not removed if that user is currently logged in.
Number of records to delete when the limit is exceeded	Specify a number. Older records are removed first. A user record is not deleted if that user is currently logged in.
Delete records now	Check whether the persistent user records limit has been exceeded. If it is, delete the number of user records specified in the option above.
Automatic deletion of user records during new user logins	Check whether the persistent user records limit will be exceeded whenever a new user record is about to be created. If true, delete the records prior to creating the user new record.

Related Documentation

- Using the System Maintenance Pages

Upgrading the System Software

This topic describes how to upgrade, downgrade, and rollback the system software. It includes the following information:

- Downloading a Software Package
- Uploading a Software Package
- Upgrading the System Software
- Downgrading the System Software
- Rolling Back the System Software

Downloading a Software Package

To download a software package:

1. Go to http://www.juniper.net/support/downloads/ and browse to the software download page for your product.

2. When prompted, log in with your Pulse SecureP customer username and password.

3. Accept the license agreement.

4. When prompted, save the software package to your local host.

Uploading a Software Package

You can upload a software package to the system without immediately initiating the upgrade process. This is known as staging the upgrade. You can stage one package. Uploading a second package overwrites the previous staging.

To upload a software package:

1. Select **Maintenance > System > Upgrade/Downgrade** to display the system software maintenance page.

   Figure 200 shows the system software maintenance page for Policy Secure. The system software maintenance page for Connect Secure is similar.

2. Under Managed Staged Service Package, select **Upload new package into staging area** and use the **Browse** button to locate and select the service package file.

3. Click **Submit** to upload the file.

The Upload Status window shows the progress of the upload operation.

Figure203: Software Upgrade Page

NOTE: If you have enabled logging for Administrator changes (System > Log/Monitoring > Admin Access > Settings page), a log is written to the Admin Access logs page.

Upgrading the System Software

Installing a service package can take several minutes and requires the system to reboot. Because existing system data is backed up during this process, you can decrease installation time by clearing your system log before trying to install a service package.

To upgrade the operating system:

1. Select **Maintenance > System > Upgrade/Downgrade** to display the system software maintenance page.

   Figure 200 shows the system software maintenance page.

2. Under Install Service Package, select one of the following options to proceed:

   - **From File**—Use the **Browse** button to locate and select the service package file.

   - **From Staged Package**—Select the service package file that was previously uploaded.

     NOTE: Do not select the **Deletes** option when you are upgrading software. The Deletes option is available to support downgrading software.

3. Click **Install**.

The system displays the Service Package Installation Status page, which provides a summary of the integrity checks and compatibility checks and other status indicators. Figure 201 shows the software upgrade status page.

Figure204: Software Upgrade Status Page

NOTE: If you have enabled logging for Administrator changes (System > Log/Monitoring > Admin Access > Settings page), a log is written to the Admin Access logs page. If you have enabled logging for System Status (System > Log/Monitoring > Events > Settings page), logs are written to the Events logs page.

Downgrading the System Software

If necessary, you can downgrade to an earlier version of the system software. When you downgrade, you must clear the system and configuration data to avoid unexpected behavior that can occur when the system has data that relates to the newer software.

If you downgrade the system, you must reestablish network connectivity before you can reconfigure it.

To downgrade the operating system:

1. Select **Maintenance > System > Upgrade/Downgrade** to display the system software maintenance page.

   Figure 200 shows the system software maintenance page.

2. Under Install Service Package, select one of the following options to proceed:

   - **From File**—Use the **Browse** button to locate and select the service package file.

   - **From Staged Package**—Select a service package file that was previously uploaded.

3. Select the **Deletes** option to delete all system and user configuration data before installing the service package, restoring the member to an unconfigured state.

4. Click **Install**.

Rolling Back the System Software

If necessary, you can rollback the system to the previous software version and configuration state. The system is rebooted and unavailable for a few minutes when you rollback.

To rollback the operating system:

1. Select **Maintenance > System > Platform** to display the system maintenance platform page.

   Figure 202 shows the system maintenance platform page for Policy Secure. The system maintenance platform page for Connect Secure is similar.

2. Click **Rollback**.

Figure205: System Maintenance Platform Page

NOTE: The rollback option appears only if you have previously upgraded the system software.

NOTE: If you have enabled logging for System Status (System > Log/Monitoring > Events > Settings page), logs are written to the Events logs page.

Related Documentation

- Using the System Maintenance Pages

Downloading Client Installer Files

You can use the system maintenance client installers page to download client installer files. The downloadable files include .exe and .msi files for use installing clients on Windows platforms, and .dmg files for installing clients on Macintosh platforms.

To download client installer files:

1. Select **Maintenance > System > Installers** to display the client installer files page.

   Figure 203 shows the client installer files for Policy Secure.

   Figure 204 shows the client installer files for Connect Secure.

2. Click **Download** to download the file to your local host.

Figure206: System Maintenance Client Installers Page – Policy Secure

Figure207: System Maintenance Client Installers Page – Connect Secure

Related Documentation

- Using the System Maintenance Pages

Restarting, Rebooting, and Shutting Down the System

You can use the admin console to perform restart, reboot, and shut down operations. The following items explain these terms:

- Restart—Kills all processes and restarts the system. The system is available again after a few minutes.

- Reboot—Power cycles and reboots the system. The system is available again after a few minutes.

- Shut Down—Shuts down the system. The system is not available again until the physical power button on the physical device is used to restart the system.

> NOTE: The restart, reboot, and shutdown operations are applied to all enabled members of a cluster. If you do not want to apply the operations to all members of the cluster, use the System > Clustering > Status page to disable members; then perform the restart, reboot, or shut down operation.

To restart, reboot, or shut down the system:

1. Select **Maintenance > System > Platform** to display the system maintenance platform page.

   Figure 205 shows the system maintenance platform page for Policy Secure. The system maintenance platform page for Connect Secure is similar.

2. Click the desired node operation:

   - **Restart Services**

   - **Reboot**

   - **Shut Down**

Figure208: System Maintenance Platform Page

NOTE: If you have enabled logging for Administrator changes (System > Log/Monitoring > Admin Access > Settings page), a log is written to the Admin Access logs page. If you have enabled logging for System Status (System > Log/Monitoring > Events > Settings page), logs are written to the Events logs page.

Related Documentation

- Using the System Maintenance Pages

Testing Network Connectivity

You can use the admin console to test network connectivity to all the servers with which the system is configured to communicate, for example network services or AAA servers.

To test network connectivity:

1. Select **Maintenance > System > Platform** to display the system maintenance platform page.

   Figure 206 shows the system maintenance platform page for Policy Secure. The system maintenance platform page for Connect Secure is similar.

2. Click **Test Connectivity**.

   Server connectivity results are highlighted in the figure.

Figure209: System Maintenance Platform Page

Related Documentation

- Using the System Maintenance Pages

CHAPTER 37 Logging and Monitoring

Logging Overview

The system generates event logs related to system performance, administrator actions, network communications, access management framework results, user sessions, and so forth. The system supports the following log collection methods:

- Local log collector and log viewer.

- Reporting to syslog servers.

- Reporting to SNMP servers.

Table 147 describes the event log severity levels.

Table147: Event Log Severity Levels

Severity Level	Description
Critical (level 10)	The system cannot serve user and administrator requests or loses functionality to a majority of subsystems.
Major (levels 8-9)	The system loses functionality in one or more subsystems, but users can still access the system for other access mechanisms.
Minor (levels 5-7)	The system encounters an error that does not correspond to a major failure in a subsystem. Minor events generally correspond to individual request failures.
Info (levels 1-4)	The system writes an informational event to the log when a user makes a request or when an administrator makes a modification.

In addition to managing system logs, you can use the admin console to configure collection of client-side logs, including:

- Host checker

- Meetings

- Windows Secure Application Manager

- Java Secure Application Manager and Applet Rewriting

- VPN Tunneling

- Terminal Services (Connect Secure only)

- Virtual Desktops (Connect Secure only)

Related Documentation

- Displaying System Logs

- Configuring Events to Log

- Configuring Syslog

- Configuring SNMP

- Enabling Client-Side Logging

- Configuring Archiving for System Logs, Configuration Files, and Snapshots

Configuring Events to Log

To configure log event categories:

1. Select **System > Log/Monitoring**.

2. Click the **Settings** tab to display the configuration page.

   Figure 207 shows the configuration page.

3. Complete the configuration as described in Table 148.

4. Save the configuration.

   NOTE: To configure log events for each local log category, you must perform this
   procedure on each local log tab: Events, User Access, Admin Access, and Sensors.

Figure210: Log Events Settings Configuration Page

Table148: Log Events Settings

Settings	Guidelines
Maximum Log Size	
Max Log Size	Specify the maximum size of the local log. The default is 200 MB. The maximum is 500 MB. The default is a good choice for logs formatted with the Standard format. If you use a more verbose format, such as WELF, specify a larger value.
	When the local log reaches the maximum log size, the current data is rolled over to a backup log file. A new, empty, file is then created for all subsequent (new) log messages. The log viewer displays the most recent 5000 log messages (the display limit). If the current log file contains fewer than 5000 log messages, older log messages from the backup log file can be displayed, up to a total of 5000 log messages. This makes the log files appear as one, even though they are stored separately.
	When you save the log messages or use the FTP archive function, the backup log file is appended to the current log file and is then downloaded as one log file. If the log files are not archived or saved by the time they are rolled over again, the oldest log messages (saved in the backup log file) are lost.
Archiving	Click the **Archiving** link to display the configuration page for Archiving jobs, including log archiving.
Select Events to Log – Events Tab	

Connection Requests	Log events related to connection requests.
System Status	Log events related to changes in system status.
Rewrite	Log events related to rewrite policies.
System Errors	Log events related to system errors.
Statistics	Log user access statistics reported on the System > Log/Monitoring > Statistics tab. If you unselect the Statistics option, the statistics are not written to the log file, but are still reported on the statistics page.
Performance	Log events related to performance, such as CPU utilization.
License Protocol Events	Log events related to licensing.
Reverse Proxy	Logs events related to reverse proxy information.
Select Events to Log – User Access Tab	
Login/logout	Log events related to sign in and sign out.
SAM/Java	Log events related to user access to SAM/Java in the local log file.
User Settings	Log events related to changes to user settings in the local log file.
Meeting Events	Log events related to meeting information.
Client Certificate	Log events related to certificate security.
IF-MAP Client User Messages	Log events related to IF-MAP.
Pulse Client Messages	Log events related to Pulse clients.
Web Requests	Log events related to user access to web.
File Requests	Log events related to user access to files.
Meeting	Log events related to user access to meetings.
Secure Terminal	Log events related to user access to secure terminal.
VPN Tunneling	Log events related to user access to VPN tunneling.
SAML	Log events related to user access to SAML.
Select Events to Log – Admin Access Tab	
Administrator changes	Log events related to configuration changes.
Administrator logins	Log events related to administrator access.
License changes	Log events related to licensing.
Select Events to Log – Sensor Tab	
Max Log Size (MB)	Specifies the maximum file size for the local log file. The default value is 200 MB. The maximum value is 500 MB.

Related Documentation

- Logging Overview

- Displaying System Logs

- Configuring Syslog

Enabling Client-Side Logging

Client-side logging is not enabled by default. If necessary, you can enable client-side logging to troubleshoot any client application issues.

To enable client-side logging:

1. Select **System > Log/Monitoring**.

2. Click the **Client Logs** tab to display the configuration page.

   Figure 208 shows the configuration page for Connect Secure.

3. Complete the configuration as described in Table 149.

4. Save the configuration.

Figure211: Client Logs Configuration Page
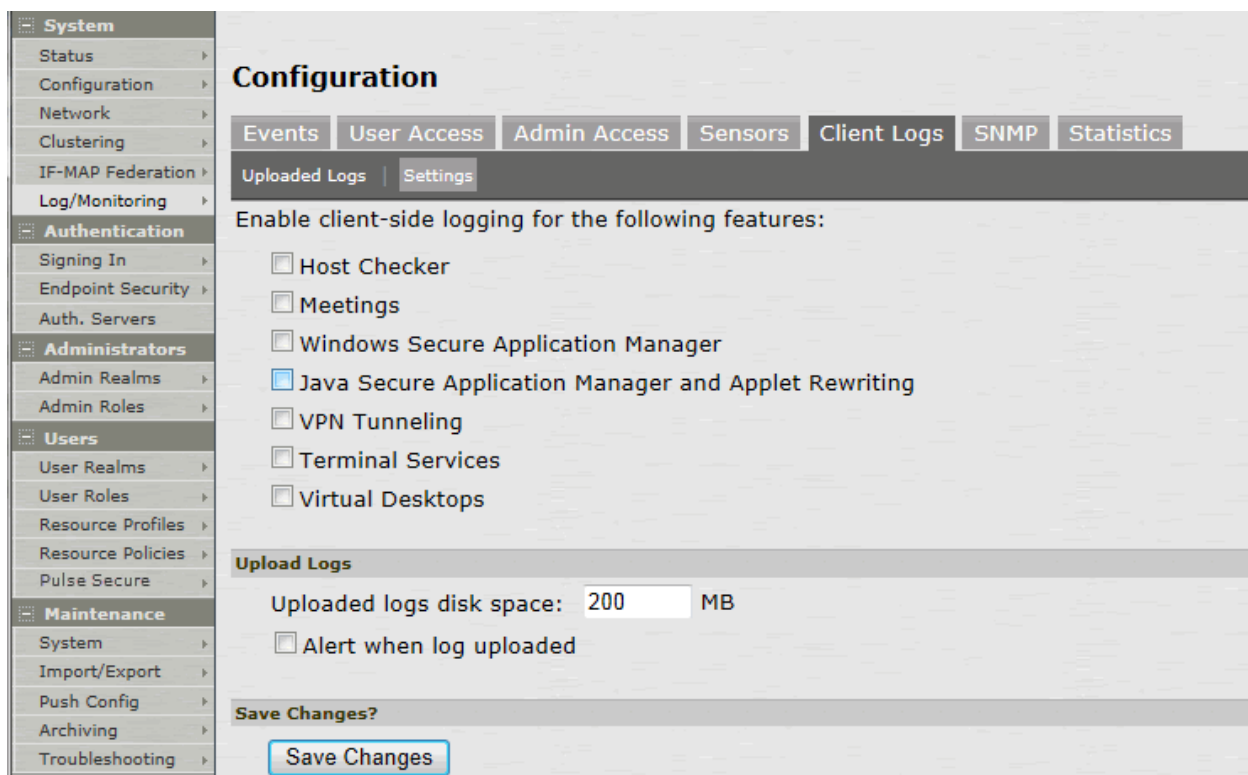
Table149: Client-Side Logs Settings

Settings	Guidelines
Host Checker	Select this option to enable client-side logging of Host Checker.
Meetings	Select this option to enable client-side logging of secure meeting.
Windows Secure Application Manager	Select this option to enable client-side logging of WSAM.
Java Secure Application Manager and Applet Rewriting	Select this option to enable client-side logging of JSAM and applet.
VPN Tunneling	Select this option to enable client-side logging of VPN tunneling.
Terminal Services	Select this option to enable client-side logging of terminal services.
Virtual Desktops	Select this option to enable client-side logging of virtual desktops.

Upload logs	
Upload logs disk space (MB)	Specify the amount of disk space (in Megabytes) you want to allocate for uploaded client log files. NOTE: You can allocate disk space from 0 to 200 MB.
Alert when log uploaded	Select this option to receive an alert message when an end user pushes a log file.

Related Documentation

- Logging Overview

Enabling and Viewing Client-Side Log Uploads

If you enable client-side logging for system features, you can also enable automatic upload of those logs at the role level. When you do, end users and Pulse Collaboration attendees who are members of the enabled roles can choose to push their log files up to the system at will. Then, you can view the uploaded files through the System > Log/Monitoring > Client Logs > Uploaded Logs page of the admin console.

When you upload log files to a device that is a node in a cluster, keep the following guidelines in mind:

- You can use the Log Node column on the System > Log/Monitoring > Client Logs > Uploaded Logs tab to view the location of existing log files collected by nodes in the cluster. This is specific to a cluster setup and does not apply to a single deployment.

- The user uploads logs to the cluster node to which he is connected.

- You can view upload log entries across all nodes in a cluster. You can save and unzip your uploaded log files from the respective nodes in the cluster where the user uploaded the logs.

- You can view upload log entries across all nodes in a cluster. You can save and unzip your uploaded log files from the respective nodes in the cluster where the user uploaded the logs.

- When a node is removed from a cluster, the system deletes the logs of that node from the Uploaded Log List in the cluster and from the node.

To enable end users to upload logs to the system:

1. Select **Users > User Roles > Select Role > General > Session Options**.

2. In the Upload logs section, select the **Enable Upload Logs** check box.

3. Click **Save Changes**.

Viewing Uploaded Client-Side Logs

If you enable end users to push log files up to the system, you can view the uploaded logs through the System > Log/Monitoring > Client Logs > Uploaded Logs page of the admin console. This page displays a list of uploaded log files from clients, featuring information such as the file name, date, associated user and/or realm, client access component type, and the log node.

NOTE: The system does not preserve uploaded logs when you upgrade the system software. To preserve the logs, you may archive them using options in the Maintenance > Archiving > Archiving Servers page of the admin console. You can also set the log-related SNMP traps to capture log events during the log upload using options in the System > Log/Monitoring > SNMP page of the admin console.

To view client log upload details:

1. In the admin console, choose **System > Log/Monitoring > Client Logs > Uploaded Logs** management page.

   Figure 209 shows the management page.

2. (Optional) Refresh uploaded client log details by clicking the **Refresh Logs** button.

3. (Optional) View or save an uploaded log by clicking on its respective link.

4. (Optional) Delete an uploaded log by clicking the trash can icon in the right side of the log's column. Note that once you delete a log from a node, those logs are lost.

Figure212: Uploaded Log Listing Page

Related Documentation

- Logging Overview

Configuring SNMP

If you prefer, you can use a third-party SNMP manager, such as HP OpenView, to monitor system health. The system supports SNMP v2c and SNMPv3.

To configure the SNMP agent:

1. Select **System > Log/Monitoring**.

2. Click the **SNMP** tab to display the SNMP configuration page.

   Figure 210 shows the configuration page for Policy Secure.

   Figure 211 shows the configuration page for Connect Secure.

3. Complete the configuration as described in Table 150.

4. Save the configuration.

Figure213: SNMP Configuration Page – Policy Secure

Figure214: SNMP Configuration Page – Connect Secure

Table150: SNMP Configuration Settings

Settings	Guidelines
MIB File	Use the **Pulse Secure MIB file** link to download the device management information base MIB file. You add this file to your SNMP manager configuration.
SNMP Version	Select your SNMP server version: • **v2c** • **v3**
Agent Properties	
SNMP Queries	Select to support SNMP queries.

SNMP Traps	Select to send SNMP traps.
System Name	Specify a system name.
System Location	Specify a location.
System Contact	Specify a system contact.
Community String	• Required only for SNMPv2c. • To query the system, your network management station must send it the community string. • To stop the SNMP system, clear the community field.

SNMPv3 Configuration

Username	Specify the SNMPv3 username. The User-Based Security Model (USM) is the default Security Module for SNMPv3. The system supports only one user at a time to be registered with an SNMP engine. Editing the SNMPv3 user attributes overwrite any already registered SNMPv3 user. The SNMPv3 user must have read-only access on all MIBs supported by the system. SNMPv3 user configuration attributes can also be used for SNMP traps.

Security Level	Selection	Auth Protocol	Auth Password	Priv Protocol	Priv Password
	No Auth, NoPriv	—	—	—	—
	Auth, NoPriv	Select MD5 (HMAC-MD5-96) or SHA (HMAC-SHA-96).	Enter an authentication password. The password can contain any ASCII characters and must be at least 8 characters in length.	—	—
	Auth, Priv	Select MD5 (HMAC-MD5-96) or SHA (HMAC-SHA-96).	Enter an authentication password. The password can contain any ASCII characters and must be at least 8 characters in length.	Select either CBC-DES or CFB-AES-128.	Enter a privacy password. The password can contain any ASCII characters and must be at least 8 characters in length.

Trap Thresholds	NOTE: Setting a threshold value to 0 disables that respective trap.
Check Frequency	Specify the frequency in seconds for sending traps. The default is 180 seconds.
Log Capacity	Specify the percent of log space used. The default is 90%.
Users	Specify the percent of user capacity used. The default is 100%.
Physical Memory	Specify the percent of physical memory used. The default is 0 (not reported).
Swap Memory (Virtual Memory)	Specify the percent of swap memory used. The default is 0 (not reported). NOTE: We recommend you monitor swap memory to alert you to potential memory issues. The threshold for traps for physical memory usage might be reached even if the system is not experiencing any difficulties.
Disk	Specify the percent of disk utilization. The default is 80%.
CPU	Specify the percent of CPU utilization. The default is 0 (not reported).
Meeting Users	Specify the percent of meeting users. The default is 100%.

Optional Traps	
Critical Log Events	Send traps when the system logs critical events.
Major Log Events	Send traps when the system logs major events.

Save SNMP Settings?	Click **Save Changes** to update the SNMP agent configuration. The page is refreshed and displays the SNMP engine ID. If the configuration is changed to move from SNMP v2c to SNMP v3, the system generates and displays two engine IDs.
SNMP Servers	
Hostname / IP address	Specify the hostname or IP address for the SNMP servers to which the system will send any traps it generates.
Port	Specify the port for the SNMP server. Typically, SNMP uses port 162.
Community	Specify the community string (if necessary).

Keep the following configuration tips in mind when you configure your SNMP manager to listen for this SNMP agent:

- Add the Pulse Secure MIB file to the SNMP manager configuration.

- If using SNMPv2c, the community string configuration for the SNMP manager and SNMP agent must match.

- If using SNMPv3, the SNMPv3 user configuration for the SNMP manager and the SNMP agent must match.

- If using SNMPv3, you must specify the Authoritative Engine ID for SNMPv3 traps that was generated when you saved the SNMP agent configuration.

Table 151 is a reference of MIB objects for the system. Some objects apply only to Connect Secure.

Table151: MIB Objects

Object	Description
logFullPercent	Returns the percentage of available file size filled by the current log as a parameter of the logNearlyFull trap.
signedInWebUsers	Returns the number of users signed in through a Web browser.
signedInMailUsers	Returns the number of users signed in to the e-mail client.
blockedIP	Returns the IP address—blocked due to consecutive failed login attempts—sent by the iveToomanyFailedLoginAttempts trap. The system adds the blocked IP address to the blockedIPList table.
authServerName	Returns the name of an external authentication server sent by the externalAuthServerUnreachable trap.
productName	Returns the licensed product name.
productVersion	Returns the software version.
fileName	Returns the file name sent by the archiveFileTransferFailed trap.
meetingUserCount	Returns the number of concurrent meeting users sent by the meetingUserLimit trap.
iveCpuUtil	Returns the percentage of CPU used during the interval between two SNMP polls. This value is calculated by dividing the amount of CPU used by the amount of CPU available during the current and previous SNMP polls. If no previous poll is available, the calculation is based on the interval between the current poll and system boot.
iveMemoryUtil	Returns the percentage of memory utilized by the system at the time of an SNMP poll. The system calculates this value by dividing the number of used memory pages by the number of available memory pages.
iveConcurrentUsers	Returns the total number of users logged in.
clusterConcurrentUsers	Returns the total number of users logged in for the cluster.

iveTotalHits	Returns the total number of hits to the system since last reboot. Includes total values from iveFileHits, iveAppletHits, meetingHits, and iveWebHits.
iveFileHits	Returns the total number of file hits to the system since last reboot.Incremented by the Web server with each GET/POST corresponding to a file browser request.
iveWebHits	Returns the total number of hits by means of the Web interface since last reboot. Incremented by the Web server for each http request received by the system, excluding file hits, applet hits, and meeting hits.
iveAppletHits	Returns the total number of applet hits to the system since last reboot.Incremented by the Web server for each GET request for a Java applet.
ivetermHits	Returns the total number of terminal hits to the system since last reboot.
logName	Returns the name of the log (admin/user/event) for the logNearlyFull and iveLogFull traps.
iveSwapUtil	Returns the percentage of swap memory pages used by the system at the time of an SNMP poll. The system calculates this value by dividing the number of swap memory pages used, by the number of available swap memory pages.
diskFullPercent	Returns the percentage of disk space used in the system for the iveDiskNearlyFull trap. The system calculates this value by dividing the number of used disk space blocks by the number of total disk space blocks.
blockedIPList	Returns a table with the 10 most recently blocked IP addresses. The blockedIP MIB adds blocked IP addresses to this table
ipEntry	An entry in the blockedListIP table containing a blocked IP address and its index (see IPEntry).
IPEntry	The index (ipIndex) and IP address (ipValue) for an entry in the blockedIPList table.
ipIndex	Returns the index for the blockedIPList table.
ipValue	A blocked IP address entry in the blockedIPList table.
logID	Returns the unique ID of the log message sent by the logMessageTrap trap.
logType	Returns a string sent by the logMessageTrap trap stating whether a log message is major or critical.
logDescription	Returns a string sent by the logMessageTrap trap stating whether a log message is major or critical.
ivsName	Returns the name of a virtual system.
ocspResponderURL	Returns the name of an OCSP responder.
fanDescription	Returns the status of the system fans.
psDescription	Returns the status of the system power supplies.
raidDescription	Returns the status of the system RAID device.
iveLogNearlyFull	The log file (system, user access, or administrator access) specified by the logName parameter is nearly full. When this trap is sent, the logFullPercent (%of log file full) parameter is also sent. You can configure this trap to be sent at any percentage. To disable this trap, set the Log Capacity trap threshold to 0%. The trap's default value is 90%. NOTE: When SNMP traps are enabled, the iveLogNearlyFull and iveLogFull traps are sent when the log files are 90% full and 100% full respectively, even if the threshold is set to 0 (disabled).
iveLogFull	The log file (system, user access, or administrator access) specified by the logName parameter is completely full. NOTE: When SNMP traps are enabled, the iveLogNearlyFull and iveLogFull traps are sent when the log files are 90% full and 100% full respectively, even if the threshold is set to 0 (disabled).

iveMaxConcurrentUsersSignedIn	Maximum number or allowed concurrent users are currently signed in. You can configure this trap to be sent at any percentage. To disable this trap, set the Users trap threshold to 0%. The trap's default value is 100%.
iveTooManyFailedLoginAttempts	A user with a specific IP address has too many failed sign-in attempts. Triggered when a user fails to authenticate according to the settings for the Lockout options on the Security Options tab. When the system triggers this trap, the system also triggers the blockedIP (source IP of login attempts) parameter.
externalAuthServerUnreachable	An external authentication server is not responding to authentication requests. When the system sends this trap, it also sends the authServerName (name of unreachable server) parameter.
iveStart	The system has just been turned on.
iveShutdown	The system has just been shut down.
iveReboot	The system has just been rebooted.
archiveServerUnreachable	The system is unable to reach the configured archive server.
archiveServerLoginFailed	The system is unable to log into the configured archive server.
archiveFileTransferFailed	The system is unable to successfully transfer files to the configured archive server. When the system sends this trap, it also sends the fileName parameter.
iveRestart	Supplies notification that the system has restarted according to the administrator's instruction.
iveDiskNearlyFull	Supplies notification that the system disk drive is nearly full. When the system sends this trap, it also sends the diskFullPercent parameter. You can configure this trap to be sent at any percentage. To disable this trap, set the Disk trap threshold to 0%. This trap's default value is 80%.
iveDiskFull	Supplies notification that the system disk drive is full.
logMessageTrap	The trap generated from a log message. When the system sends this trap, it also sends the logID, logType, and logDescription parameters.
memUtilNotify	Supplies notification that the system has met the configured threshold for memory utilization. To disable this trap, set the Physical Memory trap threshold to 0. The threshold is 0%, by default.
cpuUtilNotify	Supplies notification that the system has met the configured threshold for CPU utilization. To disable this trap, set the CPU trap threshold to 0. The threshold is 0%, by default.
swapUtilNotify	Supplies notification that the system has met the configured threshold for swap file memory utilization. To disable this trap, set the Swap Memory trap threshold to 0. The threshold is 0%, by default.
iveFanNotify	Supplies notification that the status of the fans has changed.
ivePowerSupplyNotify	Supplies notification that the status of the power supplies has changed.
iveRaidNotify	Supplies notification that the status of the RAID device has changed.
iveNetExternalInterfaceDownTrap (nicEvent)	Supplies the type of event that brought down the external interface. The nicEvent parameter can contain values of "external" for an external event and "admin" for an administrative action.
iveNetInternalInterfaceDownTrap (nicEvent)	Supplies the type of event that brought down the internal interface. The nicEvent parameter can contain values of "external" for an external event and "admin" for an administrative action.
iveClusterDisableNodeTrap (clusterName,nodeList)	Supplies the name of the cluster that contains disabled nodes, as well as a string containing the names of all disabled nodes. Node names are separated by white space in the string.

iveClusterChangedVIPTrap(vipType, currentVIP, newVIP)	Supplies the status of a virtual IP for the cluster. The vipType indicates whether the changed VIP was external or internal. The currentVIP contains the VIP prior to the change, and newVIP contains the VIP after the change.
iveNetManagementInterfaceDownTrap (nicEvent)	Supplies the type of event that brought down the management port. The nicEvent parameter can contain values of "external" for an external event and "admin" for an administrative action.
iveClusterDelete(nodeName)	Supplies the name of the node on which the cluster delete event was initiated.

Related Documentation

- Logging Overview

Configuring Syslog

If desired, you can configure the system to send logs to a syslog server.

To configure reporting to a syslog server:

1. Select **System > Log/Monitoring**.

2. Click the **Settings** tab to display the configuration page.

   Figure 212 shows the configuration page for Policy Secure.

   Figure 213 shows the configuration page for Connect Secure.

3. Specify the maximum log size and select the events to be logged.

4. Specify the server configuration as described in Table 152, and click **Add**. You can specify multiple syslog servers.

5. Save the configuration.

   NOTE: To enable syslog reporting for each local log category, you must perform this procedure on each local log tab: Events, User Access, Admin Access, and Sensors.

Figure215: Syslog Server Configuration Page – Policy Secure

Figure216: Syslog Server Configuration Page – Connect Secure
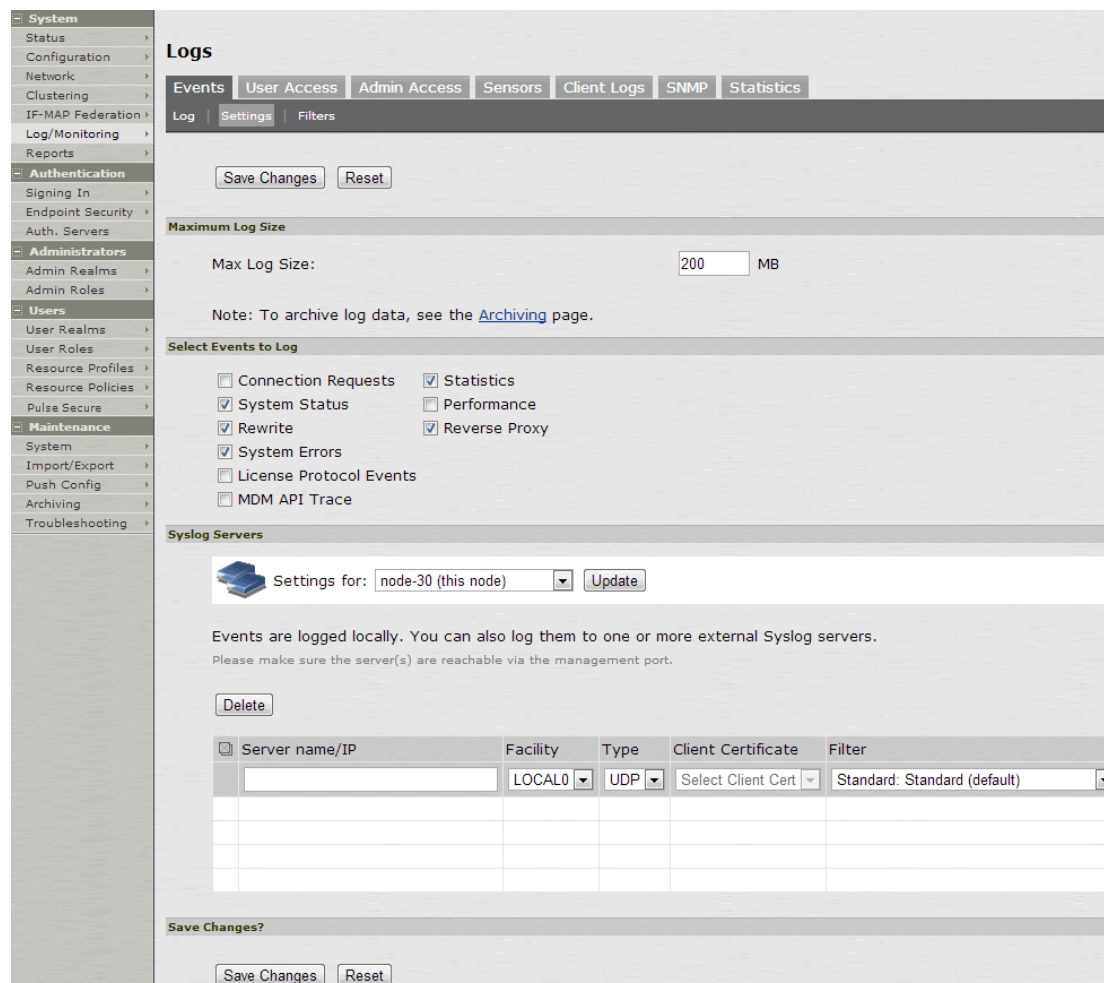
Table152: Syslog Server Configuration Guidelines

Settings	Guidelines
Server name/IP	Specify the fully qualified domain name or IP address for the syslog server. NOTE: If you select **TLS** from the Type list, the server name must match the CN in the subjectDN in the certificate obtained from the server.
Facility	Select a syslog server facility level (LOCAL0-LOCAL7). Your syslog server must accept messages with the following settings: facility = LOG_USER and level = LOG_INFO.
Type	Select the connection type to the syslog server. You can select: • UDP (User Datagram Protocol) - A simple non-secure transport model. • TCP (Transmission Control Protocol) - A core protocol of the Internet Protocol suite (IP), but lacks strong security. • TLS (Transport Layer Security) - Uses cryptographic protocols to provide a secure communication.
Client Certificate	(optional) If you select TLS from the Type menu and your remote syslog server requires client certificates, select the installed client certificate to use to authenticate to the syslog server. Client certificates are defined in the Configuration > Certificates > Client Auth Certificates page. Client certificates must be installed on the device before they can be used. NOTE: There is no fallback if a connection type fails.

Filter	Select a filter format. Any custom filter format and the following predefined filter formats are available:

- **Standard** (default)—This log filter format logs the date, time, node, source IP address, user, realm, event ID, and message.
- **WELF**—This customized WebTrends Enhanced Log Format (WELF) filter combines the standard WELF format with information about the system realms, roles, and messages.
- **WELF-SRC-2.0-Access Report**—This filter adds access queries to the customized WELF filter. You can use this filter with NetIQ's SRC to generate reports on user access methods.

Related Documentation

- Logging Overview

- Using Log Filters

Displaying System Status

The System Status page is a dashboard of system version information, system capacity utilization, uptime, and summary user information. The System Status page is the "home" page that is displayed when you log into the admin console as an administrator. To navigate to the System Status page from other admin console pages, select **System > Status**.

Figure 214 shows the System Status page for Policy Secure. The table that follows describes the numbered figure callouts.

Figure 215 shows the configuration page for Connect Secure.

Figure217: System Status Page – Policy Secure
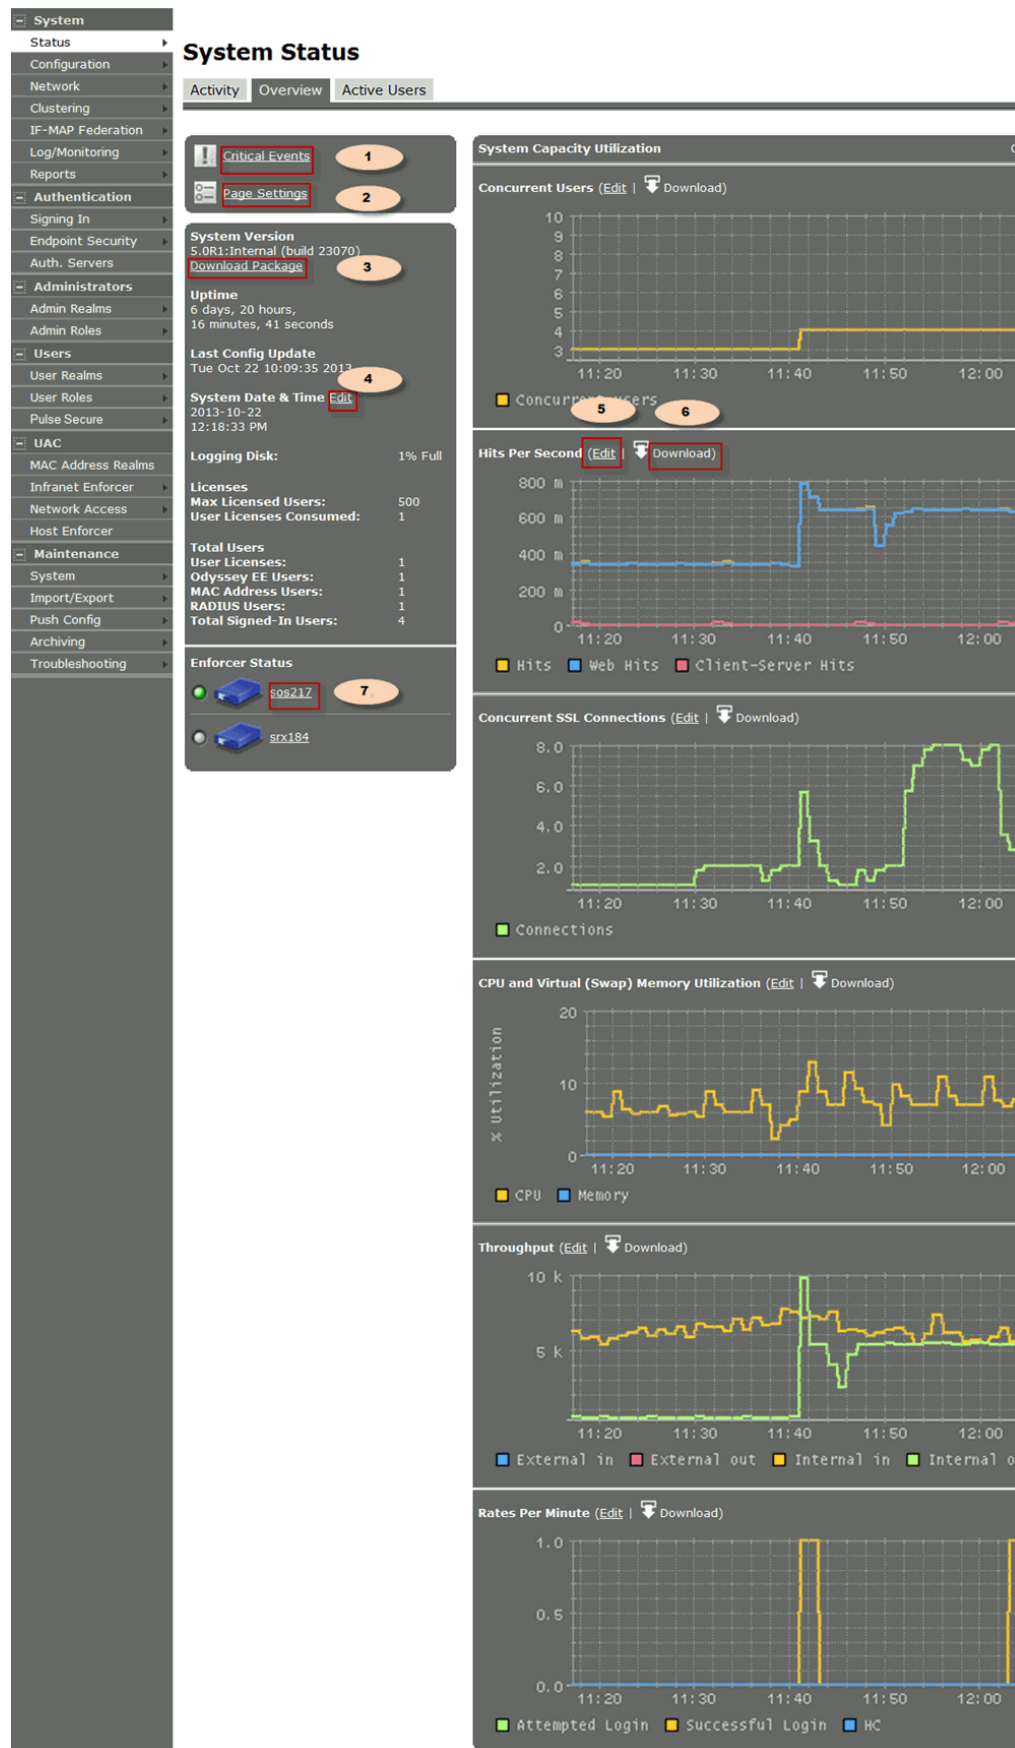
Figure218: System Status Page – Connect Secure
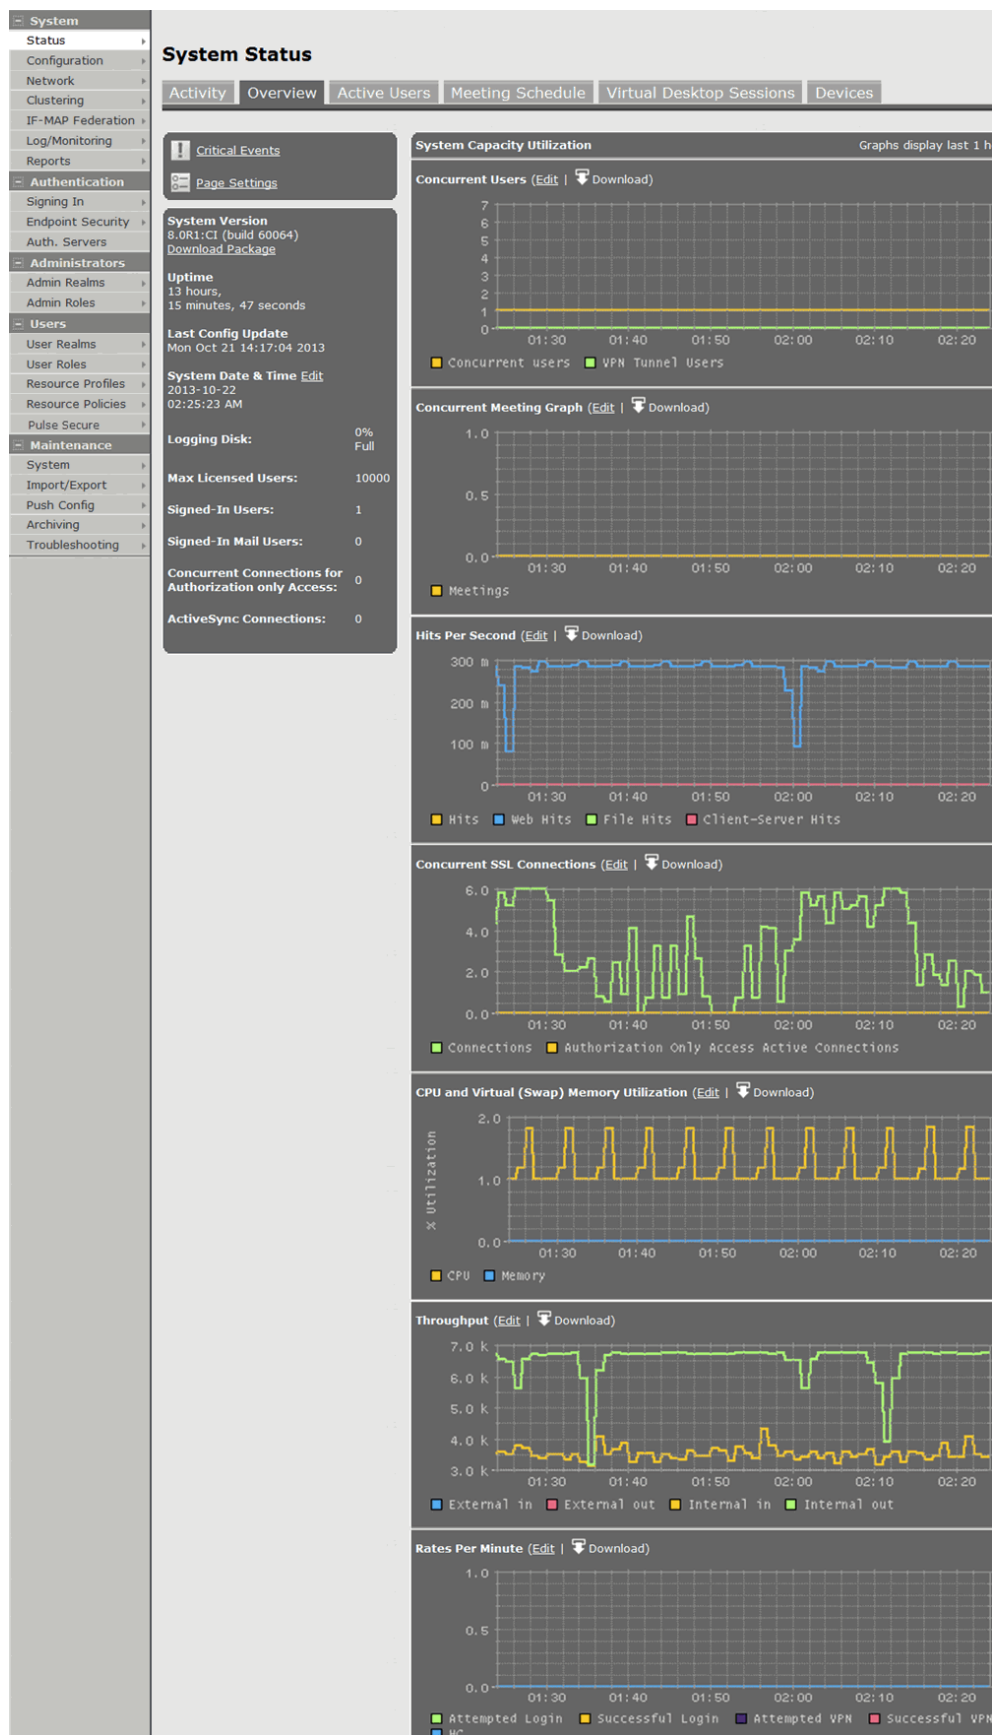
Callout	Description
1	Click the **Critical Events** link to display a new window with a table of the last 10 critical system events.
2	Click the **Page Settings** link to display a new window with the System Status Settings page shown in Figure 216.
3	Click the System Version **Download Package** link to download the software version running on the system. You might do this when you need to synchronize software on another node to the software version running on this system.
4	Click the System Date and Time **Edit** link to display the System Date and Time configuration page. See *Configuring the System Date and Time.*
5	Click a System Capacity Utilization report **Edit** link to display a new window with controls to customize the appearance of the report graphs.
6	Click a System Capacity Utilization report **Download** link to download graph data in XML format.
7	Click an Enforcer Status link to navigate to its configuration page.

Table153: Licenses and Total Users – Policy Secure

Item	Description
Licenses This applies only to Policy Secure.	
Max Licensed Users	Displays the maximum number of licensed users by supported platform type.
User Licenses Consumed	Displays how user licenses are being used. Pulse Secure client connections and agentless connections count as user licenses.
Total Users This applies only to Policy Secure.	
User Licenses	Displays the number of Pulse Secure client connections and agentless connections. These types of connections count as user licenses. The maximum number of licenses is the sum of the capacity provisioned by the licenses that have been added to the system.
Odyssey EE Users	Displays the number of OAC enterprise edition (EE) connections. These connections do not count as user licenses, and there is no maximum number of licenses set by rule.
MAC Address Users	Displays the number of connections through MAC Address authentication realms. This number is reported only if a MAC Address realm has been configured. These connections do not count as user licenses, and there is no maximum number of licenses set by rule.
RADIUS Users	Displays the number of RADIUS user connections. These connections do not count as user licenses, and there is no maximum number of licenses set by rule.
Total Signed-In Users	Displays the sum of user licenses in use. • User Licenses count • Odyssey Users count • MAC Address Users count • RADIUS Users count

Table154: Licenses and Total Users – Connect Secure

Item	Description

Max Licensed Users	Displays the maximum number of licensed users by supported platform type.
Signed-In Users	Displays the number of signed-in users.
Signed-In Mail Users	Displays the number of signed-in mail users.
Concurrent Connections for Authorization only Access	Displays the concurrent connections for authorization only access.
ActiveSync Connections	Displays the number of ActiveSync connections.

Figure 216 shows the System Status Settings configuration page. The settings configuration page for Connect Secure is similar.

Figure219: System Status Settings Configuration Page

You can use this page to select the reports displayed on the System Status page, as well as data properties, such as the time dimension and refresh rate.

The following reports are available:

- Concurrent Users—Shows a count of users signed into the system. In clustered environments, the graph includes two lines. The first line displays the number of local users signed into the node selected from the list, and the second line displays the number of concurrent users signed into the entire cluster.

- Hits per Second—Shows a count of hits currently being processed by the system. In a clustered environment, you may select a node from the list to determine which node's data is displayed in the graph. The graph includes three lines: total number of hits, number of Web hits, and number of client/server hits.

- CPU and Memory Usage—Shows the percentage of the CPU and memory being used. In a clustered environment, you may select a node from the list to determine which node's data is displayed in the graph.

- Throughput—Shows the amount of data (in KB) being processed. In a clustered environment, you may select a node from the list to determine which node's data is displayed in the graph. The graph includes four lines: external in, external out, internal in, and internal out.

- Meetings—Shows the count of concurrent meetings. This option is available only on Connect Secure.

- Connections—Shows a count of concurrent SSL connections.

- Rates—Shows the rate of attempted logins, successful logins, and Host Checker updates.

Related Documentation

- Logging Overview

Viewing and Cancelling Scheduled Meetings

You can view all of the meetings currently scheduled on Connect Secure or cancel meetings.

To view and cancel scheduled meetings:

1. Select **System > Status > Meeting Schedule**. The system displays real-time information about all of the meetings that are currently running or scheduled, including:

   - **Time and Status**—Displays the time and duration that the meeting is scheduled to run, as well as the current status of the meeting.

   - **Meeting Details**—Displays the meeting name, ID, and password requirements. This column also includes a Details link that you can use to view information about the meeting and to join the meeting.

   - **Meeting Role**—Displays the role of the meeting creator. If the creator was signed into multiple roles when he created the meeting (that is, he is a member of multiple roles and the appliance is configured for a permissive merge).

   - **Attendee Roles**—Displays the roles of the attendees who are signed into the meeting, the number of attendees signed into each role, and each role's meeting attendee limit. Note that non-Connect Secure attendees are displayed under the meeting creator's user role.

2. Use either of the following methods to change the meeting view (optional):

   - Select a time frame (Daily, Weekly, In Progress, Scheduled) from the drop-down list to control which meetings are displayed.

   - Click on any of the underlined column headers to control the order in which currently displayed meetings are sorted.

3. Click the **Details** link under a meeting to view information about the meeting and optionally to join the meeting (optional).

4. Choose **MyMeeting URLs** from the View drop menu to view personal URLs your users have created.

5. Click the delete icon in the right column to cancel a meeting or to delete a MyMeeting URL (optional).

Cancelling a meeting permanently deletes from the system. You cannot restore a meeting after cancelling it.

Related Documentation

- Pulse Collaboration Overview

Displaying Hardware Status

You can use the Maintenance > System > Platform page to display the hardware health status, including information about hard drives, fans, and power supplies.

To display hardware health status:

1. Select **Maintenance > System > Platform** to display the System Maintenance page.

   Figure 217 shows the system maintenance page for Policy Secure.

   Figure 218 shows the system maintenance page for Connect Secure.

2. Review the hardware status information described in Table 155.

Figure220: System Maintenance Page – Policy Secure

Figure221: System Maintenance Page – Connect Secure
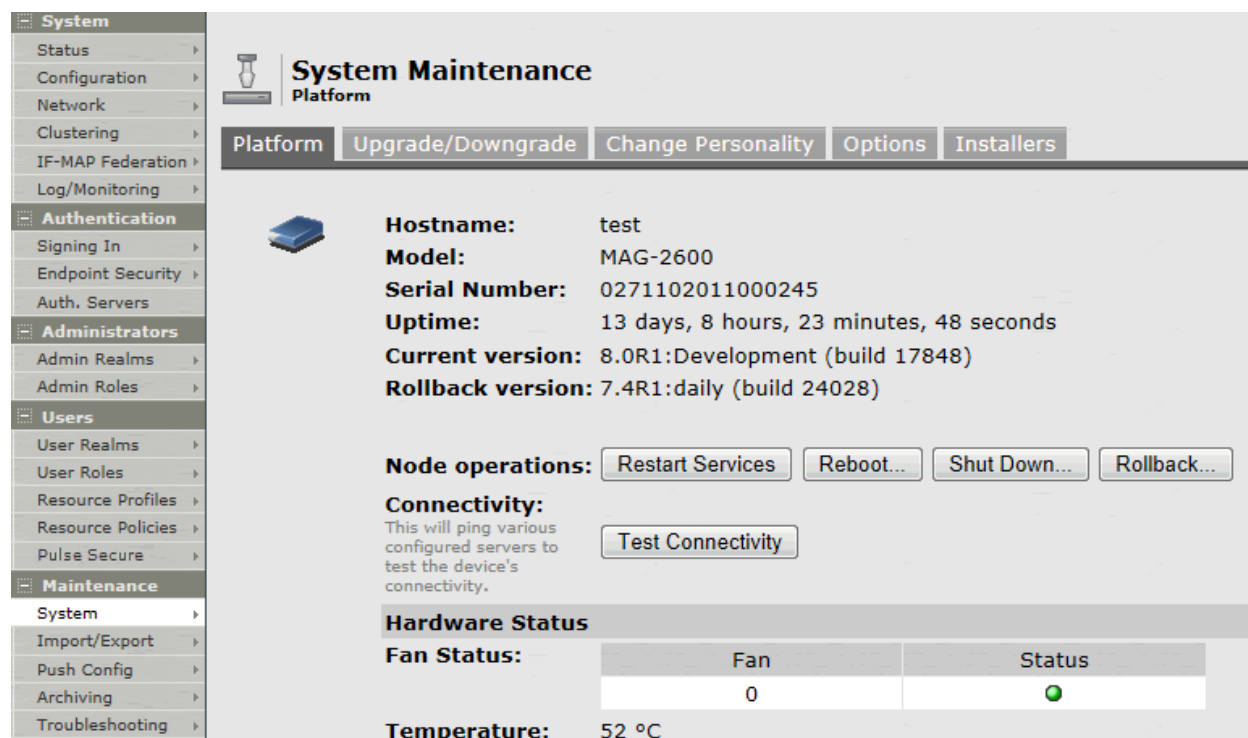
Table155: Hardware Status Information

Hardware Component	Status Message
Hard Disk Status	Displays a health statement for the device disk drive. See Table 156 and Table 157 for details.
Fan Status	Displays a health statement for the device fan(s).
Power Supply	Displays a health statement for the device power supply.

Table 156 lists the RAID status and hard drive status. Depending on your system, you may or may not see all these possible statuses.

Table156: RAID and Hard Drive Status

RAID Status	Drive 1	Drive 2
Hard Disk RAID is operational	Active	Active
Hard Disk RAID is in single drive mode	Missing	Active
Hard Disk RAID is in single drive mode	Active	Missing
Hard Disk RAID has failed	Failed	Active
Hard Disk RAID has failed	Active	Failed
Hard Disk RAID is in the process of recovering	Active	Reconstructing
Hard Disk RAID is in the process of recovering	Reconstructing	Active
Hard Disk RAID is in the process of recovering	Active	Verifying
Hard Disk RAID is in the process of recovering	Verifying	Active
Hard Disk RAID status is unknown	Unknown	Active

Hard Disk RAID status is unknown	Active	Unknown
Hard Disk RAID status is unknown	Unknown	Unknown
Not available	n/a	n/a

Table 157 lists all the possible RAID status and hard drive status for the MAG-SM360. You can also view the RAID and hard drive status in log messages and in SNMP.

Table157: RAID and Hard Drive Status for the MAG-SM360

RAID Status	Drive 1	Drive 2
Optimal	Optimal	Optimal
Degraded	Rebuilding	Optimal
Degraded	Optimal	Rebuilding
Degraded	Offline	Optimal
Degraded	Optimal	Offline

Related Documentation

- Displaying System Status

Using Software RAID on the PSA7000

This section describes the use of software RAID on the PSA7000. It includes the following information:

- Overview of Software RAID on the PSA7000
- Configuring RAID Controller on the PSA7000
- Checking RAID Statuses

Overview of Software RAID on the PSA7000

All large devices of PCS have two hard disks but, have no hardware RAID controller. In PCS 8.2, RAID functionality is enabled through the software RAID available in Linux. Software is used to create RAID devices in the RAID 1 layer. These individual disk partitions can also be used as regular partitions. All partitions, which include boot, root, data, var, tmp, swap, and so on are created out of the software RAID. Software RAID does not affect any of the other hardware models except PSA7000. It works the same way either with hardware RAID or no RAID at all.

Configuring RAID Controller on the PSA7000

To configure a RAID controller on the PSA7000, disks are hot plugged on the PSA7000. Configuring RAID on the PSA7000 also involves some manual configuration steps in the admin console. RAID configuration is carried out by the following steps:

- To remove a disk from the second slot in the system:
    1. Select Menu option: **4. System Operations**
    2. Select Menu option: **20. Manage RAID**
    3. A prompt appears: **Are you sure you want to manage RAID? (y/n)**
    4. Enter **y** for the prompt.
    5. Choose Menu option: **2. Remove Disk2**. This will detach the disk in the second slot from software RAID.
    6. Physically unplug the disk from the second slot.

- To insert a disk to the second slot in the system:
    1. Physically plug in the disk in the second slot.
    2. Select Menu option: **4. System Operations**
    3. Select Menu option: **20. Manage RAID**
    4. A prompt appears: **Are you sure you want to manage RAID? (y/n)**
    5. Enter **y** for the prompt.
    6. Choose Menu option: **4. Add Disk2.** This will attach the disk in the second slot to the software RAID.

Checking RAID Statuses

To check the status of RAID and individual disks:

1. Go to **system > platform** page of the web interface.
2. Status of RAID and individual disks are displayed. Table 158 shows the hard disk status and the hard disk RAID status.

Table 158: Hard Disk and Hard Disk RAID Statuses

	Status	Guidelines
Hard Disk	Active	A disk that is present and part of RAID
	Inactive	A disk that is present but not part of RAID
	Missing	A disk that is removed.
Hard Disk RAID	Operation	Both disks are active and part of the software RAID
	Recovering	Both disks are active and syncing with each other
	Failed	Both disks are active but one of them is not part of RAID

ℹ️ NOTE: In the single drive mode, only one disk is active. The other disk is either missing or inactive.

LCD Display

This section describes the addition of LCD to PCS devices. It includes the following information:

- Overview of adding LCD for PCS
- Modes Supported by the LCD

Overview of adding LCD for PCS

The addition of a LCD screen allows field technicians to quickly gauge the health of the system without logging into the device. The buttons on the LCD panel allow navigation through the display menus. The directional buttons are used to access the menu modes and find device information. The LCD can display two line of text. Figure 222 shows the LCD screen with navigation buttons.

Figure 222: LCD with Navigation Buttons
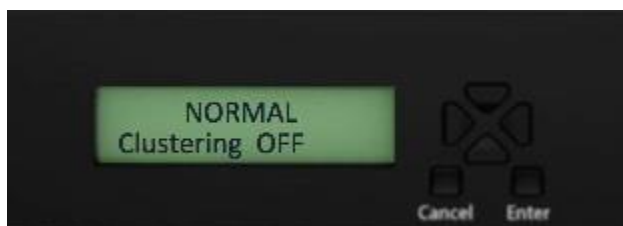
Modes Supported by the LCD

The LCD supports two modes namely the display mode (default) and the menu mode. Pressing any button in the display mode will change the mode to menu mode. If a user presses the cancel button, the LCD immediately changes back to display mode and shows the appropriate state. The LCD remains in display mode. If the LCD is in menu mode and the user does not press any button for more than two minutes, then the LCD changes back to display mode. Figure 223 shows the two modes supported by the LCD

Figure 223: Two Modes Supported by the LCD
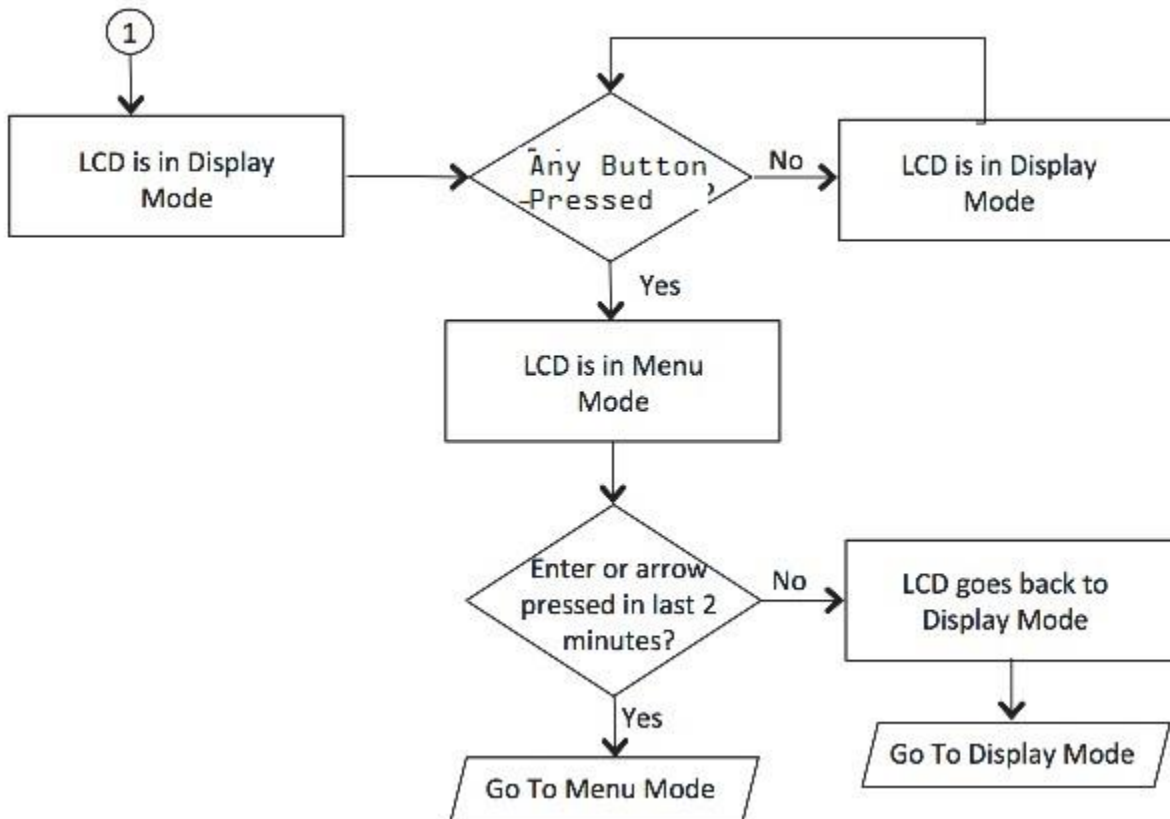
Display Mode

The display mode describes the current state of the system, such as normal state or error conditions (e.g., fan speed and overheat). It represents the default status. The LCD goes into display mode after boot-up is complete. In display mode, the LCD is either set to NORMAL or shows a label that describes an error condition. If all systems are functioning normally, then the LCD shows NORMAL. The second line in the NORMAL state is used to show whether the appliance is configured as part of a cluster. The valid states in the display mode are Clustering OFF and Clustering ON. Figure 224 shows the two valid states in display mode.

Figure 224: Valid States in Display Mode
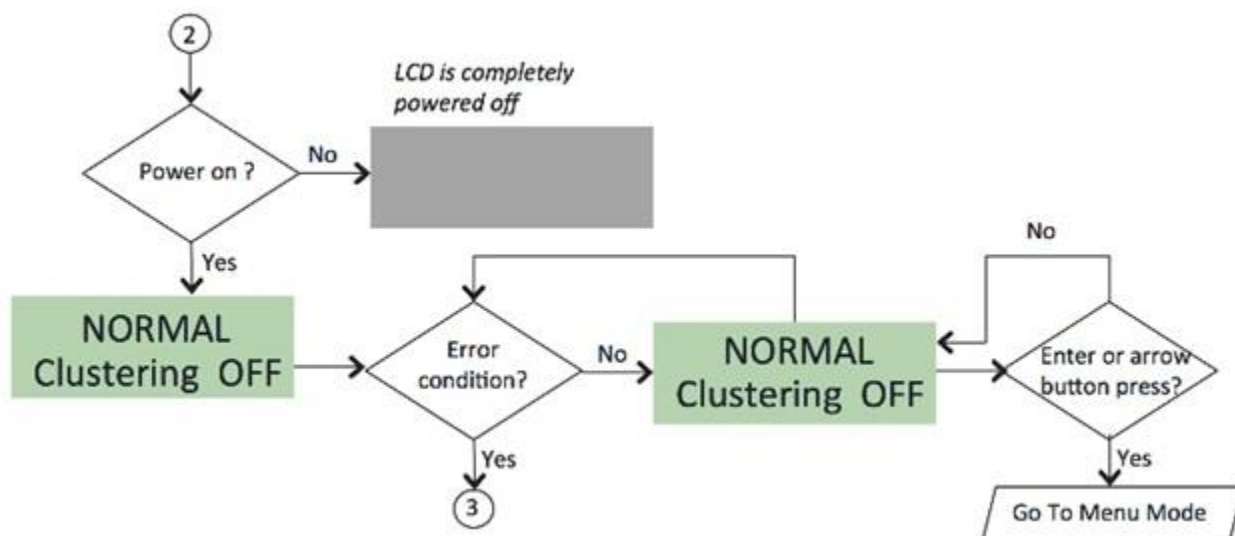
Detecting Error Conditions in Display Mode

If more than one error condition is detected, all error conditions will be displayed in sequence with a 2 second pause before switching to the next one. All error conditions need to be cleared before the status returns back to the NORMAL state. Error conditions include:

- Overheat
- Fan Failure
- RAID Errors

Figure 225 shows the various error conditions that are detected.

Figure 225: Error Conditions
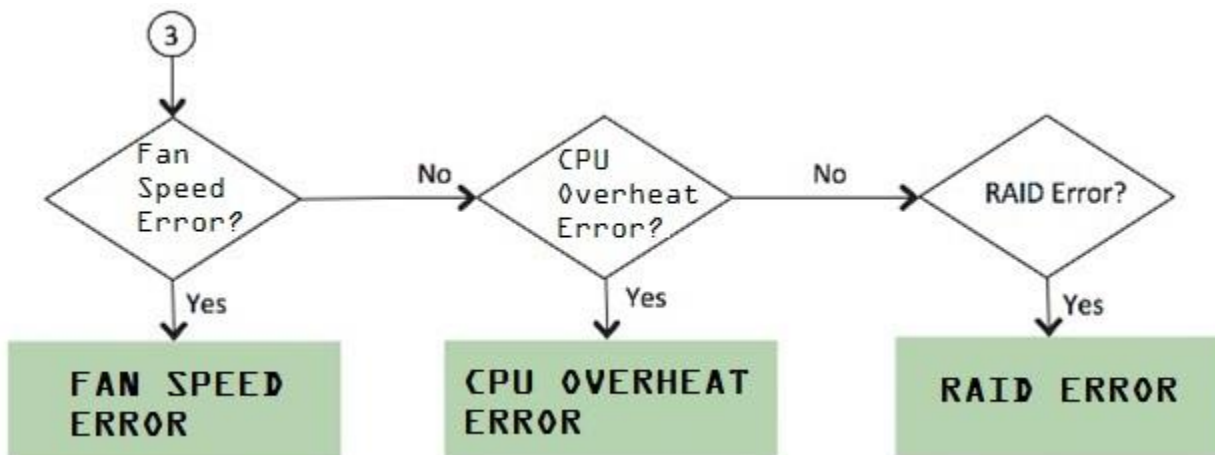
To detect the error conditions in the display mode:

If there are any error conditions, they are automatically shown on the LCD screen when it is in the display mode. The types of errors displayed are: Fan Failures, CPU overheating and RAID errors. If there are multiple errors they would be displayed in the order shown in Figure 225 with a two second pause between successive displays.

The error message is automatically cleared when the underlying error condition is resolved. For example: the CPU overheat message disappears when CPU temperature is lowered. The user can enter the Menu mode at any point, even if an error message is being displayed.

Menu Mode

The menu mode is activated when the user presses any button. A single press of the button changes to menu mode and loads the last selected menu selection.

To view information in the menu mode:

1. Press any button. This puts the LCD into menu mode.
2. Press the right and left arrows keys to obtain the available system configuration data.
3. View information starting with the Internal IP and moving in a clockwise direction.
4. The menu screens loop back in a cycle.
5. Press Cancel at any point to exit to display mode.

NOTE: Any button, even cancel will put the user in the menu mode.

Figure 226 show the available system configuration data available in menu mode.

Figure 226: System Configuration Data Available in Menu Mode
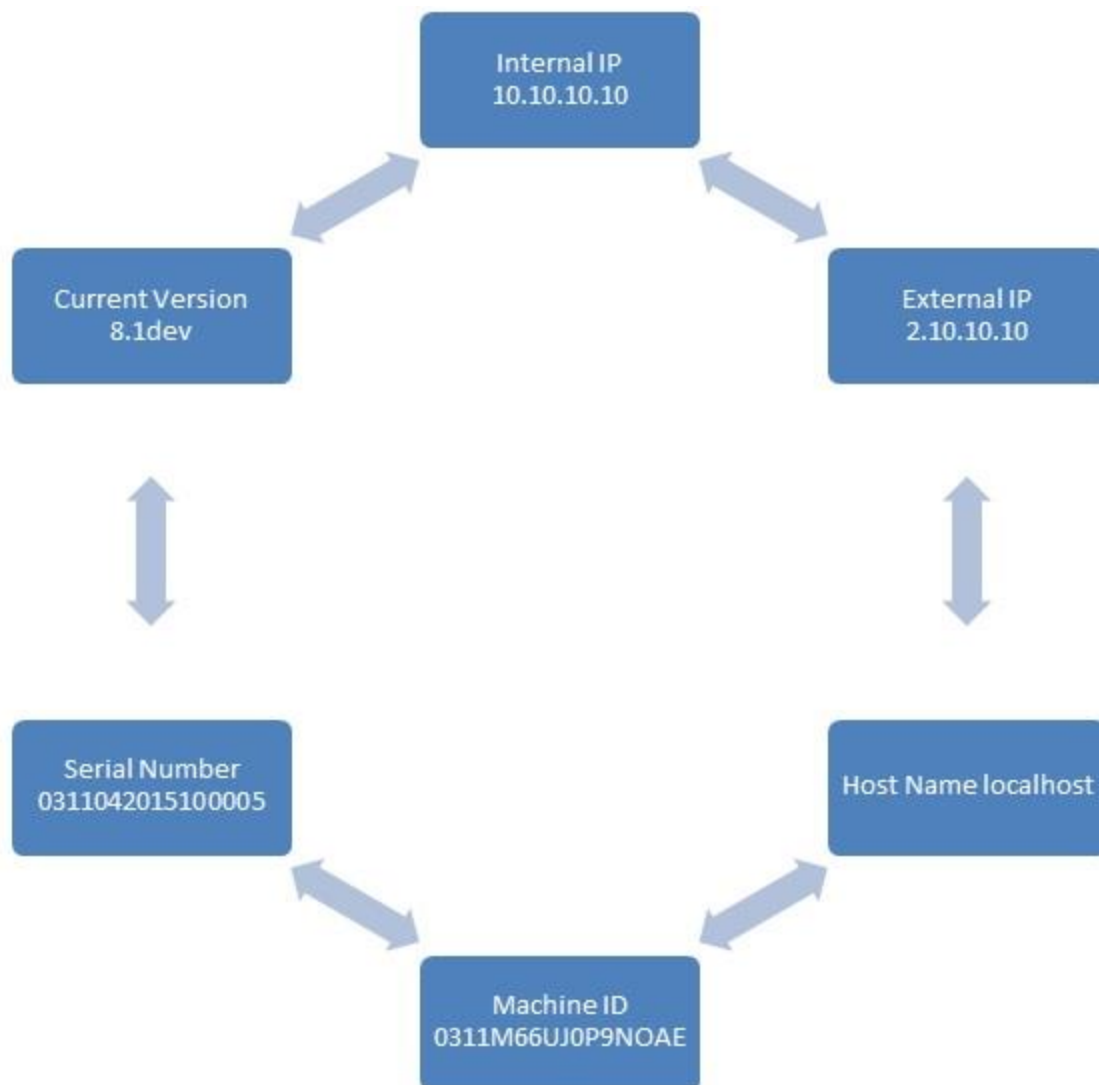
Displaying Active Users

You can use the Active Users page to display the system active users table and to perform administrative actions pertaining to active sessions.

The system active users table displays all users who have an active session (in contrast to the users tables that appear on the authentication server configuration pages, which display session records for active and inactive sessions that were authenticated by the particular authentication server).

If a user signs in and is placed in a VLAN without an IP address, the table does not display an IP address under Signed in IP.

If there is a NAT device between the user's computer and the Infranet Enforcer, the table displays both the NAT device's IP address and the endpoint's virtual source IP address under Signed in IP. For example, if the NAT device's IP address is 10.64.9.26, and the endpoint's virtual source IP address is 192.168.80.128, the following information is displayed under Signed in IP: **10.64.9.26 (192.168.80.128 behind NAT)**.

To display the system Active Users page:

1. Select **System > Status**.
2. Click the **Active Users** tab to display the system active users page.
3. Use the controls described in Table 158 to perform administrative actions pertaining to active sessions.

Table159: Active Users Page

Buttons	Administrative Actions
Update	Refresh records displayed on the page: • To refresh the page, click **Update**. • To display a specific user, enter the username in the Show Users Named box and click **Update**. If you do not know the exact username, use the asterisks (*) as a wildcard character. • To change the table size, enter a number in the Show N users box and click **Update**. TIP: To sort the table of currently signed-in users and administrators, click a column header.
Delete Session	Select the check box next to the appropriate names and then click **Delete Session** to immediately delete the session. The user is signed out by your action.
Delete All Sessions	Use this option to immediately delete all sessions. Users are signed out by your action. NOTE: If you want to sign out administrators, you must choose them individually and use the Delete Session button.
Refresh Roles	Manually evaluate all authentication policies, role-mapping rules, role restrictions, user roles, and resource policies for all currently signed-in users. Use this button if you make changes to an authentication policy, role-mapping rules, role restrictions, or resource policies and you want to immediately refresh the roles of all users.
Disable All Users	Sign out all end users who are currently signed-in and also prevent any other users from signing in. To allow users to sign in again after you disable all users, click **Enable All Users**.

Related Documentation

- Displaying System Status

Displaying System Logs

This topic describes how to display local system logs. It includes the following information:

- Displaying Events Logs
- Displaying User Access Logs
- Displaying Admin Access Logs
- Displaying Sensor Logs

Displaying Events Logs

The Events logs include system events, such as session timeouts, system errors and warnings, requests to check server connectivity, and system restart notifications. The local log viewer displays the most recent 5000 log messages (the display limit).

To display Events logs:

1. Select **System Log/Monitoring**.

2. Click the **Events** tab.

3. Click the **Log** tab to display the log page.

   Figure 219 shows the log page for Policy Secure.

   Figure 220 shows the log page for Connect Secure.

4. Use the features described in Table 159 to examine log records or manage the log collection.

Figure227: Events Logs Page – Policy Secure

Figure228: Events Logs Page – Connect Secure

Table160: Log Management Features

Controls	Description
Filter	Select a filter format. Any custom filter formats and the following predefined filter formats are available: • **Standard** (default)—This log filter format logs the date, time, node, source IP address, user, realm, event ID, and message. • **WELF**—This customized WebTrends Enhanced Log Format (WELF) filter combines the standard WELF format with information about the system realms, roles, and messages. • **WELF-SRC-2.0-Access Report**—This filter adds access queries to the customized WELF filter. You can use this filter with NetIQ's SRC to generate reports on user access methods. NOTE: Format filters change only the data displayed (or columns exported), and do not affect the log data that has been collected.
Query	In the log display, several fields are hyperlinks. The hyperlinks function as dynamic queries on the local log collection. For example, if you click the log ID, the date, or an IP address or username, the log viewer queries the log collection for records that match the value you clicked, and redisplays the log collection. You can apply additional query filters by clicking additional hyperlinked values, essentially creating a Boolean AND query (for example, date AND IP address). Use the **Reset Query** button to clear the query filters and redisplay the unfiltered log collection. Use the **Save Query** button to save the dynamic log query as a custom filter. When you click the Save Query button, the system displays the Filters tab displays with the Query field prepopulated with the variables you selected from the log. NOTE: Query filters change only the display (or rows exported), and do not affect the log data that has been collected.
Save Log As	Save the local log collection to a file. We recommend you retain the system generated log name, which follows a consistent convention: **juniper.logtype.nodename.log**. The local log viewer displays the most recent 5000 log messages (the display limit). If the current log file contains fewer than 5000 log messages, older log messages from the backup log file are displayed, up to a total of 5000 log messages. This makes the log files appear as one, even though they are stored separately. When you save the log messages or use the FTP archive function, the backup log file is appended to the current log file, and is then downloaded as one log file. If the log files are not archived or saved by the time they are rolled over again, the oldest log messages (saved in the backup log file) are lost.
Clear Log	Clear the local **log** and **log.old** file. When you clear the local log, events recorded by the syslog server are not affected. Subsequent events are recorded in a new local log file.
Save All Logs	The Save All Logs button appears on the Events, User Access, Admin Access, and Sensors tabs. When you click **Save All Logs**, the system generates a file that includes event, user access, admin access, sensor logs, and XML data for all of the system statistics and graphs shown on the Status > Overview page. After you click **Save All Logs**, you are prompted to download a file named **juniperlogs-graphs.tar.gz** to your local host.
Clear All Logs	The Clear All Logs button appears on the Events, User Access, Admin Access, and Sensors tabs. It clears event, user access, admin access, sensor logs, and XML data for all of the system statistics and graphs shown on the Status > Overview page. When you clear the local log, events recorded by the syslog server are not affected. Subsequent events are recorded in a new local log file.

Displaying User Access Logs

The User Access logs include information about user access, such as the number of simultaneous users at each one hour interval (logged on the hour) and user sign-ins and sign-outs. The local log viewer displays the most recent 5000 log messages (the display limit).

To display User Access logs:

1. Select **System Log/Monitoring**.

2. Click the **User Access** tab.

3. Click the **Log** tab.

4. Use the features described in Table 159 to examine log records or manage the log collection.

Displaying Admin Access Logs

The Admin Access logs include information about administrator actions, such as administrator changes to user, system, and network settings. It includes a log entry whenever an administrator signs in, signs out, or changes licenses on the appliance. The local log viewer displays the most recent 5000 log messages (the display limit).

To display Admin Access logs:

1. Select **System Log/Monitoring**.

2. Click the **Admin Access** tab.

3. Click the **Log** tab.

4. Use the features described in Table 159 to examine log records or manage the log collection.

Displaying Sensor Logs

The Sensor logs include information related to communication with an IDP sensor if you have deployed a coordinated threat control solution. The local log viewer displays the most recent 5000 log messages (the display limit).

To display Sensor logs:

1. Select **System Log/Monitoring**.

2. Click the **Sensor** tab.

3. Click the **Log** tab.

4. Use the features described in Table 159 to examine log records or manage the log collection.

Related Documentation

- Logging Overview
- Configuring Events to Log
- Using Log Filters

Using Log Filters

This topic describes how to use log filters. It includes the following information:

- Reviewing the Configuration of Predefined Log Format Filters
- Creating a Custom Log Collection Filter
- Example: Using the Source IP Address Filter

Reviewing the Configuration of Predefined Log Format Filters

To view the configuration of predefined log format filters:

1. Select **System > Log/Monitoring**.

2. Click the **Events** tab.

3. Click the **Filter** tab to display the log filters page.

Figure 221 shows the log filters page for Policy Secure.

Figure 222 shows the log filters page for Connect Secure.

4. Click the hyperlinked name of the filter to display its configuration page. You cannot edit the predefined filter named Standard, but you may edit the predefined WELF filters and any other custom filters that appear in the list.

Figure229: Log Filters Page – Policy Secure

Figure230: Log Filters Page – Connect Secure

Creating a Custom Log Collection Filter

If desired, you can create custom log collection filters to change the records displayed or exported. For example, it is common to see administrators use a filter for RADIUS accounting logs. This filter allows only the accounting log message, and it puts the entire message in a comma separated list. The order of the filtered message is: Date, Time, User, Realm, "List of Roles", NAS-ID, Acct-Status, Auth-Type, Attr-Value1, Attr-Value2, Attr-Value3.

Accounting attribute messages are different from authentication attribute messages in that the attribute name is not printed in the log message, but a comma is inserted for every attribute to be logged, even if it is not present.

To create a custom log collection filter:

1. Select **System Log/Monitoring**.

2. Click the **Events** tab.

3. Click the **Filter** tab.

4. Click **New Filter** to display the configuration page.

   Figure 223 shows the configuration page for Policy Secure. The configuration page for Connect Secure is similar.

5. Complete the configuration as described in Table 160.

6. Save the configuration.

Figure231: New Filter Page – Policy Secure

Table161: Filter Settings

Settings	Guidelines
Filter Name	Specify a name that is helpful to you and other administrators in understanding usage for your customer filter.
Make default	Make the filter the default on syslog and archiving configuration pages.
Query	
Start Date	Enter a start date. Click **Earliest Date** to write all logs from the first available date stored in the log file.
End Date	Enter an end date. Click **Latest Date** to write all logs up to the last available date stored in the log file.
Query	Use the Filter Variables Dictionary to insert query expressions in the Query box. Enclose the query value in single quotes. For example, insert the query expression **sourceip=**. Then complete the expression by adding the value **'192.168.0.1'**.
Export Format	Select an export format: • **Standard** (default)—This log filter format logs the date, time, node, source IP address, user, realm, event ID, and message. • **WELF**—This customized WebTrends Enhanced Log Format (WELF) filter combines the standard WELF format with information about the system realms, roles, and messages. • **Custom**—Use the Standard as a template for your custom selection of columns to be included in exports (when log collections are saved to files).

NOTE: Log query filters change only the data displayed (or rows exported). Log format filters change only the data displayed (or columns exported). Use of filters does not affect the log data that has been collected.

Example: Using the Source IP Address Filter

When drilling into logs to verify behavior or troubleshoot an issue with a dual-stack device, it is helpful to redisplay the log collection filtered on the IP address.

To filter on an IP address:

1. Select **System > Log/Monitoring**.

2. Create the filter:

   1. Select **User Access** and then **Filter**.

   2. Define the filter expression, name the filter, and click **Save**. In this example, we create a filter based on source IP address and name it IPv6_Address_Filter:Standard.

3. Use the filter:

   1. Select **Logs** to display the user logs table.

   2. Under View by filter, select **IPv6_Address_Filter:Standard**, as shown in Figure 224.

   3. If desired, under Edit Query, edit the value of the **sourceip=** variable expression to filter on different source IP addresses.

   4. Click **Update** to apply the filter and redisplay the log collection.

Figure232: Using IP Address Filters

Related Documentation

- Displaying System Logs

Displaying User Access Statistics

Every hour, the system logs the peak count of Web users in the previous hour. It displays the hourly counts for the past week on the Statistics page. It writes the report to the system log once a week.

To display user statistics:

1. In the admin console, select **System > Log/Monitoring**.

2. Click the **Statistics** tab to display the page.

   Figure 225 shows the configuration page for Policy Secure.

   Figure 226 shows the configuration page for Connect Secure.

3. Scroll the page to view the data.

Figure233: User Statistics Page – Policy Secure

Monitoring Statistics

Events | User Access | Admin Access | Sensors | Client Logs | SNMP | Statistics

Signed-In Users
Hourly peak load of users

	Sunday 2/24/2013	Monday 2/25/2013	Tuesday 2/26/2013	Wednesday 2/27/2013	Thursday 2/21/2013	Friday 2/22/2013	Saturday 2/23/2013
12:00 am	0	0	0	2	0	0	0
01:00 am	0	0	0	2	0	0	0
02:00 am	0	0	0	2	0	0	0
03:00 am	0	0	0	2	0	0	0
04:00 am	0	0	0	2	0	0	0
05:00 am	0	0	2	2	0	0	0
06:00 am	0	0	2	2	0	0	0
07:00 am	0	0	2	2	0	0	0
08:00 am	0	0	2	2	0	0	0
09:00 am	0	0	2	2	0	0	0
10:00 am	0	0	2	0	0	0	0
11:00 am	0	0	2	0	0	0	0
12:00 pm	0	0	2	0	0	0	0
01:00 pm	0	0	2	0	0	0	0
02:00 pm	0	0	2	0	0	0	0
03:00 pm	0	0	2	0	0	0	0
04:00 pm	0	0	2	0	0	0	0
05:00 pm	0	0	2	0	0	0	0
06:00 pm	0	0	2	0	0	0	0
07:00 pm	0	0	2	0	0	0	0
08:00 pm	0	0	2	0	0	0	0
09:00 pm	0	0	2	0	0	0	0

Figure234: User Statistics Page – Connect Secure

Monitoring Statistics

Events | User Access | Admin Access | Sensors | Client Logs | SNMP | Statistics

Signed-In Users
Hourly peak load of users

	Sunday 5/19/2013	Monday 5/20/2013	Tuesday 5/14/2013	Wednesday 5/15/2013	Thursday 5/16/2013	Friday 5/17/2013	Saturday 5/18/2013
12:00 am	0	0	1	1	1	2	2
01:00 am	0	0	1	1	1	2	2
02:00 am	0	0	1	1	1	2	2
03:00 am	0	0	1	1	1	2	2
04:00 am	0	0	1	1	1	2	2
05:00 am	0	0	1	1	1	2	2
06:00 am	0	0	1	1	1	2	2
07:00 am	0	0	1	1	1	2	2
08:00 am	0	0	1	0	1	2	2
09:00 am	0	0	1	0	0	2	2
10:00 am	0	0	1	0	1	1	0
11:00 am	0	1	2	0	1	1	0
12:00 pm	0	1	2	0	1	1	0
01:00 pm	0	1	3	0	2	1	0
02:00 pm	0	1	2	0	2	1	0
03:00 pm	0	2	2	1	3	1	0
04:00 pm	0	0	2	1	2	2	0
05:00 pm	0	0	2	1	2	2	0
06:00 pm	0	0	1	1	2	2	0
07:00 pm	0	0	1	1	2	2	0
08:00 pm	0	0	1	1	2	2	0
09:00 pm	0	0	1	1	2	2	0

NOTE: Upgrading software clears all statistics. If you configure the system to log statistics hourly, however, older statistics are still available in the log file after an upgrade.

Related Documentation

- Displaying System Status

CHAPTER 38 Troubleshooting Tools

Using the Admin Console Troubleshooting Tools

You can use the admin console troubleshooting tools to investigate user access issues and system issues. The following tools are available through the Maintenance > Troubleshooting pages:

- Policy tracing—Diagnose user access issues.

- Simulation—Connect Secure only. Diagnose user access issues.

- Session recording—Connect Secure only. Work with Pulse Secure Global Support Center (PSGSC) to diagnose user access issues.

- Debug logs—Work with Pulse Customer Support to diagnose system issues.

- RADIUS diagnostic log—Policy Secure only. Diagnose issues with Policy Secure RADIUS server.

- tcpdump—Sniff packet headers to diagnose networking issues.

- Network troubleshooting commands—Use standard network commands, such as ping, traceroute, NSlookup, and other commands to diagnose networking issues.

- Kerberos debugging—Diagnose issues with Kerberos communication.

- System snapshots—Work with Pulse Customer Support to reproduce and diagnose system issues.

- Remote debugging—Enable Pulse Customer Support to access your system directly to help you diagnose system issues.

If the admin console is unavailable, you can use the serial port console to perform some troubleshooting operations, such as use ping and traceroute commands, view logs, create system snapshots, and perform configuration rollbacks and factory resets.

Related Documentation

Using Policy Tracing

It is common to encounter a situation where the system denies a user access to the network or to resources, and the user logs a trouble ticket. You can use the policy tracing utility and log to determine whether the system is working as expected and properly restricting access, or whether the user configuration or policy configuration needs to be updated to enable access in the user's case.

To create a policy trace log:

1. Select **Troubleshooting > User Sessions > Policy Tracing** to display the configuration page.

   Figure 227 shows the policy tracing configuration page for Connect Secure.

Figure235: Policy Tracing Configuration Page

2. Complete the configuration as described in Table 161.

Table162: Policy Trace Configuration Guidelines

Settings	Guidelines
Record Trace File	
User	Specify the username to trace. If you are tracing anonymous access, you can use the asterisks wildcard character (*) because you might not know the internal username the system assigns to the next anonymous session.
Source IP	Specify the source IP address if you know it. If you are able to provide the source IP address, the policy trace log can include events that occur before the user ID is entered into the system.
Realm	Select the realm to trace.
Events to Log	
Pre-Authentication	Logs events related to evaluation of realm rules.
Authentication	Logs events related to authentication.

Role Mapping	Logs events related to role mapping.
Web Policies	Logs events related to web policies.
File Policies	Logs events related to file policies.
SAM Policies	Logs events related to SAM policies.
Telnet/SSH Policies	Logs events related to Telnet/SSH policies.
Terminal Services Policies	Logs events related to terminal services.
VPN Tunneling Policies	Logs events related to VPN tunneling.
Sensor Event Policies	Logs events related to sensor policies.

3. Click **Start Recording**.

Figure 228 shows the policy tracing page with the recording indicator.

Figure236: Policy Tracing Page During Recording

4. Initiate the action you want to trace, such as a user sign in.

5. Click **View Log** to display the policy trace results log.

6. Click **Stop Recording** when you have enough information.

   Figure 229 shows the page with policy trace results.

Figure237: Policy Tracing Results Page

Table 162 describes options for managing the policy trace results log file.

Table163: Post-Trace Options

Control	Guidelines
Delete Trace	Under Events to Log, click **Delete Trace** to clear the results displayed on this page.
Update	Specify a number of rows to display and click **Update** to change the number of rows that are displayed.
Save Log As	Click this button to save the trace results log to a file. This is useful particularly when you are working with the Pulse Secure Global Support Center (PSGSC) to troubleshoot a case.
Clear Log	Click this button to clear the log file from the system.

Related Documentation

- Using the Admin Console Troubleshooting Tools

Using the Simulation Utility

Connect Secure allows you to troubleshoot problems by simulating the events causing the problem. Using the Maintenance > Troubleshooting > User Sessions> Simulation page, you can create virtual user sessions without requiring actual end users to sign in to the device and recreate their problems. In addition, you can also use the Simulation tab to test new authentication and authorization policies before using them in a production environment.

To use the simulator, you must specify which events you want to simulate (for example, you can create a virtual session in which "John Doe" signs into the "Users" realm at 6:00 AM from an Internet Explorer browser). Then, you must specify which events you want to record and log in the simulation. You can log three major types of events to the simulation log:

- **Pre-Authentication**—The system events that are captured will not include any other system related events. Events are merely used as a filtering mechanism to reduce the number of logs and highlight the problem.

- **Role Mapping**—The system events that are captured will not include any other system related events. Events are merely used as a filtering mechanism to reduce the number of logs and highlight the problem.

- **Resource Policies**—The system events that are captured will not include any other system related events. Events are merely used as a filtering mechanism to reduce the number of logs and highlight the problem.

To simulate a user session:

1. In the admin console, choose **Maintenance > Troubleshooting > User Sessions > Simulation**.

   Figure 230 shows the configuration page for Connect Secure.

2. In the Query Name field, enter a name for the query.

3. In the Username field, enter the username of the user whose experience you want to simulate. Note that you may use a wildcard character (*) in place of a username. For example, if your users are signing into an anonymous server, you may want to use the wildcard character (*) since you cannot know the internal username that the system will assign to the user.

4. From the Realm drop-down menu, select the realm of the user whose experience you want to simulate.

5. If you want to determine whether to apply a specific type of resource policy to a user's session, enter the specific resource you want to simulate in the Resource field and select a policy type from the Resource drop-down list. Then:

- If you want to determine whether a user can successfully sign in to the device, select the **Pre-Authentication** check box.

- If you want to determine whether a user can successfully map to a specific role, select the Role Mapping check box. Note that this option controls whether role mapping results are logged to the simulator log, not whether to run role mapping rules. The system always runs role mapping rules, even if you do not select this check box.

- Specify the types of policies you want to log using the check boxes in the Events to Log section.

  For example, if you want to test whether a user can access the Yahoo web site, enter "http://www.yahoo.com" in the Resource field, select Web from the drop-down list, and select the Access check box in the Events to Log section.

6. In the Variables section, use a combination of text and variables to create a custom expression that reflects the exact same values as in the real session of the user who is facing a problem. For example, if you want to create a session in which the user signs in to the device at 6:00 AM, enter "time = 6:00 AM" in the Variables field. For complete instructions on how to create a custom expression. You may also view the syntax for a given variable by clicking the arrow next to it in the Variables Dictionary.

   If you fail to create a custom expression that includes the virtual user's IP address, the system uses your current IP address instead. Also note that if you use the role variable to specify the role of the virtual user (for example, role="Users"), the system ignores results from role mapping rules and assigns the virtual user to the role(s) you specify.

7. Choose one of the following options:

- **Run Simulation**—Runs the specified simulation and creates an on-screen log file.

- **Save Query**—Saves the query.

- **Save Query and Run Simulation**—Runs the specified simulation and also saves it for later use.

8. After running the simulation, choose **Save Log As** to save the simulation results to a text file.

Figure238: Simulation Configuration Page

Related Documentation

- Using the Admin Console Troubleshooting Tools

965

Using the Session Recording Utility

You can use the Session Recording utility to record a trace file that lists a user's actions when accessing a resource or connecting to a client/server application. You do this to troubleshoot issues users might report regarding the Web access or client access.

When you start recording a trace file, the system signs out the specified user and then starts recording all user actions after the user signs in again and is authenticated. Note that the system notifies the user after authentication that user actions are being recorded.

To record a trace file:

1. In the admin console, choose **Maintenance > Troubleshooting > User Sessions >Session Recording**.

   Figure 231 shows the session recording page.

2. Enter the username of the user whose session you want to record.

3. Select the **Web (DSRecord)** check box to record the user's web session and then select the Ignore browser cache check box if you want to ignore cached copies of the problem web site, which the system would not otherwise record as a part of the trace file (optional).

4. Select the **Client/Server (for JCP)** check box to record Java Communication Protocol client/server application sessions (optional).

5. Click **Start Recording**. The system signs out the user.

6. Instruct the user to sign in again and browse to the problem web site or connect to the client/server application.

7. Click **Stop Recording**.

8. Download the trace file(s) from the Current Trace File section:

   1. Click the **DSRecord Log** link to download the Web trace file.

   2. Click the **JCP** or **NCP Client-Side Log** link to download the client/server application trace file.

   3. E-mail the file(s) to Pulse Secure Support for review.

Figure239: Session Recording

Related Documentation

- Using the Admin Console Troubleshooting Tools

Using the Debug Log

The Pulse Secure Global Support Center (PSGSC) might direct you to create a debug log to assist them in helping you debug an issue with the system. The debug log is used only by Pulse Customer Support.

To use debug logging:

1. Select **Troubleshooting > Monitoring > Debug Log** to display the configuration page.

   Figure 232 shows the configuration page for Policy Secure.

   Figure 233 shows the configuration page for Connect Secure.

2. Complete the configuration as described in Table 163.

3. Click **Save Changes**. When you save changes with Debug Logging On selected, the system begins generating debug log entries.

4. Initiate the action you want to debug, such as a user sign in. You can reset the debug log file to restart debug logging if it takes you too long to initiate the action.

5. Click **Save Debug Log** to save the debug log to a file that you can send to Pulse Customer Support. You can clear the log after you have saved it to a file.

6. Unselect **Debug Logging On** and click **Save Changes** to turn off debug logging.

Figure240: Debug Logging Configuration Page – Policy Secure

Figure241: Debug Logging Configuration Page – Connect Secure

Table164: Debug Log Configuration Guidelines

Settings	Guidelines
Current Log Size	Displays the size of the current log file. If it is large, use the controls to save, reset, or clear the log file.
Debug Logging On	Select to turn on debug logging.
Debug Log Size	Specify a maximum debug log file size. The default is 2 MB. The maximum is 250 MB.
Debug Log Detail Level	Specify the debug log detail level. Obtain this from Pulse Customer Support.
Include logs	Select this option to include system logs in the debug log file. Recommended.
Process Names	Specify the process name. Obtain this from Pulse Customer Support.
Event Codes	Specify the event code. Obtain this from Pulse Customer Support.

Related Documentation

- Using the Admin Console Troubleshooting Tools

Using the tcpdump Utility

You can run the tcpdump utility from the admin console.

To use tcpdump:

1. Select **Troubleshooting > Tools > TCP Dump** to display the configuration page.

   Figure 234 shows the configuration page for Policy Secure.

   Figure 235 shows the configuration page for Connect Secure.

2. Complete the configuration as described in Table 164.

3. Click **Start Sniffing** to start the tcpdump process.

4. Initiate the action you want to debug, such as a user sign in.

5. Click **Stop Sniffing** to write the tcpdump output to the screen.

6. Click **Get** to save the output to a file, or click **Delete** to clear the output.

Figure242: TCP Dump Configuration Page – Policy Secure

Figure243: TCP Dump Configuration Page – Connect Secure

Table165: Debug Log Configuration Guidelines

Settings	Guidelines
TCP Dump Status	Displays whether the utility is stopped or running.
Interface	Select the ports on which to sniff.
VLAN Port	Select the VLAN port.

Promiscuous mode	Select a promiscuous mode option.	
Filter	Specify a filter expression. For information about TCP dump filter expressions, see the UNIX man page.	
	Example	**Result**
	tcp port 80	Sniffs packets on TCP port 80.
	port 80	Sniffs packets on TCP or UDP port 80.
	ip	Sniffs the IP protocol.
	tcp	Sniffs the TCP protocol.
	dst #.#.#.#	Sniffs the destination IP address specified, where #.#.#.# is a valid IP address.
	src #.#.#.#	Sniffs the source IP address specified, where #.#.#.# is a valid IP address.
	port 80 or port 443	Sniffs on port 80 or port 443.
	src #.#.#.# and dst #.#.#.#	Sniffs the source and destination IP addresses or hosts specified, where each #.#.#.# represents a valid IP address.
	tcp port 80 or port 443 and dst #.#.#.# and src #.#.#.#	This example shows how to specify multiple parameters to create a filter that sniffs on TCP port 80, or on TCP or UDP port 443, and on the destination and source ports, where each #.#.#.# represents a valid IP address.

Related Documentation

- Using the Admin Console Troubleshooting Tools

Using the Samba Diagnostic Log

The Samba diagnostic log utility allows you to view trace and debug the samba troubleshooting messages on the new AD authentication server. When samba diagnostic logging is enabled, the internal logs related to AD authentication server is generated.

Observe the following guidelines:

- Diagnostic logging affects system performance.

- Must be used only when the admin UI error messages, event logs and admin logs are not very useful.

- Enabling/Disabling samba logs will restart certain modules and user logins may fail during the restart.

- The default debug log setting will generate minimal logs. Enabling debug log with event AAA or AAA::samba along with this feature can generate more logs based on the debug log level.

- Enabling samba logs will cause logs to be generated from all configured AD authentication servers. Logs from multiple AD servers are interleaved and can be identified by the header in each line of the logs.

To use samba diagnostic logging:

1. Select **Troubleshooting > Monitoring > Diagnostic Logs** to display the configuration page.

   Figure 236 shows the configuration page for Policy Secure.

   Figure 237 shows the configuration page for Connect Secure.

2. Complete the configuration as described in Table 165.

3. Click **Save Changes**. When you save changes with Samba Diagnostic Logging On selected, the system begins generating diagnostic log entries.

4. Initiate the action you want to debug, such as a user sign in.

5. Manage the resulting log:

- Click **Save Log** to save the log files in a zipped format.

- Click **Clear Log** to remove previous logs and start diagnostic logging with a fresh file.

- Click **Save And Clear Log** to save the diagnostic log to a file that you can send to Pulse Customer Support. The existing logs in the device will be cleared after saving.

6. Unselect **Samba Diagnostic Logging On** and click **Save Changes** to turn off diagnostic logging.

Figure244: Samba Diagnostic Logging Configuration Page – Policy Secure

Figure245: Samba Diagnostic Logging Configuration Page – Connect Secure

Table166: Samba Debug Log Configuration Guidelines

Settings	Guidelines
Samba Diagnostic Logging On	Select this option to generate logs related to AD server.
Max Diagnostic Log Size	Specify a maximum log file size. Policy Secure • Default (MB)- 10 • Maximum (MB)-100 Connect Secure • Default (MB)- 1000 • Maximum (MB)-1000

Related Documentation

- Using Active Directory

- Using the Admin Console Troubleshooting Tools

Using Network Troubleshooting Commands

You can run common network troubleshooting commands such as arp, ping, ping6, traceroute, traceroute6, NSlookup, and AvgRTTs from the admin console. You can use these connectivity tools to see the network path from the system to a specified server. If a client can ping or traceroute to the access system, and the access system can ping the target server, any remote users should be able to access the server through the access system.

To run network troubleshooting commands:

1. Select **Troubleshooting > Tools > TCP Commands** to display the configuration page.

   Figure 238 shows the configuration page for Policy Secure.

   Figure 239 shows the configuration page for Connect Secure.

2. Complete the configuration as described in Table 166.

3. Click **OK** to run the command and write the output to the screen.

4. Click **Clear** to clear the output.

Figure246: Network Troubleshooting Commands Configuration Page – Policy Secure
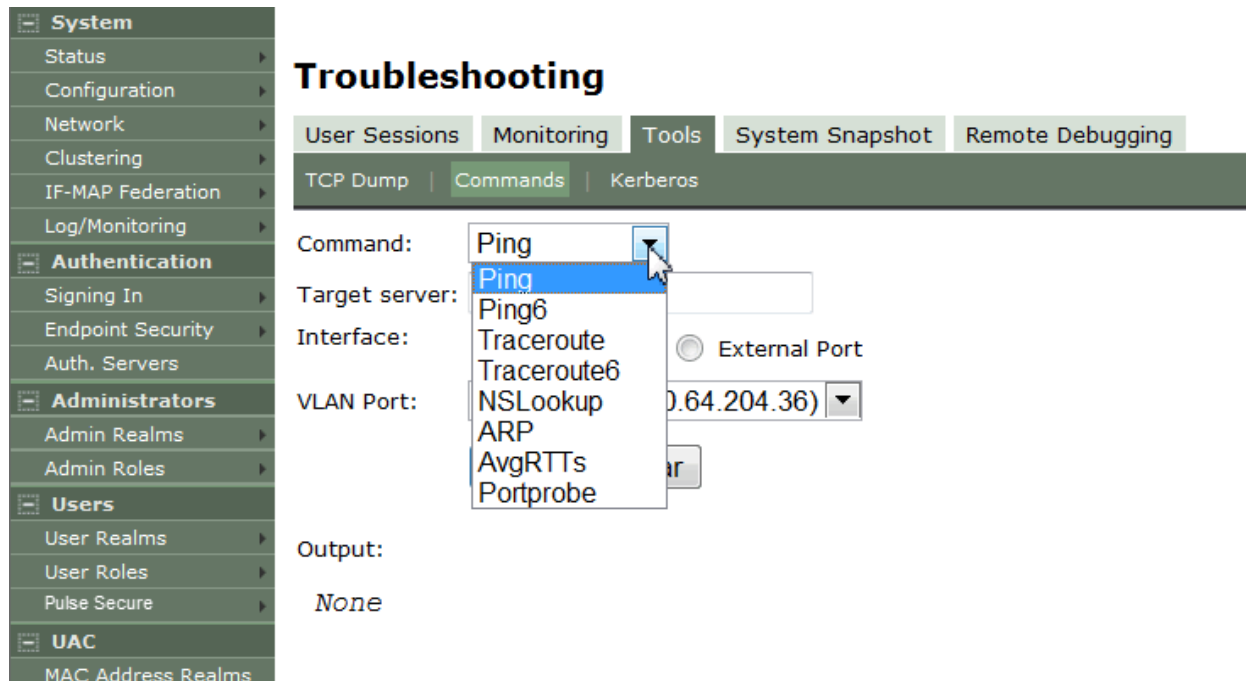
Figure247: Network Troubleshooting Commands Configuration Page – Connect Secure
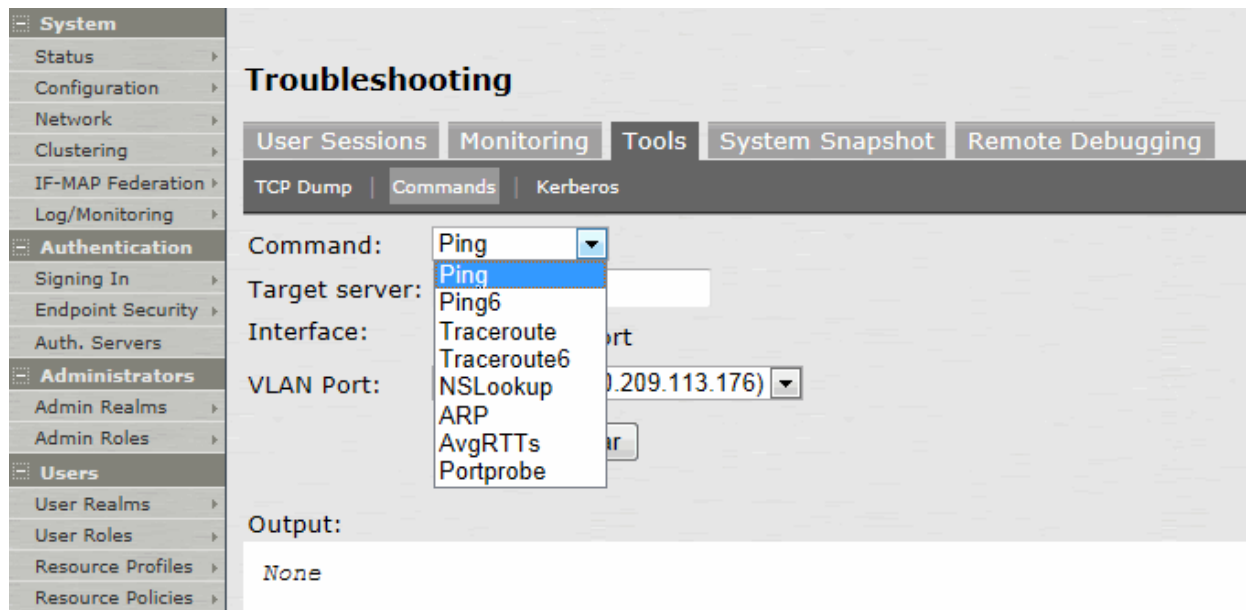
Table167: Network Troubleshooting Commands Configuration Guidelines

Settings	Guidelines
Command	Select a network troubleshooting command: • **Ping/Ping6**—Use the ping command to verify that the system can connect to other systems on the network. In the event of a network failure between the local and remote nodes, you do not receive a reply from a pinged device. In that case, contact your LAN administrator for help. The ping command sends packets to a server and returns the server response, typically a set of statistics including the target server's IP address, the time spent sending packets and receiving the response, and other data. You can ping unicast or multicast addresses, and you must include the target server name in the request. Select **ping** to ping an IPv4 address or hostname. Select **ping6** to ping an IPv6 address. We do not support DNS resolution for hosts with IPv6 addresses. Hence, ping6 does not support pings to hostnames. • **Traceroute/Traceroute6**—Use the traceroute command to discover the path that a packet takes from Connect Secure to another host. Traceroute sends a packet to a destination server and receives an ICMP TIME_EXCEEDED response from each gateway along its path. The TIME_EXCEEDED responses and other data are recorded and displayed in the output, showing the path of the packet round-trip. Select **traceroute** to target an IPv4 address or hostname. Select **traceroute6** to target an IPv6 address. We do not support DNS resolution for hosts with IPv6 addresses. Hence, traceroute6 does not support traceroute to hostnames. • **NSLookup**—Use NSlookup to get detailed information about a name server on the network. You can query on several different types of information, including a server's IP address, alias IP address, start-of-authority record, mail exchange record, user information, well-known services information, and other types of information. • **ARP**—Use the arp command to map IP network addresses to the hardware addresses. The Address Resolution Protocol (ARP) allows you to resolve hardware addresses. To resolve the address of a server in your network, a system sends information about its unique identifier to a server process executed on a server in the intranet. The server process then returns the required address to the client process. • **AvgRTTs**—Use AvgRTTs to display the average round-trip time (RTT) to the localhost. • **Portprobe**—Display the Transmission Control Protocol (TCP) or the User Datagram Protocol (UDP) port status (open or closed).
Target server	Specify the IP address or hostname for the target server.
Interface	Select the interface from which to send the command.
VLAN Port	If you are operating on an IVS license, you can select a VLAN port, to test connectivity to a subscriber intranet.
Output	Displays command output.

Related Documentation

- Using the Admin Console Troubleshooting Tools

- Troubleshooting TCP and UDP Port Status

- Running NSLookup to Test Name Server Connectivity

Troubleshooting TCP and UDP Port Status

Problem	**Description:** The system makes several connections to back-end servers using various port numbers. If communication between the system and the back-end servers stops, it can be difficult to determine the source of the problem.
Solution	You can use the Portprobe command to display the Transmission Control Protocol (TCP) or the User Datagram Protocol (UDP) port status (open or closed).

> ℹ️ **NOTE:** Only the system internal ports, management port and internal VLAN ports support the Portprobe command.

A TCP port can be closed under two conditions:

- The system sends a connection request to the back-end server port and the back-end server closes the connection (sends an RST packet).

- The connection request times out because the back-end server is not found or the back-end server is too busy to respond to the connection request.

If either of these conditions occurs, the system sends a ping command to the back-end server. If the ping command is successful, the back-end server is considered reachable but the back-end server port is closed. If the ping command fails, the back-end server is considered unreachable.

For UDP ports, the system sends a UDP datagram with a ping to the back-end server port. If the back-end server responds with Internet Control Message Protocol (ICMP) port unreachable or ICMP unreachable, the back-end port is considered unreachable. If the back-end server responds with ICMP host unreachable then the back-end server is considered unreachable.

To troubleshoot the TCP or UDP port:

1. Select **Maintenance > Troubleshooting > Tools > Commands**.

2. Select the **Portprobe** command.

3. Select either **TCP** or **UDP**.

4. Enter the target server and port number. You can enter an IP address, hostname or FQDN for the target server.

5. Enter the probe count. This is the number of times the system attempts to communicate with the back-end server port. The default for TCP is one; the default for UDP is five.

6. Enter the probe timeout. This is the number of seconds the system waits for a response from the back-end server port.

7. Select either the internal port or the management port. If the management port is not configured, it is not displayed.

8. If using an internal port, select the internal VLAN port from the list.

9. Click **OK**.

Figure 240 and Figure 241 show an example of a successful and an unsuccessful port probe.

Figure248: Successful TCP Port Probe

Figure249: Unsuccessful UDP Port Probe

Related Documentation

- Using the Admin Console Troubleshooting Tools

Running NSLookup to Test Name Server Connectivity

To run NSLookup to test name server connectivity:

1. In the admin console, choose **Maintenance > Troubleshooting >Tools > Commands**.

   Figure 242 shows the configuration page for Policy Secure.

   Figure 243 shows the configuration page for Connect Secure.

2. From the Command list, select **NSLookup**.

3. Select the type of query to use from the Query Type drop down menu.

4. Enter the query, which is a hostname, an IP address, or other information, depending on your selection of query type.

5. Enter the DNS server name or IP address.

6. If you are operating on an IVS license, you can select a VLAN port, to test connectivity to a subscriber intranet (Connect Secure only).

7. Enter other options.

8. Click **OK** to run the command.

Figure250: Network Troubleshooting Commands Configuration Page – Policy Secure
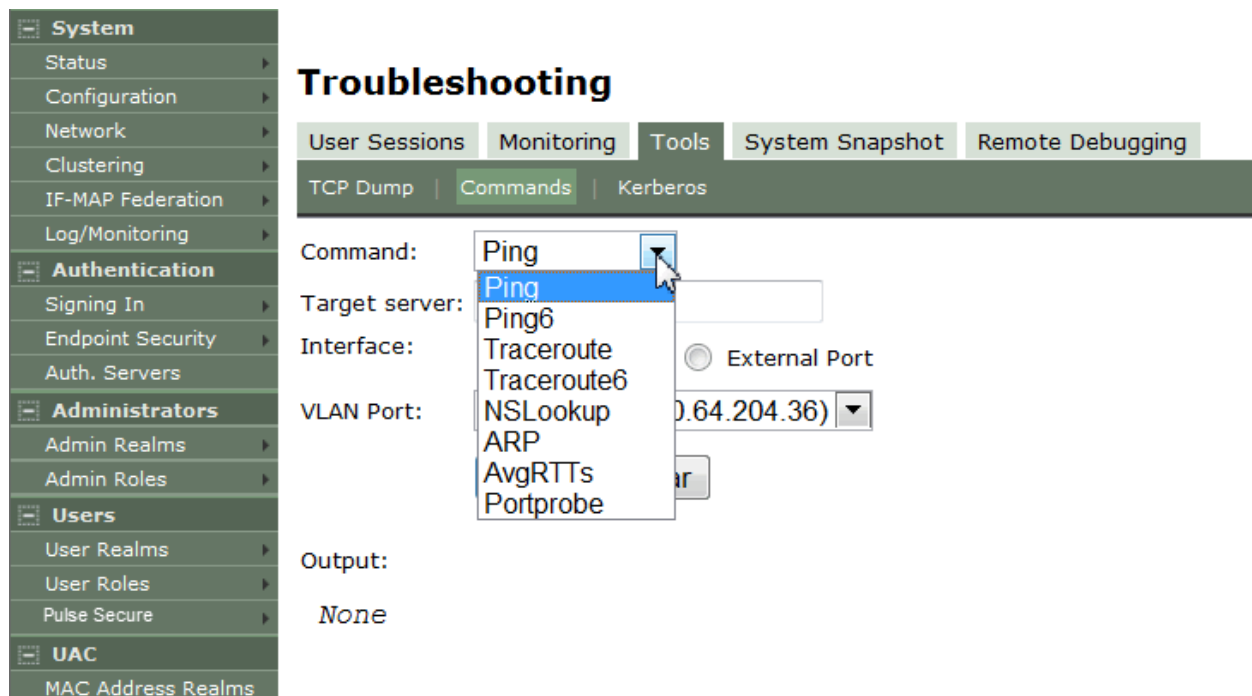
Figure251: Network Troubleshooting Commands Configuration Page – Connect Secure
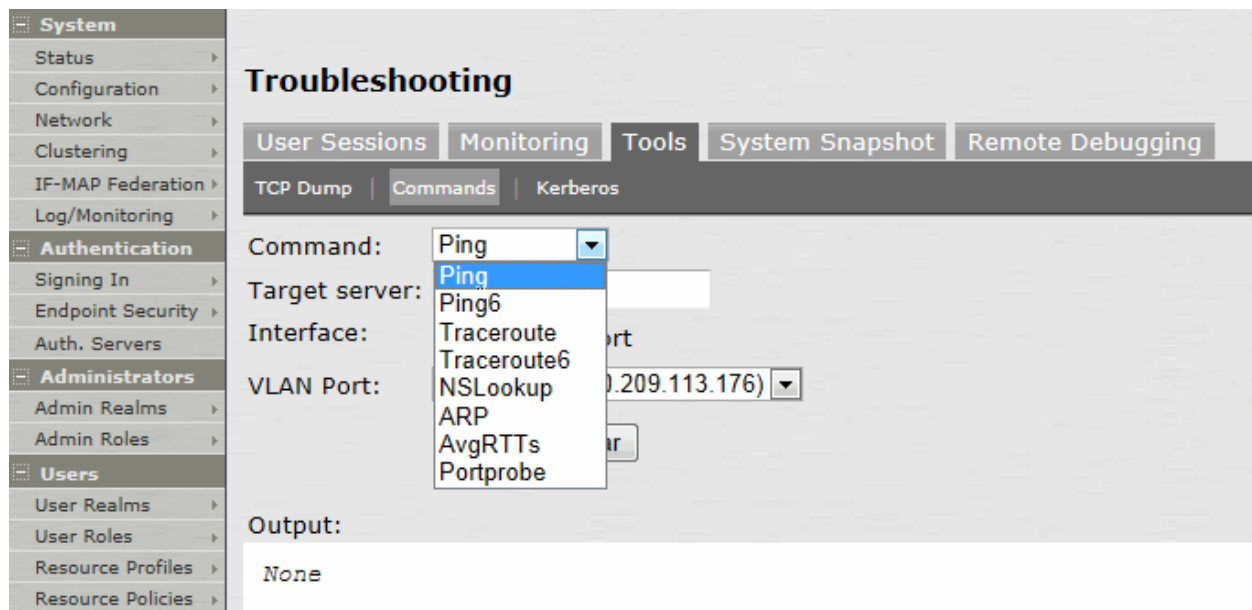
Related Documentation

- Using the Admin Console Troubleshooting Tools

- Troubleshooting VLANs

Using the Kerberos Debugging Utility

You can run the Kerberos debugging utility from the admin console. The utility checks the DNS infrastructure for validity of the Kerberos realms and defined credentials.

To use the Kerberos debugging utility:

1. Select **Maintenance > Troubleshooting > Tools > Kerberos** to display the configuration page.

   Figure 244 shows the configuration page.

2. Complete the configuration as described in Table 167.

3. Click **Run** to start the debugging process.

4. Click **Get** to save the output to a file, or click **Delete** to clear the output.

Figure252: Kerberos Debugging Utility Configuration Page
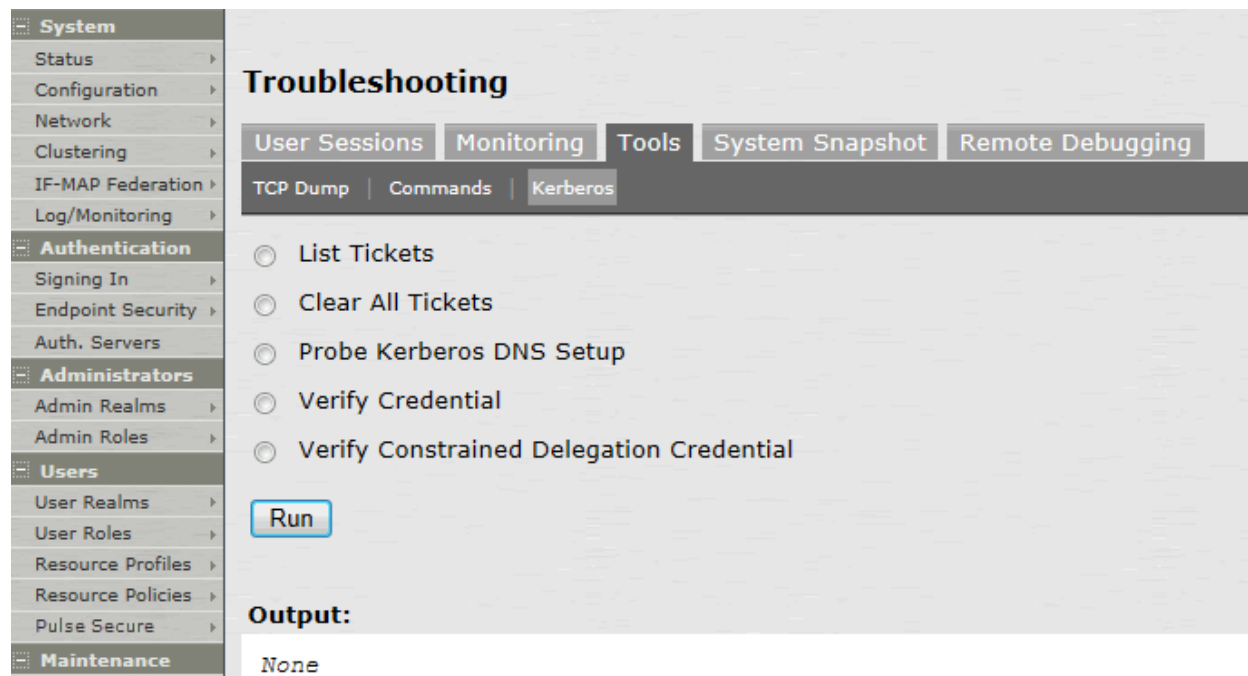
Table168: Kerberos Debugging Utility Configuration Guidelines

Settings	Guidelines
List Tickets	Select this option to list all tickets. Specify the username and the realm name.
Clear All Tickets	Select this option to remove all tickets. Specify the username and realm name.
Probe Kerberos DNS Setup	Select this option to display the configuration elements for the Kerberos DNS test. Specify the realm name and the fully qualified domain name.
Verify Credential	Select this option to verify the Kerberos ticket is valid. Specify the following: - Kerberos Client - Server - Client Realm - Server Realm (Optional) - Client KDC - Server KDC (Optional) - Password For example, if you use Kerberos to verify the username and password provided by the user, this option verifies the credentials it obtains to make sure they belong to a trusted KDB site.

Verify Constrained Delegation Credential	Select this option to verify the Constrained Delegation ticket is valid.
	Specify the following:
	• Kerberos Client
	• Delegation Account
	• Server
	• Client Realm
	• Server Realm (Optional)
	• Client KDC
	• Server KDC (Optional)
	• Password
Output	Displays results of the probe, for example:
| | ```
KDCs for realm matrix.net:
top.matrix.net,top.matrix.net
Operation complete
``` |

Related Documentation

• Using the Admin Console Troubleshooting Tools

Using System Snapshots

A snapshot of the system state captures details that can help Pulse Secure Global Support Center (PSGSC) diagnose system performance problems. The system stores up to ten snapshots, which are packaged into an encrypted "dump" file that you can download and then e-mail to Pulse Customer Support.

To create and manage system snapshots:

1. Select **Maintenance > Troubleshooting > System Snapshot** to display the configuration page.

   Figure 245 shows the configuration page for Policy Secure.

   Figure 246 shows the configuration page for Connect Secure.

2. Complete the configuration and actions as described in Table 168.

Figure253: System Snapshot Configuration Page – Policy Secure

- System
  - Status
  - Configuration
  - Network
  - Clustering
  - IF-MAP Federation
  - Log/Monitoring
- Authentication
  - Signing In
  - Endpoint Security
  - Auth. Servers
- Administrators
  - Admin Realms
  - Admin Roles
- Users
  - User Realms
  - User Roles
  - Pulse Secure
- UAC
  - MAC Address Realms
  - Infranet Enforcer
  - Network Access
  - Host Enforcer
- Maintenance
  - System
  - Import/Export
  - Push Config
  - Archiving
  - Troubleshooting

Troubleshooting

User Sessions | Monitoring | Tools | System Snapshot | Remote Debugging

A snapshot of the system state captures details that can help Pulse Secure   diagnose system performance probl stores up to ten snapshots, which are packaged into an encrypted "dump" file that you can download to a networl email to Pulse Secure.

[Take Snapshot] [Delete]

	Snapshot	Size	Da
☐	Admin generated snapshot (with debuglog)	350262 bytes	20

System snapshot options

☑ Include system config

☑ Include debug log

☑ Schedule automatic snapshots

You should enable automatic scheduled snapshots only when asked to do so by Pulse Secure   Support as part of a troubleshooting operation. Enabling this feature can affect system performance. In most situations, a four-hour snapshot schedule captures the needed data without impacting system performance. Do not set a schedule interval of less than 30 minutes as this can affect system performance.

Take a snapshot every: [____] hours   (0 - 336 hours)

[____] minutes   (0 - 59 minutes)

Max size allocated for periodic snapshots: [____] MB   (60 - 500 MB)

Stop taking snapshots after: Date: [____]   (mm/dd/yyyy)

Time: [____] [AM ▼]   (hh:mm)

Disable debug logs at stop time: ☐

[Save Changes]

Figure254: System Snapshot Configuration Page – Connect Secure

Table169: System Snapshot Configuration Guidelines

Settings	Guidelines
Include system config	Include the system configuration file in the snapshot.
Include debug log	Include debug logs (if any).
Schedule Automatic Snapshots	Enable automatic scheduled snapshots only when asked to do so by Pulse Secure Support as part of a troubleshooting operation. Enabling this feature can affect system performance. In most situations, a four-hour snapshot schedule captures the needed data without impacting system performance. Do not set a schedule interval of less than 30 minutes as this can affect system performance
	Frequency—Specify a frequency in hours and minutes.

	Maximum size—Specify a maximum file size.
	Stop taking snapshots—Specify a date and time to stop the automatic snapshot job.
	Disable debug logs at stop time—Specify that you also want to turn off debug logging when you stop the automatic snapshot job.
Save	Save the configuration.
Take Snapshot	Generate a snapshot now.
Delete	Delete a snapshot file.

Related Documentation

- Using the Admin Console Troubleshooting Tools

Using Remote Debugging

Remote debugging allows Pulse Secure Global Support Center (PSGSC) to directly access this system over a secure connection. You should enable this feature only if you have been requested to do so by Pulse Customer Support in response to an issue that you have reported.

To enable remote debugging:

1. Select **Maintenance > Troubleshooting > Remote Debugging** to display the configuration page.

   Figure 247 shows the configuration page for Policy Secure.

   Figure 248 shows the configuration page for Connect Secure.

2. Complete the configuration and actions as described in Table 169.

Figure255: Remote Debugging Configuration Page – Policy Secure

Figure256: Remote Debugging Configuration Page – Connect Secure

Table170: Remote Debugging Configuration and Action Guidelines

Settings	Guidelines
Debugging Status	Displays whether remote debugging is enabled or disabled.
Debugging Code	Specify a code as instructed by Pulse Customer Support.
Connect to	Specify the fully qualified domain name as instructed by Pulse Customer Support.
Enable Debugging	Click this option to allow remote debugging.

Related Documentation

- Using the Admin Console Troubleshooting Tools

CHAPTER 39 Clustering

Clustering Feature Overview

Clusters define a collection of servers that operate as if they were a single machine. A cluster pair is used to refer to a cluster of two units and a multiunit cluster refers to a cluster of more than two units. Once two or more units are joined in a cluster, they act as one unit.

Figure 249 shows two MAG series devices deployed as a cluster pair.

Figure257: Clustering

- Deployments
- Requirements and Limitations

Deployments

Pulse Secure access management framework supports two types of clusters:

- Failover clusters or active/active clusters– Failover clusters provide high availability (HA). The primary purpose of HA clusters is to provide uninterrupted access to data, even if a server loses network or storage connectivity, or fails completely, or if the application running on the server fails.

  Figure 250 shows an example of an active/active deployment. You can deploy up to 4 nodes in an active/active deployment. MAG-SM360 and SA6500 can support up to 8 nodes.

- Load balancing clusters or active/passive clusters– Load balancing clusters provide scalability and increase availability of Web-based services.

  Figure 251 shows an example of an active/passive deployment. You can deploy up to 2 nodes in an active/passive deployment.

  i NOTE: Pulse Secure access management framework also supports an IPv6 configuration for active/active and active/passive clusters.

Figure258: Active/Active Deployment

Figure259: Active/Passive Deployment

Requirements and Limitations

You must follow these considerations when deploying a cluster:

- Cluster members must run the same software version.

- Cluster members must use the same hardware platform.

- State synchronization must occur only through the internal Network Interface Card (NIC).

- Ensure the cluster communication and resource access must take place over an internal network.

- You can deploy active/active or active/passive clustering only within the same IP subnet.

Related Documentation

- Deploying an Active/Active Cluster

- Deploying an Active/Passive Cluster

- Cluster Licensing

- Cluster Group Communication and Node Monitoring

- Example: Creating an Active/Active Cluster That Supports IPv6 Client Access

- Example: Creating an Active/Passive Cluster that Supports IPv6 Client Access

Cluster Licensing

Starting with system software 7.0, clustering works as follows:

- Clustering does not require a license.

- CL licenses are no longer necessary but are still supported.

- For third-party features, each node in a cluster should have similar license counts.

- A 10 day *cluster grace period* provides license flexibility when a node crashes or loses connectivity with the rest of the cluster.

- Pulse Secure recommends that you distribute your ADD licenses equally across the cluster to avoid losing large number of licenses when a node disconnects from the cluster.

> **NOTE:** You can run Connect Secure with an IVS license in a cluster.

If you are content with your current licensing process, you can carry that forward with the new license scheme. Upgrading may result in a few changes if you have mismatched concurrent user and CL licenses. A cluster that was once "unqualified" may become "qualified" and will be able to support user loads. Your user capacity should not decrease when upgrading to the new licensing scheme.

The maximum number of concurrent users allowed in a cluster is the sum of all user licenses of all connected nodes. If a node disconnects from the cluster (either A/A or A/P), up until the grace period ends the maximum license per each remaining node is their current license plus the minimum of their own license and the licenses of the other nodes. The following examples explain this in more detail.

Example 1: Licenses Distributed Equally Among Nodes

Suppose Node A and Node B are part of a cluster and each node has 500 concurrent user licenses. As long as both nodes are connected, the maximum number of licenses is 1000.

Suppose Node B disconnects from the cluster. Up until the clustering grace period ends, the maximum number of licenses on Node A is 500 (from Node A's original license) + (minimum between Node A (500) and Node B (500)) = 500 + 500 = 1000.

After the grace period ends, the maximum number of licenses on Node A reverts to its original license of 500.

Example 2: Licenses Distributed Unequally Among Nodes

Suppose Node A and Node B are part of a cluster. Node A has 600 ADD licenses and Node B has 400 ADD licenses. As long as both nodes are connected, the maximum number of licenses is 1000.

Suppose Node B disconnects from the cluster. Up until the clustering grace period ends, the maximum number of licenses on Node A is 600 (from Node A's original license) + number of licenses borrowed (minimum between Node A (600) and Node B (400)) = 600 + 400 = 1000. After the grace period ends, the maximum number of licenses on Node A is 600.

Suppose Node A disconnects from the cluster. Up until the clustering grace period ends, the maximum number of licenses on Node B is 400 (from Node B's original license) + number of licenses borrowed (minimum between Node A (600) and Node B (400)) = 400 + 400 = 800. After the grace period ends, the maximum number of licenses on Node B is 400.

Example 3: Licenses Distributed Unequally Among Nodes (Extreme Case)

Suppose Node A and Node B are part of a cluster. Node A has 1000 ADD licenses and Node B has 0 ADD licenses. As long as both nodes are connected, the maximum number of licenses is 1000.

Suppose Node B disconnects from the cluster. Up until the clustering grace period ends, the maximum number of licenses on Node A is 1000 (from Node A's original license) + number of license borrowed (minimum between Node A (1000) + Node B (0)) = 1000 + 0 = 1000. After the grace period ends, the maximum number of licenses on Node A is 1000.

Suppose Node A disconnects from the cluster. Up until the clustering grace period ends, the maximum number of licenses on Node B is 0 (from the Node's original license) + number of license borrowed (minimum between Node A (1000) + Node B (0)) = 0 + 0 = 0. After the grace period ends, the maximum number of licenses on Node B is 0.

Given the scenarios in Examples 2 and 3, we recommend you distribute the licenses equally amongst the nodes.

ADD and CL Licenses

Prior to Connect Secure system software version 7.0, to create an *n*-node cluster supporting *cccc* concurrent users, you were required to purchase one ADD-*cccc*E license for one cluster node, and *n*-1 CL licenses (one for each of the remaining cluster nodes). For example, to create a 4-node cluster supporting 2000 concurrent users, you needed to purchase one ADD-2000E license and 3 CL licenses.

When upgrading to system software version 7 and later, your existing licenses will continue to work. The total concurrent user capacity is still the sum total of all user licenses as long as all nodes are connected. However, when a node disconnects the capacity computation changes as follows:

- Licenses on connected nodes count towards the total cluster capacity.

- If user licenses are present on the computing node, that same number of user licenses can be borrowed from each disconnected node that falls within the cluster grace period.

- If CL licenses are present on the computing node, user licenses can be borrowed from the disconnected nodes so that they total the number of CL licenses.

The following example explains this in more detail.

Suppose you have the following four-node cluster configuration:

- Node A with 1000 user licenses is connected to the cluster.

- Node B with 400 user licenses and 200 CL licenses is connected to the cluster.

- Node C with 500 user licenses and 500 CL licenses is disconnected from the cluster for 17 hours.

- Node D with 1000 user licenses is disconnected from the cluster for 11 days.

The total cluster capacity from Node B's point of view is as follows:

- 1000 licenses from Node A because it is connected.

- 400 licenses from Node B because that's its own license.

- Node C falls within the cluster grace period of 10 days. Using bullet 2 from the above computation notes, since Node B has 400 user licenses it can borrow 400 licenses from Node C's 500 licenses.

Node B also has 200 CL licenses. However, it already borrowed 400 of Node C's 500 user licenses so only 100 of Node C's user licenses remain to be used towards Node B's CL license count.

Node B counts 400+100=500 licenses from Node C.

- Node D has been disconnected from the cluster longer than the cluster grace period. Node B cannot borrow Node D's user licenses.

The total cluster capacity from Node B's point of view is 1000+400+500=1900.

Related Documentation

- Clustering Feature Overview

Deploying an Active/Active Cluster

This example describes the tasks involved in deploying an active/active cluster. It includes the following information:

- Overview
- Network Topology
- Limitations
- Before You Begin
- Configuring an Active/Active Cluster
- Joining Nodes to the Cluster
- Verifying

Overview

An active/active clustering provides high availability and load balancing when deployed with an external load balancer. An active/active cluster deployment requires an external device to distribute the load among the members because the cluster does not have a VIP address. The load balancing devices are equipped with algorithms that balance the load, as well as detect whether a device is down.

Active/active configuration allow for increased aggregate system throughput as well as seamless failover, which is achieved by state synchronization between the two devices for all the configurations so that the devices are virtually identical. Figure 252 shows active/active clustering deployed with an external load balancer.

> NOTE: This feature provides increased throughput and performance for peak load characteristics; however, it does not provide increased scalability beyond the total licensed users.

Network Topology

Active/active clustering can support up to four nodes in a cluster, but are also supported in a LAN environment. Within an active/active cluster, no VIP address is present, and each cluster member has its own network settings. Figure 252 shows an example of active/active deployment.

Figure260: Active/Active Clustering

Limitations

You can deploy active/active clustering only within the same IP subnet.

Before You Begin

We recommend that you deploy a cluster in a staging environment first and then move to a production environment after testing the authentication realm, user role, and resource policy configurations, as well as any applications your end users might access.

You must follow these considerations when deploying a cluster:

Cluster members must run the same software version.

Cluster members must use the same hardware platform.

State synchronization must occur only through the internal Network Interface Card (NIC).

Ensure the cluster communication and resource access must take place over an internal network.

When choosing and configuring a load balancer for your cluster, we recommend that you ensure the load balancer:

- Supports IPsec
- Listens for traffic on multiple ports
- Can be configured to manage traffic using assigned source and destination IP addresses (not destination port)

Configuring an Active/Active Cluster

You use the primary node admin GUI to create the cluster and add members. The primary node is added as part of the cluster creation operation. When you add members, you are prompted for settings unique to the member, such as the name and IP address configuration for the internal and external interfaces. A few additional settings are also unique, namely the

management port and VLAN port settings, so you add these manually after the add node procedure that follows, but before the join cluster operation.

To create a cluster and add members:

1. Select **System > Clustering > Create Cluster** and enter a name for the cluster, a cluster password, and a name for this node, such as Node-*X*.

   You need to enter the password again when specifying additional nodes to join the cluster. All nodes in the cluster use this password to communicate.

   Figure 253 shows the Create New Cluster page for Connect Secure.

Figure261: Create New Cluster Page

Create New Cluster

Join | Create

Type:	MAG-SM360	
Cluster Name:	cluster-1	Name of the cluster to create. Must be alphanumeric, "-", or "_"; must start with a letter and have a maximum of 19 character
Cluster Password:	••••••••	Shared secret among the nodes in the cluster. Must be at least 6 characters long
Confirm Password:	••••••••	Shared secret among the nodes in the cluster. Must match the password you typed in the previous line
Member Name:	node-30	Name of this node in the cluster Must be alphanumeric, "-", or "_"; must start with a letter and have a maximum of 19 character

Create Cluster

? Confirm Create Cluster

Are you sure you want to create a new cluster *SA-Cluster* ?

Please click **Create** to create a new cluster and add this appliance with member name *Node1* to the cluster. Click **Cancel** if you do not want to create a cluster.

Create | Cancel

2. Click **Create Cluster**. When prompted to confirm the cluster creation, click **Create**. After the device initializes the cluster, the Clustering page displays the Status and Properties tabs.

3. Click **Properties**.

   Figure 254 shows the Clustering page for Connect Secure.

Figure262: Clustering Page– Active/Active Configuration

4. Select Active/Active configuration and complete the configuration as described in Table 170. Active/Active configuration is selected by default.

Table171: Clustering Property Settings

Settings	Guidelines
Cluster Name	Specifies a name to identify the cluster.
Configuration Settings	
Active/Passive configuration	Select this option to run a cluster pair in active/passive mode. Then, specify an internal VIP (virtual IP address) and an external VIP if the external port is enabled.

Active/Active configuration	(Default) Select this option to run a cluster pair in active/active mode. Active/Active runs a cluster of two or more nodes in active/active mode using an external load balancer. NOTE: To change a two-unit active/passive cluster to an active/active cluster with more than two nodes, first change the configuration of the two-unit cluster to active/active and then add the additional nodes.
Synchronization Settings	
Synchronize log messages	Select this option to propagate all log messages among the devices in the cluster.
User/Session Synchronization	Select this option to
Configuration only cluster	
Synchronize user sessions	Select this option to synchronize all user session information (for example, instances of access to intranet services) among all the devices in the cluster.
Synchronize last access time for user sessions	Select this option to propagate the latest user access information across the cluster.

NOTE:
- If you select both Synchronize log messages and Synchronize user sessions check boxes, everything is replicated on the cluster nodes, including networking information. Even though networking information, including syslog and SNMP settings, can be configured per node or per cluster, all of the networking information is synchronized between nodes when these two options are set.
- If your cluster node configurations diverge because of changes made to one node while another is disabled or unavailable, the system manages the remerging of the configurations automatically, for up to 16 updates. Beyond the maximum number of allowable updates, you might need to intervene and remerge the configurations manually. In some instances, the system might be unable to remerge the configurations if there is not enough overlapping configuration information between two nodes to manage the internode communication.

  For example, for a two-node cluster in which the two nodes are partitioned from each other because of a network outage, if the internal network IP address of one of the nodes changes in one of the partitions, the two partitions are unable to rejoin, even when the network is repaired. In such a case, you must remerge the configurations manually.

Network Healthcheck Settings	
Number of ARP Ping Failures	Specify the number of ARP ping failures allowed before the internal interface is disabled.
Disable external interface when internal interface fails	Select this option to disable the external interface of the device if the internal interface fails.
Advanced Settings	
	Select the Advanced Settings check box to specify the timeouts for the underlying cluster system. Do not change any values under this setting unless instructed to do so by Pulse Secure Technical Support.

5. Click **Save Changes**.

6. Click **Add Members** to specify additional cluster nodes.

Joining Nodes to the Cluster

The primary node joins the cluster as part of the creation process.

To join additional nodes to the cluster:

1. From an existing cluster member, select the **System > Clustering > Cluster Status** tab and specify the node you want to add to the cluster.

2. From the admin console of the node you want to add to a cluster:

   1. Select the **System > Clustering > Join** tab and enter:

      - The name of the cluster to join

      - The cluster password you specified when defining the cluster

- The IP address of an active cluster member

Figure 255 shows the configuration page for Connect Secure.

Figure263: Join Existing Cluster

2. Click **Join Cluster**. When prompted to confirm joining the cluster, click **Join**.

While the new node synchronizes its state with the existing cluster member, each node's status indicates Enabled, Enabled, Transitioning, or Enabled, Unreachable.

When the node finishes joining the cluster, its Clustering page shows the Status and Properties tabs. After the node joins the cluster, you might need to sign in again.

Verifying

Purpose   Verifying the configuration on **System > Clustering > Cluster Status** page.

Action   Select **System > Clustering > Cluster Status**.

Figure 256 shows the status on the Clustering page for Connect Secure.

Figure264: Clustering Page –Status

Table 171 describes the information displayed on the Status tab and the various management tasks you can perform, such as disabling, enabling, and removing a node from a cluster.

Table172: Clustering Status

GUI Element	Description
Status Information labels	Displays the cluster name, type, configuration, internal VIP, and external VIP for an active/active cluster.
Add Members button	Click this button to specify a node you intend to add to the cluster. You can add multiple nodes at the same time.
Enable button	Click this button to add a node that was previously disabled. When you add a node, all state information is synchronized on the node.
Disable button	Click this button to disable a node within the cluster. The node retains awareness of the cluster but does not participate in state synchronizations or receive user requests unless members sign in to the node, directly.
Remove button	Click this button to remove the selected node or nodes from the cluster. After removal, the node runs in standalone mode.
Member Name column	Lists all nodes belonging to the cluster. You can click on a node's name to modify its name and network settings.
Internal Address column	Shows the internal IP address of the cluster member using Classless Interdomain Routing (CIDR) notation.
External Address column	Shows the external IP address of the cluster member using CIDR notation. Note that this column shows only the external IP address of the cluster leader unless you specify a different address for the node on its individual network settings page, which is accessible by clicking its name in the Member Name column. If you change the external IP address on the Network > Network Settings page, the change affects all cluster nodes.
Status column	Shows the current state of the node: • Green light, Leader—The node is the active member of an active/active cluster and is handling user requests. • Green light, Enabled—The node is handling user requests and participating in cluster synchronization. • Yellow light,Transitioning—The node is joining the cluster. • Red light, Disabled—The node is not handling user requests or participating in cluster synchronization. • Red light/Enabled, Unreachable —The node is enabled but because of a network issue, it cannot be reached. NOTE: A node's state is considered standalone when it is deployed outside of a cluster or after being removed from a cluster.

Notes column	Shows the status of the node's connection to the cluster:
	
	• OK—The node is actively participating in the cluster.
	• Transitioning—The node is switching from the standalone state to the enabled state.
	• Unreachable—The node is not aware of the cluster. A cluster member might be unreachable even when it's online and can be pinged.
	Possible reasons include:
	
	  • Incorrect password.
	  • No information about all cluster nodes.
	  • Configured with a different group communication mode.
	  • Running a different service package version, or the machine is turned off.

Sync Rank column	Specifies the synchronization order for nodes when a node rejoins a cluster. Accepts sync ranks from 0 (lowest rank) to 255 (highest rank). The highest rank takes precedence. If two nodes have identical sync ranks, the alphanumeric rank of the member name is used to determine precedence.
Update button	Updates the sync rank after you change the precedence of the nodes in the Sync Rank column.

Related Documentation

- Clustering Feature Overview

- Cluster Network Connectivity

- Cluster Group Communication and Node Monitoring

- Using a Load Balancer

Deploying an Active/Passive Cluster

This example describes the tasks involved in deploying an active/passive cluster. It includes the following information:

- Overview
- Topology
- Requirements
- Guidelines and Limitations
- Configuring an Active/Passive Cluster
- Joining Nodes to the Cluster
- Verifying

Overview

Active/passive clustering is supported only if the members of the cluster pair are in the same subnet because the VIP address must be shared by both the members. An active/passive cluster configuration provides high availability. Active/passive configurations allows seamless failover without the need to set up any external equipment, which is achieved by state synchronization between the two devices for all the configurations so that the devices are virtually identical. The Pulse Secure access control service uses a virtual IP (VIP) address to address the cluster pair in addition to addressing each device. The IP address takeover (IPAT) approach is used for the VIP address. If the active node fails, the passive node takes over the VIP address and sends a gratuitous Address Resolution Protocol (ARP) message notifying other networking devices that it now owns the VIP address. You should check that other devices in your network, especially the next-hop gateways, will honour the gratuitous ARP messages.

Topology

Figure 257 shows active/passive clustering.

> **(i)** NOTE: This feature provides increased throughput or capacity, but does create redundancy in the case of a failure.

Figure265: Active/Passive Clustering

Requirements

Before you begin:

We recommend that you deploy a cluster in a staging environment first and then move to a production environment after testing the authentication realm, user role, and resource policy configurations, as well as any applications your end users might access.

You must follow these considerations when deploying a cluster:

Cluster members must run the same software version.

Cluster members must use the same hardware platform.

State synchronization must occur only through the internal Network Interface Card (NIC).

Ensure the cluster communication and resource access must take place over an internal network.

When choosing and configuring a load balancer for your cluster, we recommend that you ensure the load balancer:

- Supports IPsec

- Listens for traffic on multiple ports

- Can be configured to manage traffic using assigned source and destination IP addresses (not destination port)

Guidelines and Limitations

- A virtual IP address (VIP) address is shared by all the devices in the cluster. In an active/passive configuration, you configure the VIP address.

- You can deploy active/passive clustering only within the same IP subnet.

Configuring an Active/Passive Cluster

You use the primary node admin GUI to create the cluster and add members. The primary node is added as part of the cluster creation operation. When you add members, you are prompted

for settings unique to the member, such as the name and IP address configuration for the internal and external interfaces. A few additional settings are also unique, namely the management port and VLAN port settings, so you add these manually after the add node procedure that follows, but before the join cluster operation.

To create a cluster and add members:

1. Select **System > Clustering > Create Cluster** and enter a name for the cluster, a cluster password, and a name for this node, such as Node-*X*.

   You need to enter the password again when specifying additional nodes to join the cluster. All nodes in the cluster use this password to communicate.

   Figure 258 shows the Create New Cluster page.

Figure266: Create New Cluster Page

2. Click **Create Cluster**. When prompted to confirm the cluster creation, click **Create**. After the device initializes the cluster, the Clustering page displays the Status and Properties tabs.

3. Click **Properties** and select Active/Passive configuration.

   Figure 259 shows the Clustering page for Connect Secure.

Figure267: Clustering Page– Active/Passive Configuration

4. Complete the configuration as described in Table 172.

Table173: Clustering Property Settings

Settings	Guidelines
Cluster Name	Specifies a name to identify the cluster.
Configuration Settings	
Active/Passive configuration	Select this option to run a cluster pair in active/passive mode. Then, specify an internal VIP (virtual IP address) and an external VIP if the external port is enabled.
Active/Active configuration	Select this option to run a cluster pair in active/active mode. Active/Active runs a cluster of two or more nodes in active/active mode using an external load balancer. NOTE: To change a two-unit active/passive cluster to an active/active cluster with more than two nodes, first change the configuration of the two-unit cluster to active/active and then add the additional nodes.
Synchronization Settings	
Synchronize log messages	Select this option to propagate all log messages among the devices in the cluster.

User/Session Synchronization Configuration only cluster	Select this option to disable synchronization of session data and to replicate only configuration data and user records (for example, web bookmarks, NFS and windows shared files, terminal servers, telnet sessions, SAM, preferences, and passwords). NOTE: Enabling this option limits data transfers between the cluster nodes. User and Session specific limits are only enforced on the node and not across the cluster.
Synchronize user sessions	Select this option to synchronize all user session information (for example, instances of access to intranet services) among all the devices in the cluster.
Synchronize last access time for user sessions	Select this option to propagate the latest user access information across the cluster.

NOTE:

- If you configure your cluster as active/passive, the Synchronize user sessions and Synchronize last access time for user sessions options are automatically selected.

- If you select both Synchronize log messages and Synchronize user sessions check boxes, everything is replicated on the cluster nodes, including networking information. Even though networking information, including syslog and SNMP settings, can be configured per node or per cluster, all of the networking information is synchronized between nodes when these two options are set.

- If your cluster node configurations diverge because of changes made to one node while another is disabled or unavailable, the system manages the remerging of the configurations automatically, for up to 16 updates. Beyond the maximum number of allowable updates, you might need to intervene and remerge the configurations manually. In some instances, the system might be unable to remerge the configurations if there is not enough overlapping configuration information between two nodes to manage the internode communication.

  For example, for a two-node cluster in which the two nodes are partitioned from each other because of a network outage, if the internal network IP address of one of the nodes changes in one of the partitions, the two partitions are unable to rejoin, even when the network is repaired. In such a case, you must remerge the configurations manually.

Network Healthcheck Settings

Number of ARP Ping Failures	Specify the number of ARP ping failures allowed before the internal interface is disabled.
Disable external interface when internal interface fails	Select this option to disable the external interface of the device if the internal interface fails.

Advanced Settings

	Select the Advanced Settings check box to specify the timeouts for the underlying cluster system. Do not change any values under this setting unless instructed to do so by Pulse Secure Technical Support.

5. Click **Save Changes**. After Connect Secure initializes the active/passive cluster, the Clustering page displays the Status and Properties tabs.

6. Click **Add Members** to specify additional cluster nodes.

   shows the page for Connect Secure.

Figure268: Add Cluster Member Page

Joining Nodes to the Cluster

The primary node joins the cluster as part of the creation process.

To join additional nodes to the cluster:

1. From an existing cluster member, select the **System > Clustering > Cluster Status** tab and specify the node you want to add to the cluster.

2. From the admin console of the node you want to add to a cluster:

   1. Select the **System > Clustering > Join** tab and enter:

      - The name of the cluster to join

      - The cluster password you specified when defining the cluster

      - The IP address of an active cluster member

   2. Click **Join Cluster**. When prompted to confirm joining the cluster, click **Join**.

   While the new node synchronizes its state with the existing cluster member, each node's status indicates Enabled, Enabled, Transitioning, or Enabled, Unreachable.

   When the node finishes joining the cluster, its Clustering page shows the Status and Properties tabs. After the node joins the cluster, you might need to sign in again.

Verifying

Purpose	Verifying the configuration on **System > Clustering > Cluster Status** page.
Action	Select **System > Clustering > Cluster Status**.

Figure 261 shows the status on the Clustering page for Connect Secure.

Figure269: Clustering Page –Status

Table 173 describes the information displayed on the Status tab and the various management tasks you can perform, including disabling, enabling, and removing a node from a cluster.

Table174: Clustering Status

GUI Element	Description
Status Information labels	Displays the cluster name, type, configuration, internal VIP, and external VIP for an active/passive cluster.
Add Members button	Click this button to specify a node you intend to add to the cluster. You can add multiple nodes at the same time.
Enable button	Click this button to add a node that was previously disabled. When you add a node, all state information is synchronized on the node.
Disable button	Click this button to disable a node within the cluster. The node retains awareness of the cluster but does not participate in state synchronizations or receive user requests unless members sign in to the node, directly.
Remove button	Click this button to remove the selected node or nodes from the cluster. After removal, the node runs in standalone mode.
Fail-Over VIP	Click this button to failover the VIP to the other node in the active/passive cluster. Only available if cluster is configured as active/passive.
Member Name column	Lists all nodes belonging to the cluster. You can click on a node's name to modify its name and network settings.
Internal Address column	Shows the internal IP address of the cluster member using Classless Interdomain Routing (CIDR) notation.
External Address column	Shows the external IP address of the cluster member using CIDR notation. Note that this column shows only the external IP address of the cluster leader unless you specify a different address for the node on its individual network settings page, which is accessible by clicking its name in the Member Name column. If you change the external IP address on the Network > Network Settings page, the change affects all cluster nodes.
Status column	Shows the current state of the node: • Green light, Leader—The node is the active member of an active/active cluster and is handling user requests. • Green light/enabled—The node is handling user requests and participating in cluster synchronization. • Yellow light/transitioning—The node is joining the cluster. • Red light/disabled—The node is not handling user requests or participating in cluster synchronization. • Red light/enabled, unreachable —The node is enabled but because of a network issue, it cannot be reached. NOTE: A node's state is considered standalone when it is deployed outside of a cluster or after being removed from a cluster.
Notes column	Shows the status of the node's connection to the cluster: • OK—The node is actively participating in the cluster. • Transitioning—The node is switching from the standalone state to the enabled state. • Unreachable—The node is not aware of the cluster. A cluster member might be unreachable even when it's online and can be pinged. Possible reasons include: its password is incorrect, it doesn't have information about all cluster nodes, it's configured with a different group communication mode, it is running a different service package version, or the machine is turned off.
Sync Rank column	Specifies the synchronization order for nodes when a node rejoins a cluster. Accepts sync ranks from 0 (lowest rank) to 255 (highest rank). The highest rank takes precedence. If two nodes have identical sync ranks, the alphanumeric rank of the member name is used to determine precedence.
Update button	Updates the sync rank after you change the precedence of the nodes in the Sync Rank column.

Related Documentation

- Clustering Feature Overview
- Cluster Network Connectivity
- Cluster Group Communication and Node Monitoring

Using a Load Balancer

Overview

In active/active mode, you have the option of using an external load balancer with a cluster. If you do use a load balancer, all the nodes actively handle user requests sent by the load balancer or round-robin DNS. The load balancer hosts the cluster VIP and routes user requests to a node defined in its cluster group based on source-IP routing. If a node goes off line, the load balancer adjusts the load on the active nodes. Users do not need to sign in again, however some session information entered a few seconds before the active machine went off-line, such as cookies and passwords, may not have been synchronized on the current device, in which case users may need to sign in to back-end Web servers again.

The cluster itself does not perform any automatic fail-over or load-balancing operations, but it does synchronize state data (system, user, and log data) among cluster members. When an off-line device comes back online, the load balancer adjusts the load again to distribute it among all active members. This mode provides increased throughput and performance during peak load but does not increase scalability beyond the total number of licensed users.

The system synchronizes state data on all nodes if you add or delete the host entry on the Network Settings pages. If you add or delete the host entry using the Clustering tab for a cluster member, the state data affects only the node and the system does not synchronize the data across the entire cluster.

The system hosts an HTML page that provides service status for each node in a cluster. External load balancers can check this resource to determine how to effectively distribute the load among all the cluster nodes.

The following figure illustrates an active/active cluster configuration in which the devices have enabled external ports.

This active/active cluster configuration is deployed behind an external load balancer. You can deploy a cluster pair or multi-unit cluster in active/active mode. User requests are directed to the cluster VIP defined on the load balancer, which routes them to the appropriate machine.

Figure270: Active/Active Configuration

Requirements and Limitations

When choosing and configuring a load balancer for your cluster, we recommend that you ensure the load balancer:

- Supports IPsec

- Listens for traffic on multiple ports

- Can be configured to manage traffic using assigned source and destination IP addresses (not destination port)

Configuring a Load Balancer

The load balancer is configured externally.

Health Checking a Server from a Load Balancer

Purpose The system hosts an HTML page that provides service status for each node in a cluster. External load balancers can check this resource to determine how to effectively distribute the load among all the cluster nodes.

Action To perform the Layer 7 health check for a node:

- In a browser—Enter the URL: https://Pulse Connect Secure Controller-Hostname/dana-na/healthcheck/healthcheck.cgi?status=SBR

  This returns the Steel Belted Radius (SBR) status (SBR_AVAILABLE), either HTTP Status 200 OK or 500 Internal Error. If SBR_AVAILABLE is 0, the SBR is down. If SBR_AVAILBLE is 1, then SBR is up and performing transactions.

- https://Pulse Connect Secure Controller-Hostname/dana-na/healthcheck/healthcheck.cgi?status=all

  This returns either HTTP Status 200 OK or 500 Internal Error. If this returns HTTP Status 200 OK, the following additional parameters are shown:

Parameter Name	Value	Description

CPU-UTILIZATION	0-100	Specifies the CPU utilization percentage (0-100).
VPN-TUNNEL-COUNT	integer	Specifies the number of concurrent Pulse IPSec, Network Connect and IKEv2 tunnels.
SWAP-UTILIZATION	integer	Specifies the swap utilization percentage of the device (0-100).
DISK-UTILIZATION	integer	Specifies the used disk space percentage (0-100).
SSL-CONNECTION-COUNT	integer	Specifies the total number of SSL connections.
USER-COUNT	integer	Specifies the total number of licensed users logged in to the device. This does not include any MAC address users or Radius users.

The following example performs the Layer 7 health check from an external load balancer:

- GET /dana-na/healthcheck/healthcheck.cgi?status=all HTTP/1.1\r\nUser-Agent: Mozilla/4.0 (compatible; MSIE 8.0; Windows NT 6.1; WOW64; Trident/4.0; SLCC2; .NET CLR 2.0.50727; .NET CLR 3.5.30729; .NET CLR 3.0.30729; Media Center PC 6.0; .NET4.0C; MS-RTC LM 8; .NET4.0E)\r\nHost: localhost\r\n\r\n\r\n

The concept of receive string is used by health check. The receive string is configured on the load balancer to decide whether or not to mark a node as active or inactive. It is a regular expression that checks for a value present in the response. For example, Connect Secure sends a page to the load balancer that has USER-COUNT=25 indicating that the number of active licensed users on that device is 25.

A receive string of USER-COUNT\=([0-9]|[0-9][1-9]|100); means check if USER-COUNT is between 0 and 100. In this example, 25 is between 0 and 100 and the load balancer marks the device as active and considers it for load balancing. Suppose more users log in to the device and it now sends USER-COUNT=150 to the load balancer. This value is now out of the range and the load balancer marks that device as inactive and stop sending traffic to it. Active sessions will continue to pass through the device however.

Related Documentation

- Clustering Feature Overview
- Cluster Network Connectivity
- Cluster Group Communication and Node Monitoring
- Deploying an Active/Active Cluster
- *Deploying an Active/Passive Cluster*

Admin Console Procedures

- Creating a Cluster
- Adding an Node to a Cluster Through the Admin Console
- Deleting a Cluster
- Failing Over the VIP to Another Node
- Changing the IP Address of a Cluster Node
- Adding Multiple Cluster Nodes
- Re-Adding a Node to a Cluster
- Restarting or Rebooting Cluster Nodes

Creating a Cluster

To create a cluster and add members:

---

1. Select **System > Clustering > Create Cluster** and complete the configuration as described in Table 174.

   Figure 263 shows the Create New Cluster page.

2. Click **Create Cluster**. When prompted to confirm the cluster creation, click **Create**. After the device initializes the cluster, the Clustering page displays the Status and Properties tabs.

Figure271: Create New Cluster Page

Table175: Cluster Settings

Settings	Actions
Cluster Name	Specifies a name to identify the cluster.
Cluster Password	Specifies the cluster password. You need to enter the password again when specifying additional nodes to join the cluster. All nodes in the cluster use this password to communicate.
Confirm Password	Specifies the password that is confirmed.
Member Name	Specifies the name of the member.

Adding an Node to a Cluster Through the Admin Console

Before you can add a node to a cluster (through either the Web or the serial console), you need to make its identity known to the cluster. Note that if a node has a cluster license key, it has only a Clustering > Join tab.

To add a node to a cluster through its admin console:

1. From an existing cluster member, select **System > Clustering > Cluster Status**, and specify the node you want to add to the cluster.

2. From the admin console of the node you want to add to a cluster, select **System > Clustering > Join**, and enter:

   - The name of the cluster to join

   - The cluster password you specified when defining the cluster

   - The IP address of an active cluster member

3. Click **Join Cluster**. When you are prompted to confirm joining the cluster, click **Join**. After the node joins the cluster, you may need to sign in again.

Figure 264 shows the Join Cluster page.

Figure272: Join Cluster Page

While the new node synchronizes its state with the existing cluster member, each node's status on the Status page indicates Enabled, Enabled; Transitioning; or Enabled, Unreachable.

Deleting a Cluster

If you delete a cluster, all of the nodes begin running as standalone systems.

To delete a cluster:

1. From the admin console of an active cluster member, select the **System > Clustering > Properties** page.

2. Click the **Delete Cluster** button.

3. Click **Save Changes**.

Figure 265 shows the properties for the Clustering page.

Figure273: Clustering Page –Properties

Failing Over the VIP to Another Node

In an active/passive cluster, you might need to fail the VIP to the other node, regardless of which node you are currently using.

To fail over the VIP:

1. Select **System > Clustering > Cluster Status** from the admin console.

2. Click the **Fail-Over VIP** button to move to the other node. The Fail-Over VIP button is a toggle button, so you can move from one node to the other, regardless of which is the leader. The fail over occurs immediately.

   NOTE: VIP failover does not occur when the management port fails.

Figure 266 shows the fail-over VIP option on the Clustering page.

Figure274: Clustering Page –Status

Changing the IP Address of a Cluster Node

Changing the IP address of a cluster while it belongs to a cluster is not supported. In order to change the IP address, you must first remove it from the cluster, update the IP address and then add it back.

> ℹ️ NOTE: If you attempt to change the IP address of a node while it belongs to a cluster, unpredictable results might occur.

For example:

1. Select **System > Clustering > Cluster status**.

2. Select the check box for the name of the node whose IP address you want to change.

3. Click **Remove**.

4. After the node is removed, sign in to that node, change its IP address and click **Save Changes**.

5. In the main node, add the changed node to the cluster configurations.

6. Log in to the changed node and rejoin the cluster.

The following procedure is a model for changing both node IP addresses in an active/passive cluster:

1. Select **System > Clustering > Cluster status**.

2. Click **Delete Cluster**.

3. Change the IP address of each node.

4. Log in to the main node and re-create the cluster, changing it from active/active to active/passive and defining the internal and/or external VIP addresses.

5. Add the other node to the cluster configurations.

6. Log in to the passive node and add it to the cluster.

Adding Multiple Cluster Nodes

To add multiple nodes to a cluster:

1. Select **System > Clustering > Cluster Status**.

2. Click **Add Members**.

3. Enter the node name and internal IP address.

4. Modify or add the default internal netmask and internal gateway addresses, if necessary.

5. Click **Add**.

Figure 267 shows the Add Cluster Member page.

Figure275: Add Cluster Member Page

6. Repeat the process until you have added all of the nodes.

7. Click **Save Changes** to save the node configurations.

The system automatically enables the added nodes, even if they are unreachable.

Re-Adding a Node to a Cluster

With some maintenance operations, you might need to remove a node from a cluster, then re-add and re-join it to the cluster.

When a node joins a cluster, all of its node-specific settings (including network interface addresses, route tables, virtual ports, ARP caches, VLAN interface, SNMP settings) are overwritten by the corresponding configuration setting it receives from the cluster.

To populate the newly joined node with the correct node-specific settings:

1. Add the node to the cluster.

2. On any of the existing nodes in the cluster, manually configure the appropriate node-specific settings for the newly added node by selecting the node from the menu in the settings page.

3. Add the node to the cluster.

When the node joins the cluster, it receives its newly configured node-specific settings from the cluster.

> **NOTE:** You configure the node-specific settings for the newly added node manually because binary import options are not useful. The only recommended binary import option into a cluster is "Import everything except network settings and licenses" from the Maintenance > Import/Export > Configuration page, which restores cluster-wide configuration (sign-in, realms, roles, resource policies etc.) from a backup binary file. Because this option skips node-specific settings, you must perform step 2 manually to populate the newly joined node with the right set of node-specific settings.

Restarting or Rebooting Cluster Nodes

When you create a cluster of two or more nodes, the clustered nodes act as a logical entity. When you reboot one of the nodes using either the serial console or the admin console, all nodes in the cluster restart or reboot.

To reboot only one node:

1. Select **System > Clustering > Status** to disable the node you want to restart or reboot within the cluster.

2. Select **Maintenance > System > Platform**.

3. Reboot the node, then enable the node within the cluster again.

   The system reconciles session state with the Infranet Enforcer upon restart or cluster failover. If the Infranet Enforcer is running ScreenOS 6.0r2 or later, an Policy Secure restart or failover does not interrupt network traffic of existing sessions, as long as the restart or failover occurs within two minutes.

   Figure 268 shows the System Maintenance page.

Figure276: System Maintenance

Related Documentation

- Creating a Cluster

- Clustering Feature Overview

- Joining Nodes to the Cluster

- Modifying the Cluster Properties

Creating a Cluster

To create a cluster and add members:

1. Select **System > Clustering > Create Cluster** and complete the configuration as described in Table 174.

   Figure 269 shows the Create New Cluster page.

2. Click **Create Cluster**. When prompted to confirm the cluster creation, click **Create**. After the device initializes the cluster, the Clustering page displays the Status and Properties tabs.

Figure277: Create New Cluster Page

Table176: Cluster Settings

Settings	Actions
Cluster Name	Specifies a name to identify the cluster.
Cluster Password	Specifies the cluster password. You need to enter the password again when specifying additional nodes to join the cluster. All nodes in the cluster use this password to communicate.
Confirm Password	Specifies the password that is confirmed.
Member Name	Specifies the name of the member.

Related Documentation

- Clustering Feature Overview

- Joining Nodes to the Cluster

- Modifying the Cluster Properties

Joining Nodes to the Cluster

The primary node joins the cluster as part of the creation process.

To join additional nodes to the cluster:

1. From an existing cluster member, select the **System > Clustering > Cluster Status** tab and specify the node you want to add to the cluster.

2. From the admin console of the node you want to add to a cluster, select the **System > Clustering > Join** tab and enter:

   - The name of the cluster to join

   - The cluster password you specified when defining the cluster

   - The IPv4 address for the internal port of an active cluster member

   Figure 270 shows the Join cluster page.

Figure278: Join Existing Cluster

3. Click **Join Cluster**. When prompted to confirm joining the cluster, click **Join**.

The join cluster operation validates IPv4 and IPv6 settings for all the physical ports (internal, external, and management) against those present in the existing cluster. For example, the external port IPv6 settings present on Node-*Y* are compared against external port IPv6 settings that were specified for the Node-*Y* add member operation entered on the primary node (Node-*X*). If there is a mismatch, the join operation fails with an appropriate error message.

While the new node synchronizes its state with the existing cluster member, each node's status indicates Enabled, Enabled; Transitioning; or Enabled, Unreachable.

When the node finishes joining the cluster, its Clustering page shows the Status and Properties tabs. After the node joins the cluster, you might need to sign in again.

Related Documentation

- Creating a Cluster

- Clustering Feature Overview

- Modifying the Cluster Properties

Modifying the Cluster Properties

To modify the cluster properties:

1. Select **System > Clustering > Properties**.

Figure279: Clustering Properties Page

2. Complete the configuration as described in Table 176.

Table177: Clustering Property Settings

Settings	Actions
Cluster Name	Identifies the cluster.
Configuration Settings	
Active/Passive configuration	Runs a cluster pair in active/passive mode. Then specify an internal VIP (virtual IP address) and an external VIP if the external port is enabled.

Active/Active configuration	(Default) Runs a cluster pair in active/active mode. This configuration runs a cluster of two or more nodes in active/active mode using an external load balancer.
	NOTE: To change a two-unit active/passive cluster to an active/active cluster with more than two nodes, first change the configuration of the two-unit cluster to active/active and then add the additional nodes.

Synchronization Settings

Synchronize log messages	Propagates all log messages among the devices in the cluster.
User/Session Synchronization Configuration only cluster	Select this option to disable synchronization of session data and to replicate only configuration data and user records (for example, web bookmarks, NFS and windows shared files, terminal servers, telnet sessions, SAM, preferences, and passwords). NOTE: • Enabling this option limits data transfers between the cluster nodes. User and Session specific limits are only enforced on the node and not across the cluster. • Do not activate this feature when the user sessions are in progress. • Session failover is not supported in configuration only cluster mode.
Synchronize user sessions	Synchronizes all user session information (for example, instances of access to intranet services) among all the devices in the cluster.
Synchronize last access time for user sessions	Propagates the latest user access information across the cluster.

NOTE:
- If you select both the Synchronize log messages and Synchronize user sessions check boxes, everything is replicated on the cluster nodes, including networking information. Even though networking information, including syslog and SNMP settings, can be configured per node or per cluster, all of the networking information is synchronized between nodes when these two options are set.
- If your cluster node configurations diverge because of changes made to one node while another is disabled or unavailable, the system manages the remerging of the configurations automatically for up to 16 updates. Beyond the maximum number of allowable updates, you might need to intervene and remerge the configurations manually. In some instances, the system might be unable to remerge the configurations if there is not enough overlapping configuration information between two nodes to manage the internode communication.

  For example, for a two-node cluster in which the two nodes are partitioned from each other because of a network outage, if the internal network IP address of one of the nodes changes in one of the partitions, the two partitions are unable to rejoin, even when the network is repaired. In such a case, you must remerge the configurations manually.

Network Healthcheck Settings

Number of ARP Ping Failures	Specifies the number of ARP ping failures allowed before the internal interface is disabled.
Disable external interface when internal interface fails	Disables the external interface of the device if the internal interface fails.

Advanced Settings

	Specifies the timeouts for the underlying cluster system. Do not change any values under this setting unless instructed to do so by Pulse Secure Technical Support.

3. Click **Save Changes**.

Related Documentation

- Creating a Cluster

- Clustering Feature Overview

Synchronizing the Cluster State

State synchronization occurs only by means of the internal network interface cards (NICs), and each cluster member is required to possess the cluster password to communicate with other members. Cluster members synchronize data when there is a state change on any member. Cluster state data is either persistent—permanently stored on the device—or transient—stored

on the device only for the user's session. State data is divided into the following major categories:

- **System state**—This state is persistent and does not change often.

  - Network settings

  - Authentication server configurations

  - Authorization group configurations, such as access control list, bookmark, messaging, and application data

- **User profile**—This data can be either persistent or transient, depending on whether or not you have enabled persistent cookies and persistent password caching. If you have not enabled these features, then the data is transient and falls into the next category.

  - User bookmarks—persistent

  - Persistent user cookies—if the persistent cookies feature is enabled, the device stores user cookies for web sites that issue persistent cookies

  - Persistent user passwords—if the password caching feature is enabled, the user can choose to store her credentials for applications and web sites

- **User session**—This state is transient and dynamic. The user session consists of the following data:

  - The user session cookie

  - Transient user profile information, which includes cookies and passwords stored only for during the user's session

- **Monitoring state**—This persistent information consists of log messages.

  Whether you deploy a cluster in active/passive or active/active mode, the Connect Secure is responsible for synchronizing data between cluster members. The Connect Secure synchronizes all system data, user profile data, and the user session cookies immediately, so if one cluster member goes off-line, users do not need to sign in to the device again. A small amount of latency occurs when the device synchronizes user session profile and monitoring state data, so if a member goes off-line, the user may need to sign in to some back-end Web applications again and administrators may not have access to the logs on the failed machine.

  If you notice too much latency occurring on one or more nodes, you might need to change the Clustering Timeouts Settings.

  When you add the device to a cluster, the cluster leader does not send log messages to the new member. Log messages are also not synchronized between cluster members when one member restarts its services or when an offline machine comes back online. Once all machines are online, however, log messages are synchronized.

  > **NOTE:** If you are running an active/active cluster, you must not allow the cluster to switch to active/passive mode unless the active/active and active/passive clusters share compatible spread timeout settings.

You may also configure synchronization settings to improve performance:

- **Specify the synchronization protocol**—When running three or more devices in a multi-unit or multi-site cluster, you can choose to use the synchronization protocol (Unicast, Multicast, or Broadcast) that best suits your network topology.

  > **NOTE:** Multicast or broadcast for cluster communication is not supported on the MAG Series Pulse Secure Gateways.

- **Synchronize log messages**— Log messages may create a huge payload on the network and affect cluster performance. This option is disabled by default.

- **Synchronize user sessions**—This option synchronizes all user session information (instances of access to intranet services, for example) among all devices in the cluster.

  You must select this option if your cluster is an IF-MAP client. If you do not select this option, your IF-MAP client may not work as expected.

- **Synchronize last access time for user sessions**—This option allows you to propagate user access information in the cluster. If this option is the sole synchronization item among the cluster nodes, you can significantly reduce CPU impact among the cluster devices.

  ⓘ NOTE:

  - If you configure your cluster as active/passive, the Synchronize user sessions and Synchronize last access time for user sessions options are automatically checked.

  - If you select both the both Synchronize log messages and Synchronize user sessions check boxes, everything is replicated on the cluster nodes, including networking information. Even though networking information, including syslog and SNMP settings, can be configured per node or per cluster, all of the networking information is synchronized between nodes when these two options are set.

  - If your cluster node configurations have diverged due to changes made to one node while another is disabled or unavailable, the devices manages the remerging of the configurations automatically, for up to 16 updates. Beyond the maximum number of allowable updates, you may need to intervene and remerge the configurations manually. In some instances, the devices may be unable to remerge the configurations if there is not enough overlapping configuration information between two nodes to manage the internode communication.

    For example, given a two-node cluster in which the two nodes are partitioned from each other because of a network outage, if the internal network IP address of one of the nodes gets changed in one of the partitions, the two partitions are unable to rejoin, even when the network is repaired. In such a case, you must manually remerge the configurations.

Related Documentation

- Clustering Feature Overview
- Deploying an Active/Active Cluster
- Admin Console Procedures

General Cluster Maintenance

Managing Network Settings for Cluster Nodes

To modify the network settings for a cluster or each individual node in a cluster, click System > Network. You can make your changes on the Network Settings pages. After you create a cluster, these pages provide a drop-down list from which you can select the entire cluster or a specific node to modify. When you save changes on a Network page, the settings are saved for the specified cluster or cluster node. If you change network settings for an entire cluster, they propagate to every node in the cluster.

You can access a node-specific Network page by clicking System > Clustering > Cluster Status on the node's name in the Member Name column.

Upgrading Clustered Nodes

The Connect Secure offers the ability to easily upgrade every node in a cluster. You simply install a newer service package on one node and, once the installation completes and the node reboots, the node pushes the service package to all nodes in the cluster.

Upgrading the Cluster Service Package

Install a newer service package on one cluster node only. When the installation process completes and the cluster node reboots, it instructs the other nodes to upgrade.

Related Documentation

- Synchronizing the Cluster State

- Adding Multiple Cluster Nodes

Migrating Cluster Configurations to a Replacement Cluster

To migrate system and user configurations from a Connect Secure cluster (C1) to a replacement cluster (C2) using different Connect Secure devices:

1. Export the system, user and IVS configuration from C1's primary node (PN1).

   Note the following information:

   - Cluster name

   - Cluster password

   - Name of the node where the export was done (PN1)

   - Internal IP address of PN1

   - Internal network mask of PN1

   - Internal network gateway of PN1

   - Name of all other nodes in the C1 cluster, including their internal network IP address, network masks and gateways

2. Shut down all Connect Secure devices in cluster C1.

3. Power on one of the new servers (must be running software release 6.1R1 or later) that is part of cluster C2 and is on the same network to which PN1 was attached. This Pulse server device is called PN2 for the remainder of these steps.

4. When prompted, configure the internal network settings of PN2 to the same internal network settings of PN1 as noted in Step 2.

5. Install the new primary license on PN2.

6. From the admin GUI on PN2, select **System > Clustering> Create Cluster**. Create the cluster C2 using the same cluster name and cluster password that were in use at cluster C1. Node PN2 must also be assigned the same node name as PN1 (see Step 2).

7. Open the cluster status page and add the remaining nodes to the cluster configuration. Nodes being added must be assigned the same names that existed in original cluster C1. The internal network settings of the newly added nodes must also match the corresponding settings in the original cluster C1.

   NOTE: Do not join the newly added nodes to cluster C2 yet.

8. Import the data exported from PN1 (see Step 1) into PN2.

When importing the system configuration, select the option **Import everything (except Device Certificate(s))**.

9. Power on the remaining new Pulse Connect Secure devices assigned to cluster C2. Configure the bare minimal internal network settings needed to bring up the machine. The network settings must match what has already been configured on node PN2.

> **ⓘ** NOTE: Do not do make any other configuration changes on these machines as they will be lost when these machines join the cluster. Do not add licenses on these machines yet.

10. Join the Pulse Connect Secure in Step 9 to cluster C2 and wait for the cluster status to stabilize.

11. Install the CL licenses on the newly joined nodes.

Related Documentation

- Configuration File Administration Overview

Configuring the External VIP for An Active/Passive Cluster

To add an external VIP to an existing A/P cluster:

1. Create an A/P cluster with only the internal port configured.

2. Select **System > Clustering > Clustering Properties** and add the internal VIP.

3. Select **System > Network > External Port**.

4. From the Settings for menu, select "entire cluster".

5. Add the Netmask and Default Gateway but leave the external port disabled.

6. For each node, select **System > Network > External Port** and configure the external port IP address but leave the external port disabled.

7. Add the external cluster VIP.

8. Select **System > Network > External Port**, select "entire cluster" from the Settings for menu and enable the external port.

Related Documentation

- *Specifying Active/Passive, Active/Active, and Other Cluster Settings*

Monitoring Clusters

You can monitor clusters using the standard logging tools provided by the Pulse Connect Secure. In particular, you can use several cluster-specific SNMP traps to monitor events that occur on your cluster nodes, such as:

- External interface down

- Internal interface down

- Disabled node

- Changed virtual IP (VIP)

- Deleted cluster node (cluster stop)

> **ⓘ** NOTE: Generally, it is desirable to configure your SNMP traps on a cluster-wide basis, so that any given cluster node can send its generated traps to the right target. Setting up cluster-wide configuration for the traps is particularly important when you

also use a load balancer, because you may not know which node is responsible for a specific operation. In that case, the load balancer may independently determine which cluster node can manage an administrative session.

You can use SNMP traps that are included in the Pulse Secure Standard MIB to monitor these events. These traps include:

- **iveNetExternalInterfaceDownTrap**—Supplies type of event that brought down the external interface.

- **iveNetInternalInterfaceDownTrap**—Supplies type of event that brought down the internal interface.

- **iveClusterDisableNodeTrap**—Supplies the cluster name on which nodes have been disabled, along with a space separated list of disabled node names.

- **iveClusterChangedVIPTrap**—Supplies the type of the VIP, whether external or internal, and its value before and after the change.

- **iveClusterDelete**—Supplies the name of the cluster node on which the cluster delete event was initiated.

These traps are always enabled and available in the MIB. You cannot disable the traps.

Related Documentation

- Defining and Initializing a Cluster

- Synchronizing the Cluster State

- Troubleshooting Clusters

Troubleshooting Clusters

When you have problems with cluster communication, you may be directed by your Pulse Secure Support representative to use the cluster node troubleshooting tools.

To use the cluster node troubleshooting tools:

From the admin console, select **Maintenance > Troubleshooting > Monitoring > Node Monitor**, in **Maintenance > Troubleshooting > Clustering Network Connectivity**, and in **Maintenance > Troubleshooting > Clustering Group Communication**.

You can use a built-in feature on the clustering Status page to identify the status of each cluster node. Pause the mouse pointer over the Status light icon and the system displays a tool tip containing a hexadecimal number. The hexadecimal number is a snapshot of the status of the Pulse Connect Secure. It is a bit mask indicating a number of states as shown in Table 177.

Table178: Cluster Status

Value	Meaning
0x000001	Pulse Connect Secure is in standalone mode.
0x000002	Pulse Connect Secure is in cluster disabled state.
0x000004	Pulse Connect Secure is in cluster enabled state.
0x000008	Unable to communicate (because it is offline, has wrong password, has different cluster definition, different version, or a related problem).

0x00002000	The node owns the VIPs (on) or not (off).
0x000100	Pulse Connect Secure is syncing state from another Pulse Connect Secure (initial syncing phase).
0x000200	Pulse Connect Secure is transitioning from one state to another.
0x00020000	The group communication subsystems at the local and remote nodes are disconnected from each other.
0x00040000	Management interface (mgt0) appears disconnected.
0x00080000	Management gateway is unreachable for ARP ping.
0x000800	Pulse Connect Secure int0 appears disconnected (no carrier).
0x001000	This node is configured to be a cluster member.
0x002000	Pulse Connect Secure is syncing its state to another Pulse Connect Secure that is joining.
0x004000	Initial Synchronization as master or slave is taking place.
0x008000	This Pulse Connect Secure is the leader of the cluster.
0x010000	The group communication subsystem is functional.
0x020000	The gateway on int0 is unreachable for ARP pings (see log file).
0x040000	The gateway on int1 is unreachable for ARP pings (see log file).
0x080000	Leader election is taking place.
0x100000	Server life cycle process (dsmon) is busy.
0x200000	System performs post state synchronization activities.
0x30004	• The group communication subsystem is functional. • The gateway on int0 is unreachable for ARP pings (see log file). • Pulse Connect Secure is in cluster enabled state.
0x80000000	Cluster keystore or security world has not been associated with the FIPS card.

Each code, as you see it in the Pulse Connect Secure, may relate specifically to one state. However, each code may represent a combination of states, and so the actual code does not appear in Table 177. Instead, the code you see in the Pulse Connect Secure is the sum of several of the hexadecimal numbers shown in Table 177. You will need to factor out the codes, as in the following example:

- 0x3800**4**—The right-most digit (4) in this hexadecimal number corresponds to:

  - 0x000004 The Pulse Connect Secure is in cluster enabled state.

- 0x03**8**004—The digit in the fourth position from the right (8) corresponds to:

  - 0x008000 This Pulse Connect Secure is the leader of the cluster.

- 0x**3**8004—The left-most digit (3) in this hexadecimal number does not exist in the table, which indicates that it corresponds to the sum of two other digits, in this case, 1 and 2, as shown in the following codes:

  - 0x020000—The gateway on int0 is unreachable for ARP pings (see log file).

  - 0x010000—The group communication subsystem is functional.

"Management IP Address Differs From the Management IP Address" Error Message

If you receive the following error when joining a standalone SA6000/SA6500 node to a cluster even though the management port is configured and enabled:

Management IP address (x.x.x.x) for the local system differs from the Management IP address (not entered) configured for this system in the remote system.

then perform the following steps to add the node:

1. From the admin console of the primary node, select System > Network > Management Port.

2. Select the node to add from the drop down list next to the "Setting for" label.

3. Enable the management port and enter the IP address, netmask and default gateway for the joining node.

4. Click Save Changes.

5. From the admin console of the joining node, join the cluster again.

Fail-over Transactions

In the case of a fail-over (both in active/passive and active/active configurations), all transactions currently in progress (such as telnet or SSH sessions or large file downloads/uploads) must be restarted after the fail-over. There is no seamless fail-over for on-going transactions using sockets except for HTTP requests or non-stateful connections.

Related Documentation

- Monitoring Clusters

Using the Serial Console for Cluster Administration

If you are adding a factory-set device to a cluster, we recommend that you use the serial console, which enables you to join an existing cluster during the initialization process by entering minimal information. When a node joins a cluster, it receives the cluster state settings, which overwrite all settings on a device with an existing configuration and provide new machines with the required preliminary information. You can also use the serial console to disable the node. If the node is in a synchronization state, you cannot access its admin console. Therefore, if you need to upgrade or reboot the node, for example, you must first disable the node from a cluster through its serial console.

- Joining a Node to a Cluster Using Its Serial Console

- Disabling a Clustered Node Using Its Serial Console

- Restarting or Rebooting Cluster Nodes Using Its Serial Console

Joining a Node to a Cluster Using Its Serial Console

Before a configured or factory-set node can join a cluster, you must make its identity known to the cluster.

> **NOTE:**
>
> - To add a node currently running as a standalone device to a cluster through its admin console, it must be running the same or a more recent version service package on the same hardware platform as the other members.

- If you add a node running an earlier version service package to a cluster, the node automatically detects the mismatch, gets the newer package from the cluster, and joins the cluster.

To add a node to a cluster through its serial console:

1. In the admin console of an existing cluster member, select **System > Clustering > Cluster Status** and specify the node to add to the cluster.

2. Connect to the serial console of the device you want to add to the cluster.

3. Reboot the device and watch its serial console. After the system software starts, a message appears stating that the device is about to boot as a standalone node and to press the Tab key for clustering options. Press the Tab key as soon as you see this option.

    i NOTE: The interval to press the Tab key is five seconds. If the device begins to boot in standalone mode, wait for it to finish and then reboot again.

4. Enter the number instructing the node to join an existing cluster.

5. Enter the requested information, including:

    - The internal IP address of an active member in the cluster

    - The cluster password, which is the password you entered when defining the cluster

    - The name of the device to add

    - The internal IP address of the device to add

    - The netmask of the device to add

    - The gateway of the device to add

    The active cluster member verifies the cluster password and that the new device's name and IP address match what you specified in the admin console. If the credentials are valid, the active member copies all of its state data to the new cluster member, including certificate, user, and system data.

6. Enter the number instructing the node to continue the join cluster operation. When you see a message confirming that the device has joined the cluster, select **System > Clustering > Cluster Status** in the admin console of any active cluster member to confirm that the new member's Status is green, indicating that the node is now an enabled node of the cluster (status is green).

Disabling a Clustered Node Using Its Serial Console

To disable a node within a cluster using its serial console:

1. Connect to the serial console of the device you want to disable within the cluster.

2. Enter the number that corresponds to the System Operations option.

3. Enter the number that corresponds to the Disable Node option.

4. Enter **y** when the serial console prompts you to confirm that you want to disable the node.

5. Verify that the node has been disabled (status is red) within the cluster by selecting **System > Clustering > Status** in the admin console of any active cluster member.

Restarting or Rebooting Cluster Nodes Using Its Serial Console

When you create a cluster of two or more nodes, the clustered nodes act as a logical entity. When you reboot one of the nodes using either the serial console or the admin console, all nodes in the cluster restart or reboot.

To reboot only one node:

1. Connect to the serial console of the device you want to disable within the cluster.

2. Enter the number that corresponds to the System Operations option.

3. Select **System > Clustering > Status** to disable the node you want to restart or reboot within the cluster.

4. Under system operations select the appropriate menu option <Reboot this device>, <Shutdown this device>, or <Restart Services>.

5. Reboot the node, then enable the node within the cluster again.

   The system reconciles session state with the Infranet Enforcer upon restart or cluster failover. If the Infranet Enforcer is running ScreenOS 6.0r2 or later, an Policy Secure restart or failover does not interrupt network traffic of existing sessions, as long as the restart or failover occurs within two minutes.

Related Documentation

- Clustering Feature Overview

Monitoring Cluster Nodes

If you have a problem with a cluster, a Pulse Secure Support representative may ask you to create a snapshot that includes node monitoring statistics to assist with debugging the cluster problem. When you enable the node monitor on the Maintenance > Troubleshooting > Monitoring > Node Monitor tab, the Pulse Connect Secure captures certain statistics specific to the cluster nodes on your system. Using the snapshot that results, the support team can identify important data, such as network statistics and CPU usage statistics.

To enable node monitoring:

1. Enable the node monitor on the **Maintenance > Troubleshooting > Monitoring > Node Monitor** tab

2. Enter the maximum size for the node monitor log.

3. Enter the interval, in seconds, at which node statistics are to be captured.

4. Select the **Node monitoring enabled** check box to start monitoring cluster nodes.

5. For **Maximum node monitor log size**, enter the maximum size (in MB) of the log file. Valid values are 1-30.

6. Specify the interval (in seconds) that defines how often nodes are to be monitored.

7. Select the commands to use to monitor the node.

   If you select **dsstatdump**, enter its parameters as well.

8. Click **Save Changes**.

9. If you want to include the node monitoring results in the system snapshot, choose **Maintenance > Troubleshooting > System Snapshot**, and select the Include debug log check box.

10. Take a system snapshot to retrieve the results.

Related Documentation

-

-

-

Cluster Group Communication and Node Monitoring

-
-
-

Overview

If you have a problem with a cluster, a Pulse Secure Support representative might ask you to create a snapshot that includes group communication statistics to assist with debugging the cluster problem. When you enable the group communication monitor in the Group Communication tab, the system records statistics related to all of the cluster nodes on your system. As the local node communicates with other nodes in the cluster, the system captures statistics related to intra cluster communication. The Group Communication tab is displayed only when you enable clustering on your system. On a standalone system, you do not have access to the Group Communication tab.

You can also enable the cluster networking troubleshooting server on the Network Connectivity page.

> **i** NOTE:
>
> - Performing excessive node monitoring can impact system performance and stability. You should only perform extensive monitoring when directed by your Pulse Secure Support representative.
>
> - Performing log synchronization across cluster nodes can impact your system performance and stability.

Configuring Group Communication Monitoring on a Cluster

To enable group communication monitoring:

1. Enter the maximum size for the statistics log.

2. Enter the interval, in seconds, at which events are to be logged.

3. If you want to monitor all cluster nodes from the current local node, select the **Monitor all cluster nodes from this node** check box. If you do not check this option, the group communication monitor gathers statistics only for the local node.

   > **i** NOTE: If you select the Monitor all cluster nodes from this node option, the cluster nodes must be able to communicate over UDP port 6543.

4. Select the **Enable group communication monitoring** check box to start the monitoring tool.

5. Click **Save Changes**.

   Figure 272 shows the Troubleshooting page for group communication.

Figure280: Troubleshooting using Group Communication

6. If you want to include the node monitoring results in the system snapshot, choose **Maintenance > Troubleshooting > System Snapshot**, and select the Include debug log check box.

7. Take a system snapshot to retrieve the results.

Configuring Cluster Node Monitoring

If you have a problem with a cluster, a Pulse Secure Support representative may ask you to create a snapshot that includes node monitoring statistics to assist with debugging the cluster problem. When you enable the node monitor on the Node Monitor tab, the IC Series device captures certain statistics specific to the cluster nodes on your system. Using the resulting snapshot, the support team can identify important data, such as network statistics and CPU usage statistics.

To enable node monitoring:

1. Select **Maintenance > Troubleshooting > Monitoring > Node Monitor** to enable the node monitor.

2. Enter the maximum size for the node monitor log.

3. Enter the interval, (in seconds) at which node statistics are to be captured.

4. Select the **Node monitoring enabled** check box to start monitoring cluster nodes.

   Figure 273 shows the Troubleshooting page for node monitoring.

Figure281: Troubleshooting using Node Monitor

5. For **Maximum node monitor log size**, enter the maximum size (in MB) of the log file. Valid values in the range of 1 - 30.

6. Specify the interval (in seconds) that defines how often nodes are to be monitored.

7. Select the commands to use to monitor the node.

   If you select **dsstatdump**, enter its parameters as well.

8. Click **Save Changes**.

9. To include the node monitoring results in the system snapshot, select **Maintenance > Troubleshooting > System Snapshot**, and select the Include debug log check box.

10. Take a system snapshot to retrieve the results.

Cluster Network Connectivity

- Overview
- Configuring Cluster Network Connectivity Monitoring

Overview

If you have a problem with a cluster, a Pulse Secure Support representative might ask you to enable the cluster node troubleshooting server. When you enable the server on the Network Connectivity tab, the system attempts to establish connectivity between the node on which the server resides and another node you specify. As the nodes communicate, the system displays network connectivity statistics on the page. The Network Connectivity tab is displayed only when you enable clustering on your system. On a standalone system, you do not have access to the Network Connectivity tab.

Use the Network Connectivity tab to enable the cluster node troubleshooting server and to select a node on which to perform troubleshooting tasks. The troubleshooting tool allows you to determine the network connectivity between cluster nodes.

The server component of this tool runs on the node to which connectivity is being tested. The client component runs on the node from which connectivity is being tested. The basic scenario for testing connectivity is this:

- The administrator starts the server component on the passive node.

- The administrator tests the connectivity to the server node from the Active node, by starting the client component on the active node and then contacting the passive node running the server component.

  ⓘ NOTE: The server component must be run on nodes that are configured as either standalone or in a cluster but disabled. Cluster services cannot be running on the same node as the server component.

Configuring Cluster Network Connectivity Monitoring

To enable network connectivity monitoring:

1. Select the **Enable cluster network troubleshooting server** check box to enable the server component.

   Figure 274 shows the Troubleshooting page for network connectivity.

Figure282: Troubleshooting using Network Connectivity

2. Click **Save Changes**.

3. On another machine, select **Maintenance > Troubleshooting > Cluster > Network Connectivity**.

4. Perform one of the following steps:

   - Select a node from the list.

   - Enter the IP address of the server node.

5. Click **Go** to begin troubleshooting the machine on which the server component is running.

6. Click the **Details** link below the fields to view the results.

Example: Creating an Active/Active Cluster That Supports IPv6 Client Access

This example describes the tasks involved in creating a cluster that supports IPv6 client access. It includes the following information:

- Overview

Overview

Release 8.0 supports an IPv6 configuration for active/active clusters. The previous intracluster communication mechanism is preserved. The intracluster communication occurs over the IPv4 corporate network through the internal interfaces.

If you attempt to change the IP address of a node while it belongs to a cluster, you might experience unpredictable results. Whenever you change the IP address configuration for a cluster, you must re-create the cluster. Therefore, to add support for IPv6 addresses, you must re-create the cluster.

Before You Begin

We recommend that you deploy a cluster in a staging environment first and then move to a production environment after testing authentication realm, user role, and resource policy configurations, as well as any applications your end users might access.

Before you begin a cluster configuration:

1. Ensure that all intended Pulse Connect Secure nodes use the same hardware platform (for example, all are SA6500 Series Appliances).

2. Ensure that all intended Pulse Connect Secure nodes have been initially configured (for example, Pulse Connect Secure hostname is specified and the internal and external IP addresses are assigned), and they are running the same service package version.

3. Designate one node as the primary node. On the primary node, configure system and user settings. When other nodes join the cluster, the primary node propagates its configuration to the new cluster member during the join cluster operation.

Defining and Initializing a Cluster

You use the primary node admin GUI graphical user interface to create the cluster and add members. The primary node is added as part of the cluster creation operation. When you add members, you are prompted for settings unique to the member, such as the name and IP address configuration for the internal and external interfaces. A few additional settings are also unique, namely the management port and VLAN port settings, so you add these manually after the add node procedure that follows, but before the join cluster operation.

To create a cluster and add members:

1. Select **System > Clustering > Create** and enter a name for the cluster, a cluster password, and a name for this node, such as Node-1.

   You need to enter the password again when specifying additional nodes to join the cluster. All nodes in the cluster use this password to communicate.

2. Click **Create Cluster**. When prompted to confirm the cluster creation, click **Create**. After the Pulse Connect Secure initializes the cluster, the Clustering page displays the Status and Properties tabs.

3. Click **Add Members** to specify the additional cluster nodes:

   1. Enter a name for the member; for example, Node-2.

2. Enter the internal IP address. If both IPv4 and IPv6 are enabled on the internal port on Node-1, the system prompts for both IPv4 and IPv6 settings for the internal port for Node-2. Note, however, that intracluster communication uses the IPv4 corporate network.

3. Enter the external IP address. If both IPv4 and IPv6 are enabled on the external port on Node-1, the system prompts for both IPv4 and IPv6 settings for the external port for Node-2.

4. Change the netmask/prefix-length and gateway settings for the node if necessary.

5. Click **Add Node**. When prompted to confirm adding the new member, click **Add**.

   When the add node operation has completed, Node-2 is shown as an unreachable member of the cluster.

6. The add node procedure does not prompt you to configure management port or VLAN port settings. As needed, go to the node port configuration page and configure these settings. For example, after the add node operation has completed for Node-2, go to its System > Network > *Port* > Settings page and configure its management port.

7. Repeat this procedure for each node you intend to add to a cluster.

Joining Nodes to the Cluster

The primary node joins the cluster as part of the creation process. Use the following procedure to join additional nodes to the cluster.

To join a node to the cluster:

1. From an existing cluster member, select the **System > Clustering > Cluster Status** tab and specify the Pulse Connect Secure you want to add to the cluster.

2. From the admin GUI of the Pulse Connect Secure you want to join to a cluster:

   1. Select the **System > Clustering > Join** tab and enter:

      - The name of the cluster to join.

      - The cluster password you specified when defining the cluster.

      - The IPv4 address for the internal port of an active cluster member.

   2. Click **Join Cluster**. When prompted to confirm joining the cluster, click **Join**.

The join cluster operation validates IPv4 and IPv6 settings for all the physical ports (internal/external/management) against those present in the existing cluster. For example, the external port IPv6 settings present on Node-2 are compared against external port IPv6 settings that were specified for the Node-2 add member operation entered on the primary node (Node-1). If there is a mismatch, the join operation fails with an appropriate error message.

While the new node synchronizes its state with the existing cluster member, each node's status indicates **Enabled**, **Enabled, Transitioning**, or **Enabled, Unreachable**.

When the node finishes joining the cluster, its Clustering page shows the Status and Properties tabs.

After the node joins the cluster, you might need to sign in again.

Advanced Configuration

Table 178 summarizes advanced configuration guidelines.

Table179: Pulse Connect Secure Clusters: Advanced Configuration Guidelines

Topic	Guideline
Active/Active	When using Pulse clients with an active/active cluster, you must split the IP address pool across the nodes to ensure proper routing from the backend to the end user. This is a requirement whether the IP address pool is provisioned statically on the Pulse Connect Secure or dynamically by way of DHCP.
	The client IP pool configuration is synchronized among all nodes in a cluster; however, you may configure each node to use a certain subset of the global IP pool.
	If you are running Network Connect on a multisite cluster where nodes reside on different subnets:
	1. Configure an IP address pool policy on the Users > Resource Policies > VPN Tunnelling: Connection Profiles > New Profile page that accounts for the different network addresses used by each node in the cluster.
	2. For each node in the cluster, use settings in the System > Network > VPN Tunneling page of the admin GUI to specify an IP filter that filters out only those network addresses available to that node.
	NOTE: Release 8.0 does not support IP filters for IPv6 addresses.
	3. Create a static route on your gateway router that indicates the IP address of the internal port of each cluster node. Each IP address specified on the router needs to be in the same subnetwork as the corresponding cluster node.
FIPS	If you are creating a cluster of FIPS devices, manually update the security world on each of the nodes.

Related Documentation

- Clustering Feature Overview

- General Cluster Maintenance

Example: Creating an Active/Passive Cluster that Supports IPv6 Client Access

This example describes the tasks involved in creating a cluster that supports IPv6 client access. It includes the following information:

- Overview

- Before You Begin

- Defining and Initializing a Cluster

- Joining Nodes to the Cluster

- Configuring IPv6 on an Existing IPv4 Active/Passive Cluster

- Advanced Configuration

Overview

Pulse Secure access management framework supports an IPv6 configuration for active/passive clusters. The previous intracluster communication mechanism is preserved. The intracluster communication occurs over the IPv4 corporate network through the internal interfaces.

If a device belongs to an active/passive cluster, you can enable IPv6 on its ports. If a device has IPv6 enabled on its ports, it can be added to an active/passive cluster.

If you attempt to change the IP address of a node while it belongs to a cluster, you might experience unpredictable results. Whenever you change the IP address configuration for a cluster, you must re-create the cluster.

When using active/passive clustering, the members of a cluster pair must be in the same subnet because the VIP address must be shared by both members.

Before You Begin

We recommend that you deploy a cluster in a staging environment first and then move to a production environment after testing the authentication realm, user role, and resource policy configurations, as well as any applications your end users might access.

Before you begin a cluster configuration:

Note that state synchronization occurs only through the internal network interface card (NIC).

Ensure that all intended nodes use the same hardware platform (for example, all are SA6500 Series Appliances).

Ensure that all intended nodes have been initially configured (for example, the system hostname is specified and the internal and external IP addresses are assigned), and that they are running the same service package version.

Designate one node as the primary node. On the primary node, configure system and user settings. When other nodes join the cluster, the primary node propagates its configuration to the new cluster member during the join cluster operation.

Configuring IPv6 on an existing IPv4 active/passive cluster on an external port can be done seamlessly. However, If you are configuring on an internal port, you must wait for cluster synchronization completion and then do the next configuration for the remaining node. Therefore we recommended that you complete the IPv6 configurations before creating a cluster on an internal port.

Defining and Initializing a Cluster

You use the primary node admin GUI to create the cluster and add members. The primary node is added as part of the cluster creation operation. When you add members, you are prompted for settings unique to the member, such as the name and IP address configuration for the internal and external interfaces. A few additional settings are also unique, namely the management port and VLAN port settings, so you add these manually after the add node procedure that follows, but before the join cluster operation.

To create a cluster and add members:

1. Select **System > Clustering > Create Cluster** and enter a name for the cluster, a cluster password, and a name for this node, such as Node-*X*.

   You need to enter the password again when specifying additional nodes to join the cluster. All nodes in the cluster use this password to communicate.

   Figure 275 shows the Create New Cluster page.

Figure283: Create New Cluster Page

2. Click **Create Cluster**. When prompted to confirm the cluster creation, click **Create**. After Connect Secure initializes the cluster, the Clustering page displays the Status and Properties tabs.

shows the Clustering page with Status and Properties tabs.

Figure284: Clustering Page– Status and Properties

3. Click **Properties**.

shows the Clustering page with active/passive configuration.

Figure285: Clustering Page– Active/Passive Configuration

4. Under Configuration Settings, select **Active/Passive Configuration**, then specify the IPv4 and IPv6 addresses for the VIP address on the internal and external ports, depending on what is enabled for IPv4/IPv6 at **Network > Internal Port** and **Network > External Port**.

5. Click **Save Changes**. After the system initializes the active/passive cluster, the Clustering page displays the Status and Properties tabs.

6. Click **Add Members** to specify additional cluster nodes:

Figure286: Add Cluster Member Page

1. Enter a name for the member; for example, Node-*Y*.

2. Enter the internal IP address. If both IPv4 and IPv6 are enabled on the internal port on Node-*X*, the system prompts for both IPv4 and IPv6 settings for the internal port for Node-*X*. Note, however, that intracluster communication uses the IPv4 corporate network.

3. Enter the external IP address. If both IPv4 and IPv6 are enabled on the external port on Node-*X*, the system prompts for both IPv4 and IPv6 settings for the external port for Node-*Y*.

4. (Optional) Change the netmask, prefix-length, and gateway settings for the node if necessary.

5. Click **Add Node**. When prompted to confirm adding the new member, click **Add** and then click **Save Changes**.

   After the completion of add node operation, Node-*Y* is shown as an unreachable member of the cluster.

6. Verify the configuration on **System > Clustering > Cluster Status** page.

   Figure 279 shows the status on the Clustering page.

Figure287: Clustering Page –Status

The add node procedure does not prompt you to configure management port or VLAN port settings. As needed, go to the node port configuration page and configure these settings. For example, after the add node operation has completed for Node-*Y*, go to its System > Network > Port > Settings page and configure its management port.

NOTE: Only two nodes can be present in an active/passive cluster.

Joining Nodes to the Cluster

The primary node joins the cluster as part of the creation process.

To join additional nodes to the cluster:

1. From an existing cluster member, select the **System > Clustering > Cluster Status** tab and specify the Connect Secure you want to add to the cluster.

2. From the admin GUI of the Pulse Secure access management framework that you want to join to a cluster:

   1. Select the **System > Clustering > Join** tab and enter:

      - The name of the cluster to join

      - The cluster password you specified when defining the cluster

      - The IPv4 address for the internal port of an active cluster member

   2. Click **Join Cluster**. When prompted to confirm joining the cluster, click **Join**.

The join cluster operation validates IPv4 and IPv6 settings for all the physical ports (internal, external, and management) against those present in the existing cluster. For example, the external port IPv6 settings present on Node-*Y* are compared against external port IPv6 settings that were specified for the Node-*Y* add member operation entered on the primary node (Node-*X*). If there is a mismatch, the join operation fails with an appropriate error message.

While the new node synchronizes its state with the existing cluster member, each node's status indicates Enabled, Enabled, Transitioning, or Enabled, Unreachable.

When the node finishes joining the cluster, its Clustering page shows the Status and Properties tabs.

After the node joins the cluster, you might need to sign in again.

Configuring IPv6 on an Existing IPv4 Active/Passive Cluster

We recommend as a best practice that you configure IPv6 host and network settings on individual nodes before you create a cluster. In some cases, such as routine upgrade, you have already created a cluster configuration and only want to add IPv6 addresses to the existing interface configuration. If so, follow the procedures in this section precisely.

> **NOTE:** You must leave IPv6 disabled until the last step of the procedures shown below.

To modify the internal port configuration for the cluster:

1. Select **System > Network > Internal Port > Settings**.

2. Under Settings for, select **Entire cluster**.

3. Complete the configuration for the IPv6 prefix and the IPv6 gateway, but do not enable IPv6.

4. Verify that all the nodes are up and running, are in sync, and are in reachable state. Complete synchronization of the cluster pair might take a few minutes.

5. Under Settings for, select Node 1.

6. Configure the IPv6 address, but do not enable IPv6.

7. Verify both the nodes are up and running and in reachable state.

8. Repeat steps 6-8 for Node 2.

9. Select **System > Network > Internal Port > Virtual Ports**.

10. Update the cluster virtual port configuration to add the IPv6 address.

11. Select **System > Network > Internal Port > Settings**.

12. Under Settings for, select **Entire cluster**.

13. Select **Enable IPv6**.

To modify the external port configuration for the cluster:

1. Select **System > Network > External Port > Settings**.

2. Under Settings for, select **Entire cluster**.

3. Complete the configuration for the IPv6 prefix and the IPv6 gateway, but do not enable IPv6.

4. Verify that all the nodes are up and running, are in sync, and are in reachable state. Complete synchronization of the cluster pair might take a few minutes.

5. Under Settings for, select Node 1.

6. Configure the IPv6 address, but do not enable IPv6.

7. Verify both the nodes are up and running and in reachable state.

8. Repeat steps 6-8 for Node 2.

9. Select **System > Network > External Port > Virtual Ports**.

10. Update the cluster virtual port configuration to add the IPv6 address.

11. Select **System > Network > External Port > Settings**.

12. Under Settings for, select **Entire cluster**.

13. Select **Enable IPv6**.

Advanced Configuration

Table 179 summarizes advanced configuration guidelines.

Table180: Connect Secure Clusters: Advanced Configuration Guidelines

Settings	Guideline
FIPS	If you are creating a cluster of FIPS devices, manually update the security word on each of the nodes.

Related Documentation

- Clustering Feature Overview

- General Cluster Maintenance

CHAPTER 40 Delegating Administrator Roles

About Delegating Administrator Roles

The access management system enables you to delegate various management tasks to different administrators through system administrator roles and security administrator roles. System and security administrator roles are defined entities that specify management functions and session properties for administrators who are mapped to those roles. You can customize an administrator role by selecting the feature sets, user roles, authentication realms, resource policies, and resource profiles that members of the administrator role are allowed to view and manage. Note that system administrators may only manage user roles, realms, and resource policies; only security administrators can manage administrator components.

For example, you can create a system administrator role called "Help Desk Administrators" and assign users to this role who are responsible for fielding tier 1 support calls, such as helping users understand why they cannot access a Web application or system page. In order to help with troubleshooting, you may configure settings for the "Help Desk Administrators" role as follows:

- Allow the help desk administrators Write access to the System > Log/Monitoring page so they can view and filter the system logs, tracking down critical events in individual users' session histories, as well as the Maintenance > Troubleshooting page so they can trace problems on individual users' systems.

- Allow the help desk administrators Read access to the Users > User Roles pages so they can understand which bookmarks, shares, and applications are available to individual users' roles, as well as the Resource Policy or Resource Profile pages so they can view the policies that may be denying individual users access to their bookmarks, shares, and applications.

- Deny the help desk administrators any access to the remaining System pages and Maintenance pages, which are primarily used for configuring system-wide settings—such as installing licenses and service packages—not for troubleshooting individual users' problems.

  i   NOTE: In addition to any delegated administrator roles that you may create, the system also includes two basic types of administrators: super administrators (.Administrators role), who can perform any administration task through the admin console and read-only administrators (.Read-only Administrators role), who can view—but not change—the entire system configuration through the admin console.

You can also create a security administrator role called "Help Desk Manager" and assign users to this role who are responsible for managing the Help Desk Administrators. You might configure settings for the "Help Desk Manager" role to allow the Help Desk Manager to create and delete administrator roles on his own. The Help Desk Manager might create administrator roles that segment responsibilities by functional areas of the system. For example, one

administrator role might be responsible for all log monitoring issues. Another might be responsible for all Network Connect problems.

All devices allow members of the .Administrators role to configure general role settings, access management options, and session options for the .Administrators and .Read-Only Administrators roles.

> NOTE: On certain pages, such as the role mapping page, the delegated administrator can view the role names even though the administrator does not have read/write access. However, the delegated administrator cannot view the details of that role.

Related Documentation

- Creating and Configuring Administrator Roles

- Specifying Management Tasks to Delegate

Creating and Configuring Administrator Roles

You can use the Administrators > Admin Roles pages to set default session and user interface options for delegated administrator roles.

To create individual administrator accounts, you must add the users through the appropriate authentication server (not the role). For example, to create an individual administrator account, you may use settings in the Authentication > Auth. Servers > Administrators > Users page of the admin console. For detailed instructions on how to create users on the Administrators server and other local authentication servers. For instructions on how to create users on third-party servers, see the documentation that comes with that product.

To create an administrator role:

1. In the admin console, choose **Administrators > Admin Roles**.

2. Do one of the following:

   - Click **New Role** to create a new administrator role with the default settings.

   - Select the check box next to an existing administrator role and click **Duplicate** to copy the role and its custom permissions. Note that you cannot duplicate the system default roles (.Administrators and .Read-Only Administrators).

3. Enter a name (required) and description (optional) for the new role and click **Save Changes**.

4. Modify restrictions, session options, and UI options according to your requirements.

> NOTE: If you select one of the system's default administrator roles (.Administrators or .Read-Only Administrators), you can only modify settings in the General tab (since the default system administrators roles always have access to the functions defined through the System, Users, Administrators, and Resource Policies tabs).
>
> You cannot delete the .Administrators and .Read Only Administrators roles since they are default roles.

Related Documentation

- Specifying Management Tasks to Delegate

- Role Restrictions

- Specifying Role Session Options

- Customizing the Welcome Page

Specifying Management Tasks to Delegate

This topic contains information about delegating management tasks to various delegated administrator roles.

Delegating System Management Tasks

Use the **Administrators > Admin Roles >** *Select Role* **> System** tab to delegate various system management tasks to different administrator roles. When delegating privileges, note that:

- The system allows all administrators read-access (at minimum) to the admin console home page (System > Status > Overview), regardless of the privilege level you choose.

- The system does not allow delegated administrators write-access to pages where they can change their own privileges. Only those administrator roles that come with the system (.Administrators and.Read-Only Administrators) may access these pages:

  - Maintenance > Import/Export (Within this page,.Read-Only Administrators can export settings, but cannot import them.)

  - Maintenance > Push Config

  - Maintenance > Archiving > Local Backups

- Delegation access to the Meeting Schedule page is controlled through the Meetings option on the Administrators > Admin Roles > Select Role > Resource Policies page.

Delegating User and Role Management

Use the **Administrators > Admin Roles >** *Select Role* **> Users > Roles** sub-tab to specify which user roles the administrator role can manage. When delegating role management privileges, note that:

- Delegated administrators can only manage user roles.

- Delegated administrators cannot create new user roles, copy existing roles, or delete existing roles.

- If you allow the delegated administrator to read or write to any feature within a user role, the system also grants the delegated administrator read access to the Users > User Roles > Select Role > General > Overview page for that role.

- If you grant a delegated administrator write access to a resource policy through the Administrators > Admin Roles > Select Administrator Role > Resource Policies page, he may create a resource policy that applies to any user role, even if you do not grant him read access to the role.

Delegating User Realm Management

Use the **Administrators > Admin Roles >** *Select Role* **> Users > Authentication Realms** tab to specify which user authentication realms the administrator role can manage. When delegating realm management privileges, note that:

- System administrators can only manage user realms.

- System administrators cannot create new user realms, copy existing realms, or delete existing realms.

- If you allow the system administrator to read or write to any user realm page, the system also grants the system administrator read-access to the Users > User Realms > Select Realm > General page for that role.

Delegating Administrative Management

Use the **Administrators > Admin Roles >** *Select Roles* **> Administrators** tab to specify which system administrator roles and realms the security administrator role can manage. When delegating security administrative privileges, note that:

- The security administrator role provides control over all administrative roles and realms.

- You can give a security administrator control exclusively over administrator roles, over administrator realms, or over both.

- You can restrict or grant the security administrator the permission to add and delete administrator roles and administrator realms.

Delegating Resource Policy Management

Use the **Administrators > Admin Roles > Resource Policies** tab to specify which user resource policies the administrator role can manage. When delegating resource policy management privileges, note that delegated system administrators cannot modify the following characteristics of resource policies:

- The resource itself (that is, the IP address or hostname).

- The order to evaluate the resource policies.

Delegating Resource Profile Management

Use the Administrators > Admin Roles > Resource Profiles tab to specify which user resource profiles the administrator role can manage. When delegating resource profile management privileges, note that delegated system administrators cannot modify the following characteristics of resource profiles:

- The resource itself (that is, the IP address or hostname)

- The order to evaluate the resource policies.

Delegating Access to IVS Systems

If you are running an IVS license, you can also delegate administrative access and responsibilities to specific IVS systems. You can delegate read/write access or read-only access to all IVS systems, or to selected IVS systems..

Related Documentation

- About Delegating Administrator Roles

CHAPTER 41 Instant Virtual System

Instant Virtual System (IVS) Overview

The Instant Virtual System (IVS) gives managed service providers (MSPs) the opportunity to offer cost-effective secure remote access, disaster recovery and managed extranet services to small and medium sized companies. To meet this opportunity, MSPs can deliver managed security solutions from equipment that is located on the subscriber company's premises (Customer Premises Edge router-based) or within the MSP network (Carrier Edge router-based or network-based). Network-based managed security solutions centralize the security gateway equipment in the MSP network. A virtualized system allows the MSP to provide managed, network-based SSL VPN services to multiple customers from the same equipment. The basic business model might work something like this:

- The MSP manages the SSL VPN equipment at the MSP site.

- Small and medium-sized companies subscribe to monthly services from the MSP.

- The MSP is responsible for the management of the equipment, but delegates portal administration to an IVS administrator designated by each subscriber company.

- The virtual system supports and enforces an architectural and administrative separation between subscriber companies, providing a completely secure and individualized view for each subscriber.

This system provides a number of benefits to service providers:

- Expand market share—The ability to provide secure SSL VPN capabilities to many subscriber companies from one system offers the MSP economies of scale and the opportunity to expand market share with services targeting small and medium sized businesses.

- Simplify administration—Each subscriber administrator can manage their company's IVS instance with no visibility into another subscriber company's administrative interface. The MSP root administrator can manage all hosted companies and can easily monitor or configure hosted company systems.

- Enhance subscriber security—Each subscriber company maintains complete separation from other subscriber companies. As far as the subscriber administrator or subscriber users are concerned, they are operating on a completely independent and protected SSL VPN system.

- Optimize traffic management—Traffic from end users or corporate intranet servers stays within each company's VLAN. Subscriber end users never see services located on another subscriber's intranet.

The standalone client installers are not accessible directly from the admin UI of an IVS. As a workaround, the root administrator can make the following link available to IVS administrators if the IVS administrators need to download standalone installers:

https://*myive*/dana-admin/sysinfo/installers.cgi

where *myive* is the hostname of your Connect Secure device.

> **NOTE:** IVS does not support connectivity from any version of Pulse Secure client (mobile or non-mobile). IVS is not supported on the MAG Series Pulse Secure Gateways or on virtual appliances.

You must have an IVS license to create IVS systems.

Suppose your device has no IVS license. When you create VLANs, they are available for the entire system. However, when an IVS license is applied, the VLANs are not applied to the Root IVS. They must manually be assigned to the Root IVS or other IVSes that you create.

Related Documentation

- Deploying an IVS
- IVS Provisioning Process Overview
- Creating a Virtual System (IVS Profile)
- IVS Use Case: Policy Routing Rules Resolution

Deploying an IVS

For each subscriber company, the virtualized system provides a secure portal for the company's end users (mobile employees, customers, or partners) to access its internal resources. Authentication servers that reside either on the subscriber's premises or in the MSP network, authenticate end users who sign in to the IVS. Once authenticated, end users establish secure sessions through the IVS to their respective company's back-end servers.

Figure288: MSP Deployment Scenario

The following numbered list items correspond to the labeled objects in Figure 280.

1. End users sign in to different subscriber company intranets on specified IP addresses.

2. End users sign-in over an Internet connection using a standard SSL-enabled Web browser.

3. All traffic is directed into the Managed Service Provider's (MSP) network. The MSP is the customer who holds the license to the virtualized hardware and software.

4. All traffic is directed to the virtualized system. Each message is evaluated based on its sign-in IP address and, by the virtualized system, is assigned a VLAN tag containing a VLAN ID that corresponds to a subscriber company. The system supports up to 240 IVS systems, each one representing a single subscriber company device. The subscriber is any company that subscribes to hosted SSL VPN services from the MSP.

The number of VLANs supported depends on the number of IVS systems. The number of IVS systems plus the number of VLANs must be less than or equal to 240.

Although you can create up to 240 IVS systems, it is strongly recommended that you do not create and use more than 16 IVS systems without first consulting the Performance Metrics and Testing application note. Contact Pulse Secure Technical Support for more information on this application note.

5. The MSP carrier-edge (CE) router or other Layer 2 device acts as a VLAN termination point, and routes traffic over a secure tunnel to a customer premises edge (CPE) router. Based on the VLAN ID, the router directs the traffic to the appropriate subscriber intranet. During this part of the process, the CE router removes the VLAN tag containing the VLAN ID, as once the message is correctly destined for the appropriate intranet, the ID and tag are no longer needed. The term subscriber intranet is interchangeable with the term company intranet.

6. The CE router routes messages over the service provider backbone to the appropriate customers' premises edge routers through encrypted tunnels, such as IPsec, GRE, PPP, and MPLS tunnels. Untagged traffic is sent over these tunnels to the customer intranet.

7. The CPE routers within the customer intranet on the customer premises can act as a VLAN termination point and routes traffic from the secure tunnel connected to the CE Router, to the customer intranet.

8. The end-user traffic reaches the correct subscriber company's backend resources. The system processes any return messages to the end users from the subscriber intranets following a similar set of steps.

In a typical MSP deployment, firewalls are present in front of the device in the MSP's DMZ, behind the device, in the MSP network or in the customer's intranet DMZ, or both. Note that a virtualized firewall could potentially exist behind the device (a Vsys cluster, for example), in which case it should have the ability to accept VLAN tagged traffic from the device and forward it to the proper customer VLAN (and vice versa). Also, most, if not all deployments have Domain Name Server (DNS) or Application servers located either in the MSP network or on the customer intranet.

In a virtualized deployment, the front-end is considered the external interface and is the end user or Internet-facing interface. The back-end is considered the internal interface and is the subscriber company intranet-facing interface.

The system tags inbound traffic sent by end users and destined for a server in the subscriber intranet or MSP network, with VLAN tags containing the VLAN ID. Inbound traffic can arrive over the system's internal interface or external interface.

Outbound traffic, which is traffic transmitted over the system backend and destined for servers located on MSP network or subscriber intranet, can be sourced by Connect Secure itself. For example, traffic destined for authentication, DNS, or application servers, is outbound traffic, as is traffic forwarded by Connect Secure, such as Network Connect traffic.

If the traffic arrives as inbound traffic to an IP address that has been designated for an IVS system that uses a VLAN, that traffic is tagged with the VLAN tag on arrival. When it has been identified and directed to the proper backend destination, the VLAN termination device strips the VLAN tag from the Ethernet frame and forwards the traffic to the backend destination.

Related Documentation

- Signing In Directly to the IVS as an IVS Administrator

- Signing In to the Root System or the IVS

- IVS Configuration Worksheet

Virtualized Connect Secure Architecture

The virtualized framework consists of a root system and any subscriber IVS systems the MSP root administrator creates subsequently. Subscriber IVS administrators can only manage resources on their particular IVS system. The root administrator can manage resources on all IVS systems on the appliance.

The IVS license converts the system to a root system that is functionally identical, with the added capability of provisioning virtual systems. The root system consists of system-level global data and a single default root IVS, which encompasses the access management subsystem. See Figure 281.

Figure289: IVS Architecture

The root administrator (root administrator) is the super-administrator of the root system. Often, the root administrator is the same thing as the system administrator. The root administrator has administrative control over the root system and all subscriber IVS systems. The root administrator can provision IVS systems on the root system, create IVS administrators, edit IVS configuration. The root administrator can override configuration changes made by any IVS administrator.

> ℹ️ NOTE: The instructions for configuring the root and IVS systems are meant to be read by a root administrator. The pronoun you, in these sections, denotes the root administrator. If a task can be performed by someone in a role other than the root administrator, the text makes a distinct reference to the role in the task description.

As shown in Figure 281:

1. The system administrator applies an IVS license to the system.

2. The resulting system contains the root global data and a root IVS, in effect, a virtualized device.

3. From the root IVS, the root administrator can create multiple subscriber IVS systems, each IVS completely separate from any other IVS.

The root system contains a superset of all capabilities of the system. You, as the root administrator, define all global network settings and root administrator settings on the root system. For each subscriber, you provision one or more IVS systems and manage them from the root system.

The subscriber IVS contains a unique instance of the access management framework. When you create an IVS for each subscriber company, you also create an IVS administrator (IVS administrator) account. The IVS administrator has complete administrative control over the IVS.

The IVS administrator uses an administrative admin console that contains a subset of the root administrator capabilities.

Related Documentation

- Signing In to the Root System or the IVS

- Deploying an IVS

Signing In to the Root System or the IVS

You can configure sign-in URLs using different methods:

- Sign-in URL prefix per IVS

- Virtual ports

- VLAN ports

You can use all of these methods on the same IVS.

Signing-In Using the Sign-In URL Prefix

This feature enables end users to access an IVS by way of a single hostname and IVS-specific sign-in URL prefix. By using this method, administrators can ensure that users can access multiple IVS systems by way of a single IP address.

Additionally, the use of path-based URLs results in:

- Savings in certificate costs—You need only supply one device certificate.

- Fewer DNS entries—You need only one DNS entry across all IVS systems hosted on a single device.

Administrators and end users can sign into an IVS system using sign-in URLs similar to the following (assuming the managed service provider URL is www.msp.com):

- Company A sign-in URL: www.msp.com/companyA

- Company B sign-in URL: www.msp.com/companyB

- Company A IVS administrator sign-in URL: www.msp.com/companyA/admin

- Company B IVS administrator sign-in URL: www.msp.com/companyB/admin

You can continue to restrict access by implementing additional sign-in URLs that are segregated by certain criteria, as follows:

- www.msp.com/companyA/sales

- www.msp.com/companyA/finance

- www.msp.com/companyA/hr

If you do not specify a URL prefix, the system defaults to sign-in over virtual ports. If you do specify a path-based sign-in URL prefix, the following rules apply:

- You cannot specify a multilevel path for the URL prefix, by using the / character.

- End users and administrators can sign in to an IVS on the internal port, external port, VLAN interface, or virtual port that has not already been assigned to an IVS using the selected URL prefix, in other words, where the hostname is the DNS name assigned to one of the interface IP addresses.

For example, assume that your ports are assigned to specific DNS names, as follows:

- Internal Port = MSP-internal

- External Port = MSP-external

- VLAN Port 10 = MSP-vlan10

- Virtual Port X = MSP-virtualx

Now, consider that VLAN Port 10 and Virtual Port X are not assigned to an IVS. If you host the Company A IVS, and the Company A sign-in URL prefix is specified as companyA in the IVS profile, then end users can sign-in to the Company A IVS using any of the following URLs:

- MSP-internal/companyA

- MSP-external/companyA

- MSP-vlan10/companyA

- MSP-virtualx/companyA

The path-based URL feature carries a few restrictions, as follows:

- An end user or administrator can sign into only one IVS from a given browser instance. If you attempt to sign in to another IVS from a new browser window of the same browser instance, your sign in attempt is rejected. You must create a new browser instance to sign in to multiple IVS systems.

- You cannot establish multiple concurrent sessions, with all sessions using Host Checker, from the same endpoint to different devices. You cannot establish multiple concurrent sessions from the same endpoint to multiple IVS systems, regardless of the sign-in method.

- If you configure an IVS with a path-based sign-in URL prefix, you cannot use the persistent session cookie (DSID) and maintain the ability to sign in to multiple IVS systems from the same browser using the URL prefix. The limitation does not apply to users signing in to the IVS with a sign-in IP address, because the system creates a different DSID per target IVS in that case.

- Pass-through proxy based on port numbers is supported. However, you cannot specify a pass-through proxy policy when using virtual hosts, unless the virtual host DNS entry maps to the IVS sign-in IP address. If the virtual host DNS entry points to the device, when the user signs in he will sign-in to the root IVS sign-in page.

- When using Pulse Collaboration, if a user is not already signed in to their IVS and you have enabled the option to require system users, all meeting invitation e-mails will contain a link to the root IVS sign-in page.

- If an IVS user bookmarks pages while using web rewriting, signs out, then reopens the browser and selects the bookmark, he will display the root IVS sign-in page.

Signing-In Over Virtual Ports

You may have reasons for configuring virtual ports for sign in. Virtual ports provide significant segregation of traffic. If you choose to use virtual ports, keep in mind that:

- Must provide multiple certificates—You need to supply one device certificate per virtual port address.

- Must configure multiple DNS entries—You need to supply DNS entries for each IVS system hosted on a single device.

- If multiple IVS's share the same virtual port for sign-in, the security settings of either IVS (allowed SSL/TLS version or encryption strength) can be associated to the virtual port. Once an IVS's security settings are associated to the virtual port, these settings are effective for

sign-ins (SSL sessions) to that virtual port. To guarantee deterministic selection/association of security settings to virtual ports, IVSs that share the same virtual port for sign-in should have the same security settings.

- If an IVS is configured to have different security settings relative to the root system, the security settings of the root take effect for sign-in to the IVS via sign-in URL prefix, whereas the security settings of the IVS take effect for sign-in to the IVS via sign-in over virtual ports. For an IVS with strong encryption (such as AES) on the device and with a root system having weaker encryption settings the following situation occurs: Signing in to the IVS via sign-in over virtual ports from an IE6 browser fails because the effective security settings for the SSL session is AES, which is not supported by IE6. However, the same user signing in to the same IVS from the same IE 6 browser through the sign-in URL prefix can succeed if the browser security settings are compatible with the root system's weaker security settings.

The sign-in request's target IP address drives the sign-in to the root system or IVS. To sign in to the root system or an IVS, users browse to a hostname-based URL. You map the URL, by way of external DNS, to the IP address or to an IP alias of the device's external interface.

For example, consider an MSP with hostname msp.com, that provides SSL VPN gateway services to two subscribers: s1 and s2.

- Root administrator sign-in URL: http://www.msp.com

- S1 sign-in URL: http://www.s1.com

- S2 sign-in URL: http://www.s2.com

External DNS must map these URLs to unique IP addresses, which must correspond to IP addresses or aliases hosted on the device, typically a virtual port defined on either the internal or external port.

To summarize signing-in, IVS users can sign in on:

- A virtual port configured on the external interface of the device.

- A virtual port configured on the internal interface of the device (untagged).

- A VLAN interface configured on the internal interface (tagged).

Root system users can also sign in directly over the internal or external interface.

Signing-In Over a VLAN Interface

In addition to the sign-in capabilities provided over the external interface (or the internal interface, if configured) by the root administrator, end users can sign in over any VLAN interface the root administrator assigns to their IVS. In other words, the IVS administrator can provide the VLAN port IP address to end users for sign-in.

You cannot map an explicit device certificate to any IP addresses mapped to a VLAN. When signing in over a VLAN interface, the system chooses the device certificate that is already assigned to the IVS. If there is no certificate associated with the IVS, the system assigns the certificate from the top of the device certificate list. This list can be reordered when a certificate is added or removed, which can result in an unpredictable certificate during configuration. Once an IVS is in a production state, this should not present a problem, as the IVS VIP is mapped to a specific certificate.

Related Documentation

- Navigating to the IVS

Navigating to the IVS

Only root administrators can navigate to an IVS from the root system. On the virtualized system, the admin console navigation for the root system includes an additional drop-down menu listing the configured IVS systems, on all page headers. You can navigate to an IVS and administer it by selecting an IVS from the drop-down menu. IVS administrators must sign-in directly to the IVS through a standard administrative sign-in page.

The root administrator creates the initial IVS administrator account. An IVS administrator can create additional IVS administrator accounts, using the standard procedure for creating administrator accounts.

Related Documentation

- Signing In to the Root System or the IVS

IVS Configuration Worksheet

To configure the system to properly steer inbound traffic to the correct subscriber IVS, and outbound traffic to the correct VLAN, the MSP root administrator needs to compile a profile for each subscriber company.

When creating a new virtual system, you must create a number of other system objects, and specify several pieces of data, including IP addresses, VLAN IDs, virtual ports, and DNS settings. You can use this worksheet to plan and keep track of the system data while creating each IVS. The worksheet presents data in the general order in which you should define the IVS.

Depending on the specific topology of the subscribers' networks, you may need to collect additional information, or may not use all of the information listed on the form.

Date:	Created By:

Subscriber:

Account #:

Comment:

Subscriber VLAN (System > Network > VLANs)

VLAN Settings

    VLAN Port Name:

    VLAN ID (1-4094):

VLAN Port Information

    IP Address:

    Netmask:

    Default Gateway:

Subscriber Sign-in Virtual Port Configuration (System > Network > Port 1 > Virtual Ports > New Virtual Port)

    External Virtual Port Name (for sign-in):

    IP Address:

    Internal Virtual Port Name (optional):

IP Address:	

Install Device Certificate for IVS hostname

IVS Hostname:	
Internal Port:	
External Port:	

Subscriber IVS (System > Virtual Systems > New Virtual System)

Name (Subscriber):	
Description:	
IVS Hostname:	

Administrator

Username:	
Password (at least 6 characters in length):	

Properties

Max Concurrent Users:	
Default VLAN:	
Selected Virtual Ports: (Internal Interface)	
Selected Virtual Ports: (External Interface)	
Network Connect IP Pool:	

Static Routes (System > Network > Routes > New Routes)

Destination Network/IP:	
Netmask:	
Gateway:	
Interface:	
Metric:	

DNS Settings (Subscriber IVS > System > Network > Overview)

Hostname:	
Primary DNS:	
Secondary DNS:	
DNS Domain(s):	
WINS:	

Related Documentation

- Deploying an IVS

Administering the Root System

Once you apply the IVS license to a system, the new Virtual Systems tab appears in the administrator UI. After you apply the IVS license, you can see an explicit display of the root system in the drop-down menu that appears in the admin console header area.

Setting up the system requires a series of basic procedures. Once the hardware is connected:

1. Boot the system.

2. Apply the IVS license through the Maintenance > System > Upgrade/Downgrade page of the admin console.

3. Configure the root system from the admin console.

Regardless of how many subscriber administrators you define on the subscriber IVS systems, you always maintain control over the entire system and have visibility into the settings on all IVS systems.

Related Documentation

- Signing In to the Root System or the IVS

Configuring the Root Administrator

Configuring the root administrator is similar to the task of creating a new administrator on a standalone device. You can create an administrator account through the Authentication > Auth. Servers > Administrators > Users page of the admin console, or by using the serial console.

If you upgrade from an older system software version to 5.1 version or later, the system considers any administrator in the root system who maps to the Administrators role to be a root administrator for the device. If you reimage the device or install a brand new piece of hardware, you create a primary administrator during the initial configuration steps, in the serial console.

Related Documentation

- Creating and Configuring Administrator Roles
- Signing In Directly to the IVS as an IVS Administrator

IVS Provisioning Process Overview

Figure 282 illustrates the basic tasks required to provision an IVS.

Figure290: Basic Steps to Provision an IVS

Provisioning an IVS consists of the following steps:

1. Configure one or more clusters, if needed, through the System > Clustering > Create Cluster page.

2. Configure and enable external port. The external port is in a disabled state, by default. You must enable the port and configure it, to provide sign-in capabilities from outside the network.

3. Create at least one VLAN port for each subscriber company. You must define a unique ID for each VLAN. A subscriber company can have multiple VLANs on a device.

4. Configure at least one virtual port on the external port to enable end users to sign in. You can also configure virtual ports on the internal port, for signing in from behind the firewall, if needed.

5. Load one certificate server per subscriber company.

6. If you intend to use virtual ports, for example, to support IP sourcing, create them at this point in the process.

7. Create an IVS profile for each subscriber company. The IVS profile establishes the connection between the company, the VLAN, and the available virtual ports.

8. Configure static routes to backend servers. If you intend to provide shared access to resources on the MSP network, you add static routes to the VLAN route tables that point to those resources.

9. Configure DNS settings, so that any traffic destined for resources on the MSP network first goes through the MSP's DNS server.

10. Log in as the IVS administrator.

11. Configure users, roles, realms, and resource policies for the IVS.

When you create the IVS, the IVS name appears in the drop-down menu located in the header of the admin console. You can perform operations on each IVS by selecting the IVS name in the drop-down menu and clicking the Go button.

Related Documentation

- Deploying an IVS

- Signing In to the Root System or the IVS

Configuring Sign-In Ports for IVS

You must configure virtual ports by which end users can sign in to the subscriber company intranet. A virtual port activates an IP alias on a physical port and shares all of the network settings of that port.

Configuring the External Port

To enable and configure the external port:

1. Make sure you are in the root context. If the IVS drop-down menu in the admin console header bar displays something other than Root, select Root from the menu and click Go.

2. Select **System > Network > Port 1 > Settings**.

3. Select **Enabled**.

4. Enter a valid IP address for the external port.

5. Enter a valid netmask for the IP address.

6. Enter the default gateway address.

7. Click **Save Changes**.

The system enables the port.

Configuring a Virtual Port for Sign-In on the External Port

You need to configure a virtual port to enable IVS end users to sign-in from outside the network over the external port. For example, if users sign in over the Internet, they use the virtual port defined on the external port.

To configure the virtual port for sign-in:

1. Make sure you are in the root context. If the IVS drop-down menu in the admin console header bar displays something other than Root, select Root from the menu and click Go.

2. Select **System > Network > Port 1 > Virtual Ports**.

3. Click **New Port**.

4. Enter a unique name for the virtual port.

5. Enter a valid IP address, provisioned by the subscriber company's network administrator.

6. Click **Save Changes**.

The system adds the port, displays the Virtual Ports tab, and restarts the network services. You can assign this virtual port to an IVS profile. Define as many virtual ports as needed for sign-in.

Configuring a Virtual Port for Sign-In on the Internal Port

You need to enable and configure the internal port to allow IVS end users to sign in from inside the network.

To configure the virtual port for sign-in:

1. Make sure you are in the root context. If the IVS drop-down menu in the admin console header bar displays something other than Root, select Root from the menu and click Go.

2. Select **System > Network > Internal Port > Virtual Ports**.

3. Click **New Port**.

4. Enter a unique name for the virtual port.

5. Enter a valid IP address, provisioned by the subscriber company's network administrator.

6. Click **Save Changes**.

The system adds the port, displays the Virtual Ports tab, and restarts the network services. You can assign this virtual port to an IVS profile. Define as many virtual ports as needed for sign-in.

Related Documentation

- IVS Provisioning Process Overview
- Virtual Local Area Network (VLAN) on Subscriber IVS
- Configuring VLANs on the Virtualized System

Virtual Local Area Network (VLAN) on Subscriber IVS

By defining at least one Virtual Local Area Network (VLAN) on each subscriber IVS, the MSP can take advantage of VLAN tagging, by which the virtualized system tags traffic with 802.1Q

VLAN IDs before transmitting the traffic over the backend. The carrier infrastructure uses the VLAN tag to direct the packets to the appropriate subscriber intranet.

VLAN tagging provides separation of the traffic the system transmits over the backend, destined for subscriber intranets. Traffic coming in over the front-end—that is, inbound traffic—does not have VLAN tags. The IVS adds the tag to a message upon its arrival over one of the system ports.

Each VLAN is assigned a VLAN ID which is part of an IEEE 802.1Q-compliant tag that is added to each outgoing Ethernet frame. The VLAN ID uniquely identifies each subscriber and all subscriber traffic. This tagging allows the system to direct all traffic to the appropriate VLAN and to apply respective policies to that traffic.

The VLAN termination point is any device on which VLAN-tagged traffic is identified, stripped of the VLAN tag, and forwarded to the appropriate tunnel to the backend. The VLAN termination point can be a CE router, CPE router, L2 switch, firewall, or other device capable of VLAN routing.

You must define a VLAN port for each VLAN. The root administrator assigns the specific VLAN ID when defining the VLAN port.

For each VLAN you configure, the virtualized system provisions a unique, logical VLAN interface, or port, on the internal interface. There is no relationship between the internal port IP address and any VLAN port IP address. Each VLAN port has its own route table.

Each VLAN port definition consists of:

- Port Name—Must be unique across all VLAN ports that you define on the virtualized system or cluster.

- VLAN ID—An integer in the range from 1 to 4095 that uniquely identifies the subscriber/customer VLAN.

- IP Address/Netmask—Must be an IP address or netmask from the same network as the VLAN termination point, because the virtualized system connects to the VLAN termination point on a Layer 2 network connection. VLAN IP addresses must be unique. You cannot configure a VLAN to have the same network as the internal port. For example, if the internal port is 10.64.4.30/16 and you configure a VLAN as 10.64.3.30/16, you may get unpredictable results and errors.

- Default gateway—The IP address of the default router, typically the CE or CPE router. The default gateway could act as the VLAN termination point, or could reside behind the VLAN termination point.

- Other network settings—Inherited from the internal port.

> NOTE: If you do not specify a VLAN for the subscriber company, you must configure the IVS to transmit traffic over the internal interface by selecting it as the default VLAN.

Configuring VLANs on the Virtualized System

The relationship between a VLAN and a given IVS allows the root system to separate and direct traffic to different subscribers.

Configuring a VLAN Port

Before creating a new virtual system, create a VLAN port to identify the specific subscriber traffic.

To create a VLAN port:

1. Make sure you are in the root context. If the IVS drop-down menu in the admin console header bar displays something other than Root, select Root from the menu and click **Go**.

2. Select **System > Network > VLANs** to open the VLAN Network Settings tab.

3. Click **New Port**.

4. Under VLAN settings, enter a name for the VLAN port.

5. Enter a VLAN ID.

   The VLAN ID must be between 1 and 4095 and must be unique on the system. The root system uses untagged traffic and cannot be changed.

6. Enter the IP address for the VLAN.

7. Enter a netmask for the VLAN.

8. Enter a default gateway for the VLAN.

9. Click **Save Changes**.

Assigning a VLAN to the Root IVS

To assign a VLAN to a role, you must first assign the VLAN to the root IVS. If you have not assigned a VLAN to the root IVS, the VLAN is not available in the VLAN drop-down menu in the Users > User Roles > Select Role > VLAN/Source IP page.

To assign a VLAN to the root IVS

1. Select **System > Virtual Systems > Root**.

2. Under Properties, select the VLAN from the Available VLANs list.

3. Click **Add ->** to move the VLAN name to the Selected VLANs list.

4. Click **Save Changes**.

Related Documentation

- Using VLANs with Connect Secure

Using VLANs with Connect Secure

Connect Secure enables you to create VLANs for your enterprise. VLANs are used extensively in the virtual systems. You can also create a VLAN for use in an environment in which you have deployed a device for use by all of your enterprise end users.

The VLAN feature is based on VLAN tagging, by which the system tags traffic with 802.1Q VLAN IDs before transmitting the traffic over the backend. The infrastructure uses the VLAN tag to direct the packets to your appropriate VLANs/subnets.

Traffic coming in over the front-end—that is, inbound traffic—does not have VLAN tags. The system adds the tag to a message upon its arrival over one of the device ports.

Each VLAN is assigned a VLAN ID which is part of an IEEE 802.1Q-compliant tag that is added to each outgoing Ethernet frame. The VLAN ID uniquely identifies all inbound traffic. This tagging allows the system to direct all traffic to the appropriate VLAN and to apply respective policies to that traffic.

The VLAN termination point is any device on which VLAN-tagged traffic is identified, stripped of the VLAN tag, and forwarded to the appropriate tunnel to the backend.

You must define a VLAN port for each VLAN. You assign the specific VLAN ID when defining the VLAN port.

For each VLAN you configure, the system provisions a unique, logical VLAN interface, or port, on the internal interface. There is no relationship between the internal port IP address and any VLAN port IP address. Each VLAN port has its own route table.

The Internal Port must be assigned to the root system and marked as the default VLAN. Additionally, VLAN interfaces can be assigned to the root system. All authentication servers configured for the root system, however, must have routes in the Internal Port's route table, even if the servers are reachable via VLAN interfaces. Routes to servers reachable via VLAN interfaces must have the next-hop gateway set to the configured gateway for the VLAN interface, and the output port defined as the VLAN port.

For an Active/Passive clustered deployment, the root admin of an MSP network should configure all VLAN ports with at least one virtual port (System > Network > VLANs > Virtual Ports). The router administrator must configure routes for the IVS Network Connect IP ranges that point to the VLAN virtual port's IP address as the next-hop gateway. This is required for Network Connect session failover from an IVS in the Active Node to the corresponding IVS in the Passive Node.

> ⓘ NOTE: Suppose your system has no IVS license. When you create VLANs, they are available for the entire system. However, when an IVS license is applied, the VLANs are not applied to the Root IVS. They must manually be assigned to the Root IVS or other IVSes that you create.

Each VLAN port definition consists of:

- **Port Name**—Must be unique across all VLAN ports that you define on the system or cluster.

- **VLAN ID**—An integer in the range from 1 to 4094 that uniquely identifies the VLAN.

- **IP Address/Netmask**—Must be an IP address or netmask from the same network as the VLAN termination point, because the system connects to the VLAN termination point on a Layer 2 network connection. VLAN IP addresses must be unique. You cannot configure a VLAN to have the same network as the internal port. For example, if the internal port is 10.64.4.30/16 and you configure a VLAN as 10.64.3.30/16, you may get unpredictable results and errors.

- **Default gateway**—The IP address of the default router, typically the CE or CPE router. The default gateway could act as the VLAN termination point, or could reside behind the VLAN termination point.

- **Other network settings**—Inherited from the internal port.

When you create a new VLAN port, the system creates two static routes, by default:

- The default route for the VLAN, pointing to the default gateway.

- The interface route to the directly connected network.

Figure 283 shows an example of using VLANs. In this example, we are configuring a static route in the CE router to point to the end-user's IP address, with the next-hop gateway set to the IP address of the subscriber's VLAN port.

Figure291: Example VLAN Use

The following briefly describes the steps shown in Figure 283.

1. End users sign in over an Internet connection using an IP address agreed upon by the MSP and the subscriber company.

2. Specify a static route in the MSP's VLAN termination point (in this example, a CE router) route table to point to the end-user sign-in IP addresses for each subscriber company.

3. You must also specify the next-hop gateway in the CE router's route table.

4. In the CE route table, specify all end-user sign-in IP addresses as static routes, and all corresponding VLAN port IP addresses as defined in the virtualized system.

5. Define at least one unique VLAN ID for each subscriber company. Use the IP addresses of each VLAN as the next-hop gateway addresses in the CE router's route table.

6. The subscribers' CPE routers perform the proper traffic routing to the subscriber company intranets.

7. Each subscriber company must provide sign-in pages for the IP addresses defined as static routes for end-user sign-in.

Related Documentation

- Creating a Virtual System (IVS Profile)

Adding Static Routes to the VLAN Route Table

When you create a new VLAN port, the system creates two static routes, by default:

- The default route for the VLAN, pointing to the default gateway.

- The interface route to the directly connected network.

In addition, you can static routes to shared servers in the MSP network.

To add static routes to a VLAN route table:

1. Make sure you are in the root context. If the IVS drop-down menu in the admin console header bar displays something other than Root, select Root from the menu and click **Go**.

2. Select **System > Network > VLANs**.

3. Either click **New Port** or select an existing VLAN for which to add a static route.

4. At the bottom of the VLAN port page, click the Static Routes link.

5. From the drop-down menu, select the VLAN for which to create static routes, if not already selected.

6. Click **New Route**.

7. On the New Route page, enter the destination network/IP address.

8. Enter the destination netmask.

9. Enter the destination gateway.

10. Select the interface from the Interface drop-down menu.

11. Enter the metric.

   The metric is a number between 0 and 15, and usually signifies the number of hops that traffic can make between hosts. You can set the metric of each route to specify precedence. The lower the number, the higher the precedence. Therefore, the device chooses a route with a metric of 1 over a route with a metric of 2. A router that uses metrics compares a dynamic route metric with the static route metric and chooses the route with the lowest metric.

12. If you want to add static routes to shared services, for example, you should perform one of the following steps:

   - Click **Add to [VLAN] route table**, where [VLAN] is the name of an available VLAN, to add the route to a selected VLAN. This action adds the static route to a particular subscriber company's VLAN route table and excludes access from all other VLANs, including from users of the MSP network.

   - Click **Add to all VLAN route tables** to add the route to all VLANs defined on the system. For example, if the root administrator wants to share some service among all end users of all subscriber company's, select this option.

You can also use static routes if you want to configure shared services on the MSP network. To accomplish this:

- Add a static route to the shared resource in either your own VLAN route table, if the root system has a VLAN, or in the main system route table, if the root system uses the internal interface.

- Click **Add to all VLAN route tables**, which populates all VLAN route tables with the static route. When you add the static route to all VLAN route tables, all IVS profiles can access the shared services.

Related Documentation

- Managing the Routes Table

Deleting a VLAN

You cannot delete a VLAN that is associated with an IVS. First, you must either delete the IVS or remove the relationship between the IVS and the VLAN port.

To delete a VLAN:

1. Select **System > Network > VLANs**.

2. Select the check box next to the name of the VLAN to delete.

3. Click **Delete**.

- Creating a Virtual System (IVS Profile)

Loading the Certificates Server

On the root system, you can load certificates. You must associate the virtual ports that you have defined as sign-in ports for IVS end users with the device certificate. You can specify virtual ports on the Certificate Details page.

On an IVS, you can only import Trusted Client CAs and Trusted Server CAs.

> NOTE: You cannot share certificates across IVS systems. You must have a unique IP and certificate for each IVS.
>
> You can only configure the root IVS to re-sign applets/controls in the admin console. The admin consoles for subscriber IVS systems do not show the re-signing option. You should take note of the following information:
>
> - All root and subscriber end users see the same applets/controls: either all of the default Pulse Secure controls, or all of the controls signed by the root IVS.
>
> - If you do not want subscriber IVS systems to see controls signed by the certificate from the root IVS, then you should not re-sign the controls. If you re-sign the controls, the subscriber IVS systems have access to them.

- Using Trusted Client CAs

- Using Trusted Server CAs

Creating a Virtual System (IVS Profile)

The IVS profile defines the subscriber IVS and any elements required to reach the subscriber's intranet, such as DNS settings and authentication servers.

To define the IVS profile:

1. Make sure you are in the root context. If the IVS drop-down menu in the admin console header bar displays something other than Root, select Root from the menu and click **Go**.

2. Select **System > Virtual Systems**.

3. Click **New Virtual System** to display the IVS - Instant Virtual System page.

4. Enter the name of the subscriber company.

5. Enter a description (optional).

6. Select **Enabled**, if it is not already selected.

   If you ever need to prohibit a subscriber and the subscriber's end users from accessing the IVS due to billing or other problems, disable their account here. By disabling the account, you can resolve any customer issues and then enable access without having to delete the subscriber account and lose all the configuration data.

7. Under Initial Configuration, select the configuration to initialize the IVS.

   - **Default Configuration**—Populates the IVS with system-generated defaults,

- **Root**—Copies the root system configuration to the new IVS, or any of the existing IVS's on the system

- *IVS name*—Copies the configuration from the selected IVS to the new IVS.

8. Under Administrator, create a username and password for the IVS administrator.

   The IVS administrator username and password are available in the IVS profile the first time you create the IVS. Subsequently, if you edit the IVS, these fields are not available, for security purposes. However, if you need to access the IVS administrator username and password, you can do so through the IVS configuration page, by going to the Administrators authentication server.

9. Under User Allocation, you can limit the number of concurrent users on the IVS. Specify the limit values for these options:

   - **Minimum Guaranteed Users**—You can specify any number of users between one (1) and the maximum number of concurrent users defined by the system's user license.

   - **Burstable Maximum Users** (optional)—You can specify any number of concurrent users from the minimum number you specified up to the maximum number of licensed users. If left blank, the Minimum Guaranteed Users value limits the number of concurrent users on the IVS.

     > NOTE: If you later reduce the Burstable Maximum Users setting, you must also edit the realm limits. The realm limits are not automatically updated when you change the Burstable Maximum Users setting.

10. Under Properties, specify the sign-in properties, VLAN, and port settings for the IVS.

    1. Select a VLAN from the Available list box and click **Add ->** to move the name of the VLAN to the Selected VLANs list box. You can add multiple VLANs to an IVS. You can select the internal port as a VLAN even if you have added other VLANs to the Selected VLANs list. Unlike other VLAN interfaces, you can add the internal port to multiple IVS profiles. If you have not defined a VLAN, you must select the internal interface instead.

    2. To specify the default VLAN for the IVS, select the VLAN name in the Selected VLANs list box, then click **Set Default VLAN**. The IVS marks the VLAN name with an asterisk (*). The virtualized system uses the default VLAN to provide authentication server access. The system consults the default VLAN's route table to look up the route to authentication servers for a given IVS.

       You must specify a default VLAN for each IVS. The significance of the default VLAN for a given IVS is that when an end user attempts to sign into a particular realm within that IVS, the IVS sends traffic to the authentication server for that realm over the default VLAN interface.

       > NOTE: You must specify the internal port as the default VLAN for the root IVS.

    3. If you want to define a sign-in URL prefix that your end users can sign in over rather than over a virtual port, add the prefix to the Sign-in URL Prefix field. The prefix is the equivalent of the first node in the URL, for example, companyA in the following URL:

       http://www.mycompany.com/companyA

    4. If you have defined virtual ports for either the internal interface or the external interface, you can select them in the Available list boxes and click **Add ->** to move them to the Selected Virtual Ports list boxes for the respective interfaces.

5. Enter the address or range of IP addresses that are available for Network Connect clients (end users). If you intend to configure a DNS server on the IVS, for a server located on the subscriber intranet, you must add the available Network Connect IP address pool values here.

> NOTE: Unlike the NC Profile page, you cannot specify a subnet range in the IVS profile page. In other words, entering 172.19.48.0/24 is not allowed. You can however, specify arbitrarily large IP ranges such as 172.19.48.0-172.19.48.255.

6. Click **Save Changes**.

Related Documentation

- Using Virtual Ports

- Signing In Directly to the IVS as an IVS Administrator

IVS Initial Config Via Copy from the Root System or Another IVS

When creating a new IVS profile, you have the option to copy the configuration of a root system or an existing IVS to "initialize" the new IVS. The resulting IVS combines the information you specify in the new IVS profile along with the configuration copied from the root system or existing IVS. Settings in the IVS profile override the copied settings.

Note that the copy is a one-time operation, with no on-going relationship between the source of the copy and the target of the copy. Subsequent configuration changes in the copy source are not reflected in the copy target or vice-versa.

Caveats

After performing an IVS config copy, some manual reconfiguration is required on the target.

- DNS settings are always left blank regardless of whether or not the configuration was copied from another IVS.

- When copying from an existing IVS, you should check all references to VLANs and virtual ports. For example, you may need to update the association of roles to VLAN interfaces or virtual ports.

- When copying from an existing IVS, you should check all NC connection profiles and reconfigure them, as needed, to meet the requirements of the NC IP range in the IVS profile.

- When copying from a root system, you must configure the admin roles and admin realm role mappings manually. These settings are not copied from the root system.

- Based on license restrictions of the new IVS, some admin and user realm limits may get reset. Check the admin and user realms and configure as needed.

- The Initial Configuration option allows you to select the configuration to initialize the IVS. This option does not copy the configurations of the archive servers regardless of whether you select the root system or any other IVS.

The copy operation can result in populating the target IVS with unsuitable settings. When an IVS is configured by copying the configuration from the root system or another IVS, the entire configuration from the source is copied to the target, including ACLs, resource policies, resource profiles, PAC files, and so forth, which frequently refer to specific backend servers by name, URL or IP address. In the common case, these resources tend to be company/department intranet-specific (that is, private to a particular IVS), and should not be exposed to end users of other IVSs.

It is the responsibility of the root admin to go through the entire target IVS configuration manually and purge references to any backend network resources and IP addresses that are not applicable to the target IVS.

Use Cases for IVS Initial Config Via Copy

An IVS is created from a "template" IVS

An example of this use case is provisioning virtual systems. A service provider may want to provision three categories of subscribers: Gold, Silver and Bronze. The administrator creates three IVSs (called Gold, Silver, and Bronze) and manually configures the authentication servers, roles, resource policies, etc. appropriate to the subscriber category. These IVSs are not actually used in production, but as sample configuration "templates" that can be applied to real subscriber IVSs. When provisioning a new subscriber IVS, the administrator "clones" the template by copying initial configuration for the new IVS from the Gold, Silver or Bronze "template" IVS.

It is the responsibility of the root admin to carefully craft the template IVSs to contain only those settings that are meant to be shared across multiple IVSs, such as centralized auth servers, host checker policies etc. The template IVSs should not contain references to private resources that are not meant to be exposed to/reachable from the real subscriber IVSs.

The deployment is converted to an IVS deployment

This use case is for copying the config from the root system. When an existing standalone deployment is converted to an IVS deployment by applying an IVS license and then creating multiple IVSs, the new IVSs can be configured by copying the configuration from the root system or other IVSs, and then following the steps mentioned in the caveats above.

Related Documentation

- Importing and Exporting IVS Configuration Settings

Signing In Directly to the IVS as an IVS Administrator

Signing in directly to the IVS as an IVS administrator is different than picking the IVS from the virtual system drop-down menu in the Web-based administrator UI console. If you, as the root administrator, want to sign in the same way that all IVS administrators must sign in to the IVS, perform the following steps:

1. Sign-out of the root device.

2. Enter the sign-in URL in the address bar of a valid browser, using either the hostname or the IP address. For example:

   https://www.company.com/admin
   https://10.9.0.1/admin

   This example assumes that you assigned the IP address 10.9.0.1 as a virtual port for sign-in. The format depends on whether or not you defined a DNS entry for the sign-in hostname. When logging in, the administrator can enter the hostname or the IP address defined as the virtual port for sign-in. If the administrator signs in from within the network, he should use the IP address you configured for signing in over the internal port. If the administrator signs in from outside the network, he should use the IP address you configured for signing in over the external port.

3. Press **Enter**.

4. Enter the IVS administrator username.

5. Enter the IVS password.

6. Click the **Sign in** button.

Assuming the credentials are valid, the System Status page for the IVS appears.

When either the root or an IVS administrator exits the IVS, the appliance immediately severs the connection.

Related Documentation

- Signing In to the Root System or the IVS

- Navigating to the IVS

About Role-Based Source IP Aliasing

If the subscriber company employs policy evaluation devices/firewalls in their network for the purpose of separating traffic based on the source IP address as it enters the intranet from the IVS, you, the root administrator, must configure the IVS to generate traffic with different source IP addresses. The role-based source IP aliasing feature, also known as VIP sourcing, provides the capability to map end-user roles to VLANs and specific source IP addresses (the IP address of any one of the virtual ports hosted on the VLAN interface). All traffic generated by the IVS over the back-end on behalf of the end user carries the source IP address configured for the end-user's role.

For example, assume that the traffic to a particular subscriber intranet needs to be differentiated based on whether it originates from customers, partners, or employees. There are two ways to accomplish this:

- Provision three different VLANs for the subscriber, create three roles corresponding to customers, partners and employees, and map each role to a different VLAN.

- Provision a single VLAN for the subscriber, configure three virtual ports with unique IP addresses, and map customers, partners and employee roles to the same VLAN but to different source IP addresses.

You can use role-based source IP aliasing whether or not you have defined a VLAN. In the case of a non-VLAN configuration, you define a virtual port, then assign that port to a role's source IP.

When using a VLAN, you can set the source IP address of a role to either the VLAN port IP address or to an IP alias configured on a VLAN port.

Related Documentation

- Associating Roles with Source IP Addresses in an IVS

- Using Virtual Ports

Associating Roles with Source IP Addresses in an IVS

Assuming that the root administrator has already configured a VLAN, virtual ports for the VLAN, and the IVS, the IVS administrator can associate roles with the virtual ports as follows:

1. Log in to the IVS as the IVS administrator.

2. Choose **Users > User Roles**.

3. Click **New Role**.

4. Enter a name for the role.

5. Select the **Source IP** check box.

6. Select any other options and the features you want a user with this role to be able to access (Optional).

7. Click **Save Changes**.

   The page refreshes and a set of tabs now appears.

8. Select the VLAN/Source IP tab.

9. Select the VLAN, if the root administrator has defined more than one VLAN for this IVS.

10. Select the source IP from the Select Source IP drop-down menu.

11. Click **Save Changes**.

12. Repeat the process for each new role.

When creating new users, the IVS administrator can then assign each user to one of the roles, which determines what source IP address each user can access.

Related Documentation

- About Role-Based Source IP Aliasing

Configuring Policy Routing Rules on the IVS

The virtualized system uses a policy routing framework that depends on rules, route tables, and route entries that are configured on the system.

When you create a VLAN, the system provisions a new route table for that VLAN. VLAN route tables exist in addition to the main route table for the system. Only the root administrator can manage VLAN route tables. IVS administrators cannot view or access the route tables.

Each VLAN route table contains the following route entries:

- Automatically-created route entries
- Manually-created route entries

Automatically-Created Route Entries

The default route 0.0.0.0. points to the default gateway you have configured for the VLAN interface. The system creates this route internally when it creates the VLAN interface. End users can reach most of their company's resources through the default route.

The interface route is the network route corresponding to the VLAN interface IP address.

Manually-Created Route Entries

Static routes to servers within the same VLAN that are accessible through routers other than the default gateway.

Static routes to server IP addresses on other VLAN ports within the same subscriber company intranet, or VLAN ports within the MSP network. For example, you might define in a VLAN route table static routes to DNS or authentication servers in either a subscriber company intranet or in the MSP network.

Static routes to server IP addresses accessible through the internal interface. These are usually required if your MSP network is connected to the internal interface.

Routing Rules

A number of rules have been built into the system to enable the correct routing of traffic to the appropriate subscriber intranets. For example, rules exist to map:

- The Network Connect IP pool address for each Network Connect end-user session to a corresponding VLAN route table.

  To construct this rule, the system determines an end-user's role when the user establishes an Network Connect session. The system can then search the role for the associated VLAN.

- A configured source IP address to a corresponding VLAN route table.

  The system creates this rule whenever you configure a virtual port or source IP alias on a VLAN port.

There are no explicit rules governing the flow of traffic between the subscriber or MSP networks and end users. Traffic arriving at the device over the backend has a destination IP address set to the configured IP address of one of the network interfaces, either the external interface, VLAN interface, or a Network Connect tunnel interface. the system automatically handles the processing.

You cannot access the rules table. This section includes a description of the rules table and how rules are constructed to help you understand how the system operates.

Overlapping IP Address Spaces

The virtualized system supports overlapping IP addresses in subscriber intranets, and overlapping source IP addresses for Network Connect. At this time, the virtualized system does not support multiple VLAN interfaces with identical IP addresses.

The virtualized system supports overlapping IP addresses among customer networks that are tied to VLANs in different IVS systems, because IVS systems do not share route tables.

Assume that Company 1 and Company 2 both have internal networks that use IP addresses 10.64.0.0/16. Because these addresses are internal to each company's network, and because each company has a completely separate IVS, identified by a unique VLAN ID, the MSP can support them, even though, technically, they overlap.

Define Resource Policies

Both you, as the root administrator, and the IVS administrator can create policies for end users.

You can also customize which policies are visible to IVS administrators. However, you must customize each IVS independently. Also, if you are in the root IVS context and you customize the admin console, you are only customizing the console as it appears to you or other administrators who are permitted to view the root IVS console. To customize any IVS admin console, you must be in the context of that IVS.

Related Documentation

- Resource Policy Components

- Customizable Admin and End-User UIs

Clustering a Virtualized Connect Secure System

You can cluster the entire system, including all IVS systems. You cannot cluster an individual IVS system. The clustering rules and conditions in a device network also apply to clusters in an IVS network, with the following exceptions:

- Virtual port replication—Any virtual port you define on the Active node is replicated to the Passive node. The virtual port's name and address is the same on both Active and Passive nodes.

- Virtual port source IP—Given an end user who maps to a particular role, and a backend connection from any node on behalf of that end user, the source IP of the backend connection is the same as the source IP of the virtual port configured for the end-user's role.

- VLAN port replication—When you create or delete a VLAN port on an Active cluster node, the system automatically adds or deletes the VLAN port on the Passive node.

- VLAN definition—For any given VLAN port, the slot, logical name, VLAN ID, netmask, and default gateway are the same across all cluster nodes.

- VLAN port IP address—The IP address for each VLAN port is node-specific. Corresponding VLAN ports on an Active/Passive cluster are configured on the same IP network. You can only configure an IP address/netmask combination for a VLAN port on the standby node if the resulting network corresponds to the VLAN port in the Active cluster node. VLAN IP addresses must be unique. You cannot configure a VLAN to have the same network as the internal port. For example, if the internal port is 10.64.4.30/16 and you configure a VLAN as 10.64.3.30/16, unpredictable behavior and errors can occur.

- Policy routing—You can configure route settings per node and per interface, either physical or VLAN, however, those route settings are synchronized across the cluster when you edit them.

- IVS profiles—IVS profiles are replicated across cluster nodes, and are not partitioned across cluster nodes.

- Network Connect—If you deploy the virtualized system as an Active/Passive cluster, the Network Connect connection profile that you or an IVS administrator configures within each IVS is propagated to the standby node.

- Network Connect in Active/Active cluster—In an Active/Active cluster, the Network Connect IP address pool for each IVS is split across individual cluster nodes by way of role-level settings.

- Failover behavior—In the event of a failover, the VLAN interface does not disappear. Both Active and Passive nodes should contain the VLAN interface.

> NOTE: When using Network Connect, you should always define virtual ports for each VLAN port you create. If you have defined a Network Connect IP address pool, and you are running in Active/Passive cluster mode, you must configure your routers to direct traffic to one of the VLAN's virtual ports as the next-hop gateway. Otherwise, Network Connect sessions may not recover gracefully from a failover.

Related Documentation

- *About Clustering*

Accessing a DNS Server on the MSP Network

In the root system, you can configure access so that any traffic destined for resources on the MSP network goes through the DNS server on the MSP network.

To access a DNS server on the MSP network:

1. In the admin console, **choose System > Network > Overview**.

2. Under DNS name resolution, provide the primary DNS address, secondary DNS address, and the DNS domains.

   When you add the DNS addresses, each one is added to the resolv.conf file on Connect Secure, in a nameserver directive.

3. If you are using WINS, provide the WINS server address.

4. Click **Save Changes**.

5. Configure the Network Connect connection profile.

You can provide DNS services to non-Network Connect users by specifying a global DNS/WINS server in the MSP network. The global DNS/WINS server hosts DNS for all participating subscriber companies. As an alternative, you can configure a HOSTS file on Connect Secure with DNS entries for all participating subscriber companies.

When you configure a global DNS/WINS server in this way, it provides DNS services to any requesting entity, including from Network Connect users of participating subscriber companies that do not have DNS servers in their intranets.

Related Documentation

- Creating VPN Tunneling Connection Profiles

Accessing a DNS Server on a Subscriber Company Intranet

In each IVS system, you can configure access so that any traffic destined for resources on the IVS subscriber's network goes through the DNS server on their internal company network.

To access a DNS server on a subscriber intranet:

1. If you did not add a valid Network Connect IP address pool to the IVS profile when you created the virtual system, modify the IVS profile to include the Network Connect IP addresses.

2. In the admin console, select the name of the subscriber IVS from the drop-down menu in the console header bar.

3. Click **Go**.

4. On the subscriber IVS admin console page, choose **System > Network > Overview**.

5. Under DNS name resolution, provide the primary DNS address, secondary DNS address, and the DNS domains.

6. If you are using WINS, provide the WINS server address.

7. Click **Save Changes**.

8. Configure the Network Connect Connection Profiles.

   1. Choose **Users > Resource Policies > Network Connect > Network Connect Connection Profiles**.

   2. Click **New Profile**.

   3. Provide a name for the Connection Profile.

   4. In the IP Address Pool field, enter the range of IP addresses available for use by Network Connect users.

   5. Select any other connection settings, or take the defaults.

   6. Choose a role to which to apply the settings, if necessary. By default, if you do not choose a role, the policy applies to all roles.

   7. Click **Save Changes**.

   8. Click the DNS tab.

   9. Select the **Use Custom Settings** check box.

10.   Add the Primary DNS, Secondary DNS (optional), the DNS domain name, and WINS server IP addresses.

11.   Select the DNS search order. When you enter custom settings for the IVS, the root system searches the subscriber DNS server first, then the MSP DNS server, by default.

12.   Click **Save Changes**.

> *i*   NOTE: You must perform this task for every IVS.

Related Documentation

- IVS Provisioning Process Overview

Configuring VPN Tunneling for Use on a Virtualized Connect Secure

You, as the root administrator, must work with the IVS administrator to configure VPN Tunneling so that end users can send traffic to the subscriber intranet and receive traffic back from the subscriber intranet.

If you want to use VPN Tunneling on a subscriber company's IVS (rather than just by way of VPN Tunneling running on the MSP network) you must configure a DNS server on the IVS.

For clients to be able to establish VPN Tunneling sessions to an IVS, DNS settings must be set for the IVS. Otherwise, the VPN Tunneling session initiation fails and displays an error message.

Configuring the VPN Tunneling Connection Profile

Configure the VPN Tunneling connection profile using the IP addresses from the range specified in the IP pool in the IVS profile.

1. Select **Users > Resource Policies > VPN Tunneling > VPN Tunneling Connection Profiles**.

2. Click **New Profile**.

3. Enter the IP addresses in the IP Address Pool text box, one address per line. The Help text in the admin console shows examples of valid ranges.

4. Change the transport, encryption, and compression settings from the defaults, if necessary.

5. Add the appropriate role from the Available roles listbox to the Selected roles listbox.

6. Click **Save Changes**.

Configuring VPN Tunneling on Backend Routers

Both you, as the root administrator, and the IVS administrator must configure static routes on the backend to ensure that each end user can be reached from the subscriber intranet, and if needed, the MSP network.

If you want VPN Tunneling users to be able to access the MSP network's DNS server, configure a static route in the route table of each application server or DNS server to the end-user's VPN Tunneling IP pool address. Set the next-hop gateway to the IP address of the root system's internal interface. See Figure 284.

Figure292: Setting a Static Route in MSP Network DNS or Application Servers

The following numbered list items correspond to the labeled objects in Figure 284.

1. End users sign in over an Internet connection, using an IP address from a VPN Tunneling IP address pool, to reach the DNS server on the MSP network.

2. The root administrator specifies a static route in the DNS server route table to point to an IP address from the VPN Tunneling IP address pool. The subscriber company must define the IP address pool in its intranet.

3. The DNS server resides on the MSP network and serves all end users of all subscriber companies.

4. The DNS server's route table contains a static route to the VPN Tunneling IP address pool and the next-hop gateway IP address.

5. Connect Secure's internal interface is the DNS server's next-hop gateway address.

6. The subscribers' CPE routers perform the proper traffic routing to the subscriber company intranets.

7. Each subscriber company that intends its users to pass through the MSP DNS or application servers must define a corresponding VPN Tunneling IP address pool.

As shown the following figure, the IVS administrator can configure the subscriber CPE router with a static route to the end-user's IP address, with the next-hop gateway set to the IP address of the corresponding CE router on the MSP network.

Alternately, the subscriber can configure a default route on the CPE router to point to the MSP CE router as the next-hop gateway. In this case, you do not need to add individual static routes to the VPN Tunneling IP pool addresses.

1. End users sign in over an Internet connection using an IP address agreed upon by the MSP and the subscriber company.

2. Specify a static route in the subscriber company's CPE router's route table to point to the end-user sign-in IP addresses.

3. You must also specify the next-hop gateway in the CPE router's route table.

4. You use the MSP CE router's IP address as the next-hop IP in the CPE router's route table.

5. The CPE router resides on the subscriber company's intranet. Using this arrangement, each subscriber company must specify the static route to their own end-user sign-in address and must specify the MSP's CE router IP as the next-hop gateway in the CPE route table.

6. Once the MSP VLAN termination point (in this example, a CE router) determines the intended subscriber intranet, the termination point directs the traffic to the appropriate CPE router, which sends the traffic to the proper resource in the subscriber intranet.

As shown in Figure 285, you can configure a static route in the CE router to point to the end-user's IP address, with the next-hop gateway set to the IP address of the subscriber's VLAN port.

Alternately, you can configure a default route on the CE router with the next-hope gateway set to the IP address of the subscriber VLAN port. In this case, you do not need to add individual static routes to the VPN Tunneling IP pool addresses.

You can also allocate an entire network to a VPN Tunneling IP address pool.

Figure293: Allocating an Entire Network to a VPN Tunneling IP Address Pool

The following number list corresponds to the labeled objects in Figure 285.

1. End users sign in over an Internet connection using an IP address agreed upon by the MSP and the subscriber company.

2. Specify a static route in the MSP's VLAN termination point (in this example, a CE router) route table to point to the end-user sign-in IP addresses for each subscriber company.

3. You must also specify the next-hop gateway in the CE router's route table.

4. In the CE route table, specify all end-user sign-in IP addresses as static routes, and all corresponding VLAN port IP addresses as defined in the virtualized system.

5. Define at least one unique VLAN ID for each subscriber company. Use the IP addresses of each VLAN as the next-hop gateway addresses in the CE router's route table.

6. The subscribers' CPE routers perform the proper traffic routing to the subscriber company intranets.

7. Each subscriber company must provide sign-in pages for the IP addresses defined as static routes for end-user sign-in.

Once the MSP VLAN termination point (in this example, a CE router) determines the intended subscriber intranet, the termination point directs the traffic to the appropriate CPE router, which sends the traffic to the proper resource in the subscriber intranet.

Related Documentation

- About VPN Tunneling

Configuring a Centralized DHCP Server

You can configure one or more centralized DHCP servers if you want to provide VPN Tunneling IVS users with dynamic IP addressing, without requiring each IVS subscriber to support an IVS-specific DHCP server.

The DHCP server maintains separate IP address pools for each IVS, using the IVS name property, defined in the IVS profile, to uniquely identify the IVS-specific pools.

Upon receiving a request from an IVS, the DHCP server selects an IP address based on the IVS name and the IP address of the node from which the request originated, which is delivered in the giaddr field of the request. Using this combination of data points, the DHCP server picks an available IP address from the appropriate pool and returns the address in the DHCP offer.

To configure your system to support a centralized DHCP server

The following notes apply to the use of a centralized DHCP server in an IVS configuration:

- You can configure the same DHCP server IP address for VPN Tunneling roles in multiple IVS systems.

- Within a VPN Tunneling role, if you configure both an NC IP pool and a DHCP server for the same role, the DHCP server takes precedence.

- DHCP IP address assignment can coexist with IP address assignment by way of NC IP pools within an IVS.

- You can employ multiple DHCP servers in the service provider network, with different groups of IVS systems pointing to different central servers.

1. Configure the DHCP server entry in the VPN Tunneling Connection Profile, for each VPN Tunneling role that will acquire IP addresses by way of DHCP.

2. Configure the DHCP server itself, by configuring classes and subclasses on the DHCP server to distinguish between requests from different IVS systems and to provide IP addresses from IVS-specific IP address pools.

To configure the DHCP server entry in the VPN Tunneling Connection Profile

1. In the Root context, choose **Users > Resource Policies > VPN Tunneling**.

2. Click the **VPN Tunneling Connection Profiles** tab.

3. Click **New Profile**.

4. Enter a name for the profile.

5. Under IP address assignment, select the **DHCP Server** radio button.

6. Enter the DHCP server name or IP address.

7. Under Roles, select the applicable roles in the Available roles list box and click **Add** to move them to the Selected roles list box.

8. Click **Save Changes**.

9. Repeat the procedure for each IVS that should use the DHCP server., making sure to enter the same DHCP server name or IP address that you entered for the Root.

Related Documentation

- Creating VPN Tunneling Connection Profiles

About Authentication Servers

You can configure authentication servers, such as RADIUS and Active Directory, on both the MSP network and the subscriber company intranets. The authentication server authenticates the incoming traffic differently depending on whether the traffic is authenticated when it comes into the MSP network or when it reaches the customer intranet.

> **NOTE:** If you connect an authentication server to the internal port, you must set the default VLAN to the internal port when configuring the IVS.

The following authentication servers are supported on a subscriber IVS:

- Local Authentication
- LDAP Server
- RADIUS Server
- Active Directory/Windows NT
- Anonymous Server
- Certificate Server

The following authentication servers are supported on the root system:

- Local Authentication
- LDAP Server
- NIS Server
- ACE Server
- RADIUS Server
- Active Directory/Windows NT
- Anonymous Server
- SiteMinder Server
- Certificate Server

Rules Governing Access to Authentication Servers

The following rules apply to the access of authentication servers on the MSP network or on the subscriber company network. Each IVS profile must include settings for:

- The default VLAN, which can also be the internal port, if provisioned as the default VLAN.
- The default VLAN interface IP is the source IP address used to contact the authentication server.
- Static routes in the VLAN that point to the appropriate authentication servers, which can reside in the MSP network (with an assigned VLAN ID or untagged on the internal port), or on the subscriber company network.

Configuring Authentication on a RADIUS Server

You must configure the RADIUS server in each IVS. If you have a RADIUS server on the MSP network as well, all of the IVS RADIUS servers can point to the same MSP RADIUS IP address.

To configure the RADIUS server:

1. Select the context:

   - If you are in an IVS context, and you want to define a RADIUS server on the MSP network, select Root from the context drop-down menu in the admin console header bar and click Go.

   - If you are in the root context, and you want to define a RADIUS server on a subscriber intranet, select the IVS name from the context drop-down menu in the admin console header bar and click Go.

2. Follow the same steps as you would for configuring a RADIUS server instance.

   NOTE: In the current release, ACE authentication is not available for individual IVS systems. If you want to use RSA 2 factor token-based authentication, you should use RADIUS from the IVS to access RSA ACE.

Configuring Authentication on Active Directory

You must configure the AD/NT server in each IVS. If you have an AD/NT server on the MSP network as well, all of the IVS AD/NT servers can point to the same MSP AD/NT IP address.

To configure the Active Directory server:

1. Select the context:

   - If you are in an IVS context, and you want to define an AD/NT server on the MSP network, select Root from the context drop-down menu in the admin console header bar and click Go.

   - If you are in the root context, and you want to define an AD/NT server on a subscriber intranet, select the IVS name from the context drop-down menu in the admin console header bar and click Go.

2. Follow the same steps as you would for configuring an Active Directory or NT Domain instance.

Related Documentation

   - Using a RADIUS Server

   - Using Active Directory

Delegating Administrative Access to IVS Systems

As the root administrator, you can delegate administrative access and responsibilities to specific IVS systems. You can delegate read/write access or read-only access to all IVS systems, or to selected IVS systems.

To delegate administrative access to IVS systems

1. Select Administrators > Admin Roles > *Role* where *Role* indicates one of the listed administrator roles. You can also create a new administrator role, if you prefer.

2. Click the **IVS** tab.

3. To give the administrator read/write access to the IVS, select one of the following:

- To give the administrator read/write access to all IVS systems, select the **Administrator can manage ALL IVSs** check box.

- To limit the administrator's access to specific IVS systems, select the **Administrator can manage SELECTED IVSs** check box, then select the IVS systems from the Available IVSs list and click **Add** to move them to the Selected IVSs list.

4. To give the administrator read only access to the IVS, select one of the following:

- If you want to give the administrator read only access to all IVS systems, select the **Administrator can view (but not modify) ALL IVSs** check box.

- If you want to limit the administrator's access to specific IVS systems, select the **Administrator can view (but not modify) SELECTED IVSs** check box, then select the IVS systems from the Available IVSs list and click **Add** to move them to the Selected IVSs list.

5. Click **Save Changes**.

By adding these access rights to a given role, you can exercise different levels of control over different MSP administrators.

Related Documentation

- Creating and Configuring Administrator Roles

Accessing Standalone Installers on an IVS System

The IVS administrator might need to access the Host Checker, WSAM, or other standalone installers. To give IVS administrators access to the installers, which are located on the Maintenance > System > Installers page, you can delegate the access to them by way of the Administrators > Admin Roles > SelectRole > IVS page. Once you have delegated access, the IVS administrator can see the Installers page from within the context of the IVS admin console.

Related Documentation

- Downloading Client Installer Files

Exporting and Importing IVS Configuration Files

Use the binary import/export feature to export and import root system and user settings, and also to export and import subscriber IVS settings and profiles. The two types of operations are mutually exclusive: if exporting IVS settings, the exported configuration file does not contain root system settings; if exporting root system settings, the exported configuration file does not contain subscriber IVS settings.

Use the XML import/export feature to export and import system configuration. This feature enables you to make significant changes to your system configuration and provides a number of benefits, particularly when it comes to make a large number of repetitive changes, or when you want to update configuration data all at once.

You perform export and import operations from the context of the root system. On the Maintenance > Import/Export > Import/Export Configuration page, the Maintenance > Import/Export > Import XML and Export XML pages, and on the Maintenance > Import/Export > Import/Export Users page, you can find the standard controls for exporting root system and user configuration. A subscriber IVS administrator cannot export or import data from or to a subscriber IVS. Only you, as the root administrator, can perform these tasks.

Note the following:

- You can only import/export all IVS systems in a single operation. You cannot import/export an individual IVS system's configuration.

- When creating an IVS using XML import, an XML containing the IVS profile should be imported first before importing the IVS configuration.

- The schema for XML import/export of a root IVS (click the Download the schema files link in the XML export page) is different from the XML import/export schema of a non-root IVS.

- You can also use the binary archiving feature to perform local backups of your IVS system.

Exporting IVS Configurations

To export IVS configurations:

1. Select **Maintenance > Import/Export > Import/Export IVS**.

2. Enter a password to password protect the configuration file.

3. Click **Save Config As**.

4. Click **Save**.

5. Provide a file name and target location for the file.

6. Click **Save** and **Close**, if necessary.

The saved configuration file contains the following settings for all IVS systems:

- IVS Profiles

- IVS System Settings

- IVS Signing In Settings

- IVS Administrators

- IVS Users

- IVS Resource Policies

- IVS Maintenance Settings

Importing IVS Configurations

To import IVS configurations:

1. Select **Maintenance > Import/Export > Import/Export IVS**.

2. Click **Browse**.

3. Locate and select the file and click **Open**.

4. Enter a password to password protect the configuration file.

5. To import the network settings in the IVS profile, such as VLAN ports and virtual ports, select the **Import IVS Profile Network Settings** check box.

   NOTE: Importing network settings as described in above only works if you export the system and IVS configurations from the same system.

   The network settings themselves do not get imported; only the references to the network settings get imported. Network settings are only imported/exported when importing/exporting the root system settings.

6. Click **Import Config**.

The IVS provides a confirmation message if the import operation succeeds. The IVS then restarts certain services, which may require several minutes.

> **NOTE:** You can use the XML Import/Export feature to export and import XML-based configuration files on the root IVS.
>
> You can use Push Config to copy one root IVS configuration to another root IVS. You cannot use Push Config to copy configuration data between subscriber IVS systems or from a root IVS to a subscriber IVS.

Related Documentation

- Configuration File Administration Overview

Using XML Import and XML Export on IVS Systems

XML import and export provides a way to make a large number of repetitive changes all at once by importing and exporting system configuration. As with the binary import/export, a subscriber IVS administrator cannot export or import data from or to a subscriber IVS. Only the root administrator can perform this task. You cannot create or delete an IVS through XML import.

The process for importing and exporting IVS system configuration is the same as for Connect Secure and will not be described here. Note that you will see fewer XML export options on the IVS as compared to Connect Secure as not all options are applicable to IVS.

As with Connect Secure, some services might be restarted after performing an IVS XML import due to changes in configuration. The following configuration changes result in service restarts:

- NCP option changes

- Windows/UNIX file browsing config changes

- Security option changes

The IVS root administrator can restrict IVS administrators from changing configuration data that can cause service restarts (both from the admin console and through XML import) by selecting System > Virtual Systems > Virtual Systems and then selecting the Restrict IVS administrator access to configuration without system-wide impact option.

Related Documentation

- Configuration File Administration Overview

Monitoring Subscribers

Log files contain detailed information about events, user access, administrator access and more. The log entries specify the context of each entry, whether the entry was caused by a root action or an action on one of the IVS systems. The root entries contain the word Root. For example, the following entries show access by two administrators, the first being Root and the second, an administrator called Test:

```
ADM20716 2005-05-10 10:52:19 - ive - [10.11.254.160]
Root::administrator(administrator Users)[.Administrators] - User Accounts
modified. Added Unspecified Name with username testuser1 to authentication
server System Local.
```

```
Info ADM20716 2005-05-10 10:35:26 - ive - [10.11.254.160]
Test::administrator(administrator Users)[.Administrators]!Root - User Accounts
modified. Added IVE Platform Administrator with username omiadmin to
authentication server Administrators.
```

Suspending Subscriber Access to the IVS

To suspend subscriber access to the IVS:

1. Select **System > Virtual Systems**.

2. Click the **Disabled** radio button.

By performing this step, you make the IVS unavailable to any user of the IVS, including the IVS administrator. To provide access to the IVS, set the radio button to the Enabled state.

Related Documentation

- Virtual Local Area Network (VLAN) on Subscriber IVS

Troubleshooting VLANs

In addition to the standard troubleshooting features, the virtualized system provides several enhancements, specifically for managing IVS systems. You can use the following troubleshooting features on either the root system or each IVS, separately:

- Policy simulation

- Policy tracing

- Session recording

Functionally, these utilities are the same as the standard system capabilities. The key difference is a matter of context. If you initiate one of these three utilities from the root system context, you get results for users, policies, and sessions on the root system or from the MSP network. If you initiate the utilities from a subscriber IVS context, you get results for users, policies, and sessions on the IVS or the subscriber intranet.

The TCPDump, Ping, Traceroute, NSLookup, and ARP commands are enhanced for use in virtualized systems. You can initiate these commands on the internal and external ports, as well as on selected VLAN ports, which you might do if you want to troubleshoot traffic on a subscriber VLAN. The basic functionality of the commands is unchanged, except for the ability to specify a VLAN port

Running TCPDump on a VLAN

To perform a TCPDump on a VLAN:

1. If you are not in the root system context, select **Root** from the IVS drop-down menu in the admin console header, and then click **Go**.

2. Choose **Maintenance > Troubleshooting > Tools > TCP Dump**.

3. With Internal Port selected, select the VLAN from the VLAN Port drop-down menu.

4. (optional) Add a filter to the Filter text box.

5. Click **Start Sniffing**.

6. To retrieve the results, click **Stop Sniffing**.

7. Choose the type of Dump file from the Dump File drop-down menu.

8. Click **Get**.

9. Open the file with the appropriate editor.

Using Ping, traceroute, NSLookup, ARP Commands on a VLAN

1. If you are not in the root system context, select **Root** from the IVS drop-down menu in the admin console header, and then click **Go**.

2. Choose **Maintenance > Troubleshooting > Tools > Commands**.

3. Select a command from the Command drop-down menu.

4. Enter the target server.

5. Select the VLAN from the VLAN Port drop-down menu.

6. Enter the other settings, depending on the command you choose.

7. Click **OK**.

Related Documentation

- *Troubleshooting Tools Overview*

IVS Use Case: Policy Routing Rules Resolution

This use case illustrates how policy routing takes place in an MSP deployment. The first part of the use case details two subscriber company configurations and how end users access their respective subscriber company networks. The second part of the use case describes what happens when you create a VLAN on the MSP network to provide shared services to the subscriber companies' end users.

Company 1 and Company 2 are hosted companies on the MSP network. The following table shows the VLANs, VLAN IDs, interfaces and roles defined for each company. Company 1 has defined two VLANs, one for Sales and one for Human Resources. Each company has an associated role defined for each VLAN. The root administrator creates each VLAN, providing a unique VLAN ID for each, and indicating a given port. In this case, the root administrator has created all four VLANs on the internal interface

	VLAN	VLAN ID	Interface	Role
Company 1	Sales	1	int0.1	SALES
	HR	2	int0.2	HR
Company 2	Employee	3	int0.3	EMPLOYEE
	Partner	4	int0.4	PARTNER

NOTE: The labels for ports have been changed. The port name eth0 (internal port) is now called int0 and eth1 (external port) is now ext0. You can only see the route table device names (such as int0.1) from the serial console. You can view the route table by selecting menu item 1, then menu item 2 from the serial console.

Figure294: MSP and Subscriber Company Deployments.

> **NOTE:** IVS VLANs are not explicitly tied to subscriber intranets by configuration on Connect Secure. The association of a VLAN to a subscriber intranet is accomplished by mapping VLAN interfaces to private tunnels in the subscriber intranet within the CE->CPE router framework.

In figure 286, VPN Tunneling end users get their source IP addresses from configured VPN Tunneling IP address pools that the root administrator has defined for the IVS. Also, in the figure, non-VPN Tunneling users can still access specified realms based on their roles and on role-based source IP (VIP sourcing) addresses that you define as virtual ports on the VLAN.

The following list describes each item that is marked with a numbered label in the above figure.

1. VPN Tunneling end users get IP addresses from VPN Tunneling IP pools. Traffic from these users is routed through the appropriate subscriber VLAN, which you define on the internal port.

2. Non-VPN Tunneling end users get IP addresses from virtual IP (VIP) pools. Traffic from these users is sourced through the appropriate subscriber VLAN.

3. In Figure 67, this numbered box represents a subscriber IVS, which contains two VLANs that are defined on ports int0.1 and int0.2.

4. In Figure 67, this numbered box represents a second subscriber IVS, which contains two VLANs that are defined on ports int0.3 and int0.4.

5. The subscriber defines a role for "Sales" on VLAN1. End users signing in to IP 13.10.0.1 over the Internet are routed to the Company 1 intranet, to the appropriate backend resources located in the "Sales" realm at 10.10.0.0/16. End users signing in on IP 13.11.0.1 are VIP sourced to the Company 1 intranet, also to the appropriate backend resources located in the "Sales" realm at 10.10.0.0/16.

6. The subscriber defines a role for "HR" on VLAN2. End users signing in on IP 13.12.0.1 over the Internet are routed to the Company 1 intranet, to the appropriate backend resources located in the "HR" realm at 10.11.0.0/16. End users signing in on IP 13.13.0.1 are VIP sourced to the Company 1 intranet, also to the appropriate backend resources located in the "HR" realm at 10.11.0.0/16.

7. The subscriber defines a role for "Employee" on VLAN3. End users signing in on IP 14.10.0.1 over the Internet are routed to the Company 2 intranet, to the appropriate backend resources located in the "Employee" realm at 10.10.0.0/16. End users signing in on IP 14.11.0.1 are VIP sourced to the Company 2 intranet, also to the appropriate backend resources located in the "Employee" realm at 10.10.0.0/16.

8. The subscriber defines a role for "Partner" on VLAN4. End users signing in on IP 14.12.0.1 over the Internet are routed to the Company 2 intranet, to the appropriate backend resources located in the "Partner" realm at 10.11.0.0/16. End users signing in on IP 14.13.0.1 are VIP sourced to the Company 2 intranet, also to the appropriate backend resources located in the "Partner" realm at 10.11.0.0/16.

9. The Company 1 intranet supports two realms: "Sales" at 10.10.0.0/16 and "HR" at 10.11.0.0/16. These realms correspond to the roles defined on VLAN1 and VLAN2/

10. The Company 2 intranet supports two realms: "Employee" at 10.10.0.0/16 and "Partner" at 10.11.0.0/16.

NOTE: The realms are valid even though they contain overlapping IP addresses. Because the roles are defined for different VLANs, the VLAN IDs provide the separation that allows them to overlap without danger of mixed traffic.

The route tables for each VLAN appear as follows:

Table181: VLAN1 Route Table

Destination IP	Gateway	Output Port
0.0.0.0	Default gateway on VLAN1	int0.1
10.10.0.0/16	0.0.0.0	int0.1

Table182: VLAN2 Route Table

Destination IP	Gateway	Output Port
0.0.0.0	Default gateway on VLAN2	int0.2
10.10.0.0/16	0.0.0.0	int0.2

Table183: VLAN3 Route Table

Destination IP	Gateway	Output Port
0.0.0.0	Default gateway on VLAN3	int0.3
10.10.0.0/16	0.0.0.0	int0.3

Table184: VLAN4 Route Table

Destination IP	Gateway	Output Port
0.0.0.0	Default gateway on VLAN4	int0.4
10.10.0.0/16	0.0.0.0	int0.4

Now consider the situation in which the MSP decides to provide shared services to end users of Company 1 and Company 2. Assume the MSP network is also on a VLAN (VLAN5). If you want to provide services on 10.64.0.0/16 to both Company 1 and Company 2, and services on 10.65.0.0/16 to Company 2 only, you can configure either VPN Tunneling pools or virtual ports for those addresses.

Figure 287 illustrates this situation. Some details from this figure have been removed or greyed out to improve readability.

Figure295: Providing shared services

1. Company 1 end users sign-in over the Internet to the MSP network and the MSP VLAN, VLAN5 (represented as number 3 in the illustration).

2. Company 2 end users sign-in over the Internet to the MSP network and the MSP VLAN, VLAN5 (represented as number 3 in the illustration).

3. The MSP VLAN5 provides access to shared services on the MSP network.

4. You must define separate IP addresses for each subscriber company's end users, even though they share MSP services.

Once you configure routes to support users who have access to shared services on the MSP network and to support users who also have access to restricted MSP network services, the VLAN route tables appear as follows.

Table185: VLAN1 Route Table

Destination IP	Gateway	Output Port
0.0.0.0	Default gateway on VLAN1	int0.1
10.64.0.0	Router on VLAN5	int0.5

Table186: VLAN2 route table

Destination IP	Gateway	Output Port
0.0.0.0	Default gateway on VLAN2	int0.2
10.64.0.0	Router on VLAN5	int0.5

Table187: VLAN3 route table

Destination IP	Gateway	Output Port
0.0.0.0	Default gateway on VLAN3	int0.1
10.64.0.0	Router on VLAN5	int0.5
10.65.0.0	Router on VLAN5	int0.5

Table188: VLAN4 route table

Destination IP	Gateway	Output Port
0.0.0.0	Default gateway on VLAN4	int0.2
10.64.0.0	Router on VLAN5	int0.5
10.65.0.0	Router on VLAN5	int0.5

If the MSP network is connected to the untagged port (internal), the route entries are similar, but the output port is int0 only.

Related Documentation

- IVS Provisioning Process Overview

Use Case: Configuring a Global Authentication Server for Multiple Subscribers

If your subscriber companies prefer to lease or purchase authentication services from you, the service provider, you can configure a global authentication server on your network. In that case, you must perform several tasks:

1. Configure one or more authentication servers on your MSP network.

2. Configure path-based URLs or virtual ports for sign-in on your MSP network.

3. Configure VLANs and IVS systems to map to the authentication servers on the MSP network.

Related Documentation

- IVS Provisioning Process Overview

Use Case: Configuring a DNS/WINS Server IP Address per Subscriber

If you want to configure a particular DNS/WINS server IP address per subscriber, you can do so from within each IVS.

To configure a DNS/WINS server IP address:

1. Configure your IVS systems.

2. Select an IVS from the system drop-down menu in the admin console header area and then click **Go**. Within the IVS context, the header color changes and displays the name of the subscriber.

3. Select **System > Network > Overview**.

4. Enter the DNS/WINS settings that correspond to the DNS/WINS server on the subscriber intranet.

5. Click **Save Changes**.

Related Documentation

- IVS Provisioning Process Overview

Use Case: Configuring Access to Web Applications and Web Browsing for Each Subscriber

The IVS administrator may want to configure specific Web browsing policies for the IVS end users.

To configure Web browsing access, the IVS administrator needs to configure the following pages:

- Users > User Roles > *RoleName* > Web

- Users > Resource Policies > Web

Configuring Web Browsing Access

To configure Web browsing access:

1. Choose **Users > User Roles >** *RoleName* **> Web**.

2. Select the **Bookmarks** tab.

3. Click **New Bookmark**.

4. Supply settings to configure the bookmark to a given Web URL.

5. Click **Save Changes** or **Save + New** if you want to add multiple bookmarks.

    The bookmarks you define here appear in the Web bookmarks section to which end users have access.

6. Select the **Options** tab.

7. Select the Web browsing privileges you want to provide to your end users.

8. Choose the other options you want, including setting the timeout value for the HTTP connection.

9. Click **Save Changes**.

Configuring Web Browsing Access Policies

To configure Web browsing access policies:

1. Choose **Users > Resource Policies > Web**.

2. Supply the appropriate settings on each of the tabs.

Related Documentation

- IVS Provisioning Process Overview

Use Case: Configuring File Browsing Access for Each Subscriber

The IVS administrator may want to configure specific file-browsing access policies for the IVS end users. The IVS administrator can perform this type of operation based on roles.

To configure file browsing, the IVS administrator needs to configure the following pages:

- Users > User Roles > *RoleName* > General

- Users > User Roles > *RoleName* > Files

- Users > Resource Policies > Files

Configuring File Browsing Access

To configure file browsing access:

1. Choose **Users > User Roles > RoleName > General**.

2. Under Access Features, select the **Files** check box (for Windows).

3. Click **Save Changes**.

4. Select the **Files** tab.

5. Select the **Options** page.

6. Depending on the file system type, select the options that apply to the IVS end users access.

7. Click **Save Changes**.

Configuring File System Access Policies

To configure file system access policies:

1. Make sure you are in the IVS context. If the IVS drop-down menu in the admin console header bar displays Root, select the IVS name from the menu and click Go.

2. Select **Users > Resource Policies > Files > Access > Windows**.

3. Choose the role from the Show policies that apply to drop-down menu, and click **Update**.

4. Click **New Policy**.

5. Supply the appropriate settings.

6. Click **Save Changes**.

7. Select the Credentials tab.

8. Supply the appropriate settings.

9. Click **Save Changes**.

10. Repeat these steps for each role.

11. Select the **Encoding** tab to select the language encoding and click **Save Changes**.

12. Select the **Options** tab to set options, such as IP based matching for Hostname based policy resources and click **Save Changes**.

> NOTE: UNIX/NFS file resource policies are supported only on the root IVS.

Related Documentation

- IVS Provisioning Process Overview

Use Case: Setting Up Multiple Subnet IP Addresses for a Subscriber's End-Users

Assume that the subscriber wants to create subnets within the intranet to support traffic separation between subscriber end users from three different departments: Marketing, Finance, and Engineering. The procedures needed to accomplish this task are divided between those performed by the root administrator and those performed by the IVS administrator.

Tasks Performed by the Root Administrator

1. Create subscriber VLAN.

2. Create subscriber IVS.

3. Create path-based URLs or virtual ports for sign-in.

4. Create virtual ports for role-based source IP aliasing.

Tasks Performed by the IVS Administrator

1. Create users.

2. Assign roles to VLAN/Source IP.

3. Assign users to roles.

Related Documentation

- IVS Provisioning Process Overview

- Configuring VLANs on the Virtualized System

- Creating a Virtual System (IVS Profile)

- Using Virtual Ports

- About Role-Based Source IP Aliasing

- User Roles Overview

- Associating Roles with Source IP Addresses in an IVS

Use Case: Configuring Multiple IVS Systems to Allow Access to Shared Server

There may be cases in which you want to provide end users of multiple subscriber companies to access a shared server on the MSP network.

The following steps describe a simple use case and solutions.

Solution #1

To configure access to a shared server, assuming two IVS systems for two subscribers:

1. Add the internal port to the IVS1 list of selected VLANs.

2. Add the internal port to the IVS2 list of selected VLANs.

3. Edit the internal port's route table and configure a static route pointing to the shared server, with the internal interface as the output port.

Solution #2

To configure access to a shared server, assuming two IVS systems for two subscribers:

1. Add VLAN1 to the IVS1 selected VLAN list and set it as the default VLAN.

2. Add VLAN2 to the IVS2 selected VLAN list and set it as the default VLAN.

3. Edit the route tables for both VLAN1 and VLAN2 and configure a static route in each that points to the shared server, with the internal interface as the output port.

Related Documentation

- Configuring VLANs on the Virtualized System

- IVS Provisioning Process Overview

- Creating a Virtual System (IVS Profile)

CHAPTER 42 Deployments with IDP

About IDP

Securing intranet work application and resource traffic is vital to protecting your network from hostile outside intrusion. You can add levels of application security to your remote access network by integrating a Connect Secure system with a Juniper Networks Intrusion Detection and Prevention (IDP) Sensor. The IDP device may provide the following types of protection in this solution (some forms of protection depend upon the specific configuration):

The IDP sensor monitors the network on which the IDP system is installed. The sensor's primary task is to detect suspicious and anomalous network traffic based on specific rules defined in IDP rulebases.

The IDP device provides the following types of protection (some forms of protection depend upon the specific configuration):

- Protects against attacks from user to application and from application to user (from a server-side endpoint)

- Detects and blocks most network worms based on software vulnerabilities

- Detects and blocks non-file-based Trojan Horses

- Detects and blocks effects of spyware. adware, and key loggers

- Detects and blocks many types of malware

- Detects and blocks zero day attacks through the use of anomaly detection

  > NOTE: An IDP Sensor can send logs to one Connect Secure device only. However, the Connect Secure device can receive logs from more than one IDP Sensor.

You do not need a special license from Pulse Secure to enable interaction between Connect Secure and the IDP.

Using the Connect Secure admin console, you can configure and manage interaction attributes between it and an IDP, including the following:

- Global configuration parameters such as the IDP hostname or IP address, the TCP port over which the sensor communicates with Connect Secure, and the one-time password Connect Secure and IDP use to authenticate with one another.

- Dynamically changing the IDP configuration from Connect Secure and alerting the IDP of changes in the IP address pool available to remote users.

- Various levels of attack severity warnings.

The IDP sits behind Connect Secure on your internal network and monitors traffic flowing from Connect Secure into the LAN. Any abnormal events detected by the IDP Sensor are reported to Connect Secure, which you configure to take appropriate action based on the severity level of the reported events. The IDP Sensor performs reporting functions in addition to any normal logging the IDP has been configured to undertake.

You can use an IDP Sensor on the Connect Secure cluster, if the cluster is configured with a virtual IP (VIP) address.

Related Documentation

- IDP Deployment Scenarios

- Configuring Connect Secure to Interoperate with IDP

- Configuring IDP Sensor Policies

IDP Deployment Scenarios

The two most likely deployment scenarios are as follows:

- Customer use of Connect Secure for extended enterprise access and IDP for security of all perimeter traffic including but not limited to traffic from Connect Secure. Figure 288 illustrates this scenario, in which Connect Secure is deployed in the DMZ or on the LAN and the IDP is deployed in-line behind the firewall and in front of the LAN.

Figure296: Connect Secure and IDP Topology Scenario 1

- In the second deployment scenario, IDP is only used to protect traffic that comes through Connect Secure but not in-line with other perimeter traffic. Figure 289 illustrates this deployment scenario.

Figure297: Connect Secure and IDP Topology Scenario 2

Related Documentation

- About IDP

- Configuring Connect Secure to Interoperate with IDP

Configuring Connect Secure to Interoperate with IDP

The IDP Sensor is a powerful tool to counter users who initiate attacks. Integration with Connect Secure allows you to configure automatic responses as well as manually monitor and manage users.

To configure the system to interoperate with an associated standalone IDP Sensor, you must first ensure the IDP has been configured according to the instructions described in the Signaling Setup appendix of IDP Series Concepts and Examples Guide, Version 5.1rX.

Once the IDP Sensor has been set up, you can specify the events you want the IDP to watch for and the actions that Connect Secure takes once a particular event has been noted and reported.

There are two locations on Connect Secure where you can specify actions to be taken in response to users that perform attacks:

- **Sensor Event policies page**—Define the policy on this page to generate an automatic response to users who perform attacks.
- **Users page**—Manually identify and quarantine or disable users on the System > Status > Active Users page, which lists users who have performed attacks.

Interaction Between the IC Series and IDP

Connect Secure reads attack information as it is being sent by the IDP sensor. Connect Secure receives the source and destination IP addresses and port numbers of the attacking host and the resource against which the attack was launched, along with the attack identifier, severity of the attack, and the time at which the attack was launched.

Connect Secure incorporates and displays the attack information received from the IDP sensor on the System > Status > Active Users page. Based on the attackers IP address and port number, the system can uniquely identify the user's session.

You can choose automatic or manual actions for attacks detected by the IDP sensor. For manual action, you look up the information available on the Active Users page and decide on an action. For automatic action, you configure the action in advance when you define your IDP policies.

Related Documentation

- Configuring IDP Sensor Policies

Configuring IDP Sensor Policies

The Sensors tab allows you to specify the system settings used to establish a connection to a Juniper Network's Intrusion Detection and Prevention (IDP) device.

Use the System > Configuration > Sensors > Sensors tab to perform a number of tasks related to configuring and managing interaction between Connect Secure and an IDP Sensor. The main Sensor page displays the sensor, the network address, the state (enabled), the version, and the status of any configured sensors.

Creating a New IDP Sensor Entry In IDP versions prior to 5.0, Connect Secure sends only the user IP address. With IDP version 5.0 and later, Connect Secure sends session information including the user, user role, and IP address.

To enable or disable existing IDP Sensor entries on Connect Secure:

1. In the admin console, choose **System > Configuration > Sensors**.

> **NOTE:** To use the IDP sensor with Connect Secure you must enable logging for the applicable policies.

2. Click **New Sensor**. The admin console displays the New Sensor page.

3. Under Sensor Properties, specify the following information:

   - **Name**—A name Connect Secure uses to identify the new connection entry

   - **Hostname**—The hostname or IP address of the IDP Sensor to which Connect Secure connects to receive application and resource attack alert messages

   - **Port**—The TCP port on the IDP Sensor to which Connect Secure listens when receiving application and resource attack alert messages

   - **One-time password**—The encrypted password Connect Secure uses when conducting the initial Transport Layer Security (TLS) handshake with the IDP Sensor. You must enter the encrypted OTP password as displayed on the IDP ACM configuration summary screen.

   > **NOTE:** The hostname, TCP port, and one-time password must already be configured on the IDP Sensor before this configuration can be successful.

4. Under **Monitoring Options**, specify IP addresses to monitor and the minimum alert severity level the IDP Sensor will record and submit to Connect Secure:

   - In the Addresses to Monitor field, specify individual IP addresses and address ranges, one entry per line. IDP reports attack information only for the IP addresses that you specify. If you want IDP to report all events, enter 0.0.0.0/0. If you want IDP to report only selected events, enter <default> to permit IDP to report events for events with source IPs that have an active user session on Connect Secure, and /or enter one or more addresses or address ranges for any endpoint that you want the IDP sensor to report.

   - Select one of the severity options available in the Severity filter drop down list. The severity level is a number on a scale from 1 to 5, where 1 is informational and 5 is critical. This option represents the severity of messages the IDP should send.

5. Click **Save Changes**.

Enabling or Disabling IDP Sensors

To enable or disable existing IDP Sensor entries on Connect Secure:

1. In the admin console, choose **System > Configuration > Sensors**.

2. Select the check box next to one or more IDP Sensor entries you want to enable or disable.

3. Click **Enable** or **Disable** to enable or disable the specified IDP Sensor entries, respectively.

You can delete existing IDP Sensor entries that define a connection between Connect Secure and an IDP Sensor.

To delete one or more existing IDP Sensor entries from Connect Secure:

1. In the admin console, choose **System > Configuration > Sensors**.

2. Select the check box next to the IDP Sensor entry or entries you want to delete.

3. Click **Delete** and then confirm that you want to delete the sensor entry or entries.

Reconnecting to an IDP Sensor

When the connection to an IDP Sensor is down, you can use the admin console on Connect Secure to re-establish the connection. You can also use the admin console to refresh the status of existing connections between Connect Secure and the IDP Sensor.

If you need to re-establish communication with an IDP Sensor, you must generate a new One-time Password.

To reconnect to an associated IDP Sensor:

1. In the admin console, choose **System > Configuration > Sensors**.

2. Select the check box next to the IDP Sensor to which you want to reconnect.

3. Click **Reconnect**.

   The admin console displays a message informing you that Connect Secure is currently attempting to re-establish connection to the specified IDP Sensor. This page automatically refreshes each second during the reconnection process. Otherwise, the connection status page automatically refreshes once every 30 seconds.

Refreshing and Displaying the Connection Status

To refresh and display the connection status for the specified IDP Sensor:

1. In the admin console, choose **System > Configuration > Sensors**.

2. Select the check box next to one or more IDP Sensor entries for which you want to display current connection status.

3. Click **Refresh**.

Related Documentation

- About IDP

- Configuring Connect Secure to Interoperate with IDP

Defining Automatic Response Sensor Event Policies

Use the **System > Configuration > Sensors > Sensor Event Policies** tab to specify one or more rules that specify the action(s) Connect Secure takes when it receives attack alert messages from an IDP Sensor.

To create a new IDP rule:

1. In the admin console, select **System > Configuration > Sensors > Sensor Event Policies**.

2. On the Sensor Event Policies page, click **New Rules**.

3. On the Pulse Secure IDP Rule page, in the Rule: On Receiving... section:

   - Select an existing event from the Event drop-down list.

   - Click **Events** to edit an existing event or create a new type of event and add it to the options in the Events drop-down list:

   1. Specify a name for the event.

   2. Populate the Expressions field by manually entering expressions or by selecting one or more clauses from the Expressions Dictionary and clicking Insert Expression.

      For example, to check for all critical/highest severity level attacks, enter the following expression:

      idp.severity >= 4

      To check for all critical/highest severity level attacks for HTTP traffic, enter the following expression:

      idp.severity >= 4 AND idp.attackStr = "*HTTP*"

3. When you have finished entering the expressions you want to apply to this event, click **Add Expression**.

4. Click **Close**.

5. In the **Count this many times** section, specify a number between 1 and 256 to determine the number of times an event must occur before action is taken.

6. In the **...then perform this action** section, specify one of the following actions:

   - **Ignore (just log the event)**—Specifies that the system should log the event, but take no further action against the user profile to which this rule applies. This option is best used to deal with very minor "informational" attack alert messages that come from the IDP Sensor.

   - **Terminate User Session**—Specifies that Connect Secure should immediately terminate the user session and require the user to sign in again.

   - **Disable user account**—Specifies that Connect Secure should disable the user profile associated with this attack alert message, thus rendering the client unable to sign in until the administrator re-enables the user account. (This option is only applicable for users who have a local Connect Secure user account.)

   - **Replace user's role with this one**—Specifies that the role applied to this user's profile should change to the role you select from the associated dropdown list. This new role remains assigned to the user profile until the session terminates. This feature allows you to assign a user to a specific controlled role of your choice, based on specific IDP events. For example, if the user performs attacks, you might assign the user to a restricted role that limits the user's access and activities.

   - Choose to **make this role assignment:**

     - **Permanent**—User remains in the quarantined state across subsequent logins until the administrator releases the user from the quarantined state.

     - **For this session only**—Default. User can log in to another session.

7. In the **Roles** section, specify:

   - **Policy applies to ALL roles**—To apply this policy to all users.

   - **Policy applies to SELECTED roles**—To apply this policy only to users who are mapped to roles in the Selected roles list. Make sure to add roles to this list from the Available roles list.

   - **Policy applies to all roles OTHER THAN those selected below**—To apply this policy to all users except for those who are mapped to the roles in the Selected roles list. Make sure to add roles to this list from the Available roles list.

8. Click **Save Changes**.

Related Documentation

- Configuring Connect Secure to Interoperate with IDP

Identifying and Managing Quarantined Users Manually

When the system quarantines a user based on an attack, you can display and manage the states by locating the user link in the **System > Status > Active Users** page.

- A small warning icon displayed in front of the username.

- The hyperlinked username.

- An enabled Quarantined option button on the specific user's page. If the user is not quarantined, the option button is disabled.

You can manage quarantined users from either the admin GUI or by logging in as a user with administrative rights on the local authentication server.

To manage quarantined users:

1. Identify quarantined users at **System > Status > Active Users**.

2. Locate the quarantined user from the **Authentication > Auth. Servers > System Local** on the admin GUI or from the **Admin Users** window on the local authentication server. You must be logged in to the local authentication server as an administrator user in order to see the Admin User option.

3. Click the username link. The user page opens, showing a number of options. See Figure 290.

Figure298: Managing Quarantined Users

4. Click **Disabled** to disallow a user from authenticating.

5. Click **Quarantined** to leave a user in a quarantined state. The Quarantined option is only enabled if the user is already quarantined.

> **i** NOTE: The system assigns quarantined users to the quarantined role, regardless of their login realm.

6. Click **Save Changes**.

7. To re-enable previously quarantined or disabled users, select **Authentication > Auth. Servers >** *Select Server* **> Users** and click the link for the given user.

> **i** NOTE: You can also disable users from this location.

8. Click **Enabled** to release the user from quarantine.

9. Click **Save Changes**.

All Sensor events are logged at System > Log/Monitoring > Sensors > Log.

Related Documentation

- Defining Automatic Response Sensor Event Policies

- Displaying Active Users

CHAPTER 43 Dashboard and Reports

Dashboard and Report Overview

A dashboard is an interface used to manage the Pulse Secure access management framework. It provides an integrated view of all devices and users accessing the network, their device profile information, authentication methods used to gain access, device posture compliance and so on.

A report is an element of a dashboard used to convey complex data in simplified formats. Pulse Secure access management framework collects log and configuration data from across your network, and it then aggregates the data into reports for you to view and analyze. It provides a standard set of predefined reports that you can use and customize to fit your needs. The reports are grouped into logical categories for information related to authentication, session traffic, device administration, configuration and administration, and troubleshooting.

You can use the system dashboard and reports to analyze system utilization.

Related Documentation

- Enabling the Dashboard
- Using the Dashboard

Enabling the Dashboard

You can use the admin console to enable or disable the dashboard.

To enable the dashboard.

1. Select **System > Status > Activity > Settings**.

2. Select **Enable Dashboard**.

   i   NOTE: The dashboard is enabled by default.

Figure 291 shows the Dashboard Settings for Policy Secure.

Figure299: Dashboard Settings – Policy Secure

Figure 292 shows the Dashboard Settings for Connect Secure.

Figure300: Dashboard Settings – Connect Secure

Figure 293 shows the available system reports through the dashboard for Connect Secure.

Figure301: Dashboard

Related Documentation

- Dashboard and Report Overview
- Using the Dashboard

Using the Dashboard

This topic describes the dashboard. It includes the following information:

- Dashboard Overview
- Displaying the Dashboard
- Selecting a Data Timeframe
- Refreshing Data
- Drilling Down to Detailed Reports

Dashboard Overview

- About the Dashboard
- About the Dashboard Database

About the Dashboard

The dashboard contains six default graphic reports focused on security, network activity, application activity, system monitoring, and compliance.

Table 188 describes the dashboard status bar for Policy Secure and Connect Secure.

Table189: Dashboard Status Bar

Metric	Description
Policy Secure	
Total Endpoints	The total number of unique endpoints over a time period. If an endpoint provides a unique identifier (mobile device ID, client ID, and so on), then it will be used to identify the endpoint. For a browser-based session where an identifier is not available, each session is considered a unique endpoint.
Active Endpoints	The total number of unique endpoints with active sessions.
Active Guests	The total number of active guest users. A guest user is defined as a user with an expiration date. This includes administrator-created and GUAM-created guest users. This does not include GUAM-created users without expiration dates.
Active MAC Auth Users	The total number of active MAC authentication users.
Connect Secure	
Total Users (Last 24 Hours)	The total number of unique users logged in over the past 1, 7, or 30 days. (The count is based on the chart time period setting. The pertinent time period, for example, 1, 7, or 30 days, is shown within brackets along with the number of users.)
Active Users	The number of unique users currently logged in.
Current SSL Sessions	The total number of current SSL sessions.
Auth Only Sessions	The total number of authentication-only user sessions.
ActiveSync Device Count	The total number of active synchronization devices.

describes the default dashboard charts.

Table190: Dashboard Charts

Dashboard Chart	Description
Authentication Success	The number of successful authentications over the selected time period (1, 7, or 30 days). The 7-day chart is a bar graph. The 1-day and 30-day charts are line graphs.
Authentication Failure	The number of failed authentications over the selected time period (1, 7, or 30 days). The 7-day chart is a bar graph. The 1-day and 30-day charts are line graphs.
Session OS Count	Pie chart showing the number of the sessions per operating system.
Top Roles	Pie chart showing the number of top user roles assigned during the selected time period.
Compliance Results	Pie chart showing Host Checker posture assessment results: Compliant, Not Compliant, Not Assessed, or Remediated. Compliance results are reported for all instances in which Host Checker is run. The posture assessment chart is clubbed with the compliance results chart for Policy Secure. To view the posture assessment chart select **Posture Assessment** from the dropdown list.
Posture Assessment	Pie chart showing Host Checker policy violations. Policy violations are reported only for instances in which Host Checker is run at initial sign in.
Auth Mechanism	Pie chart showing the number of sessions per authentication mechanism: 802.1x, Layer 3, MAC address. It applies to Policy Secure only.

About the Dashboard Database

The dashboard monitoring service collects and stores data in a database for 30 days. The total number of records stored in the database can be up to 300,000 records.

The dashboard database is created only after upgrading to Policy Secure version 5.0 or later or Connect Secure version 8.0 or later and enabling the dashboard option. Note that only new sessions are added to the database and changing the Timeframe filter or clicking refresh sends queries to the database. The data is collected only when the dashboard option is enabled.

Table 190 describes the different actions and their results.

Table191: Dashboard Database

Action	Description
Disable and then reenable the dashboard.	The data collection stops when your dashboard is disabled.
Restore the data from backup, snapshot, or import config.	The data is not exported and the data is retained during upgrades.

Displaying the Dashboard

To display the dashboard, select **System > Status > Activity > Dashboard**.

Figure 294 shows the dashboard for Policy Secure.

Figure302: Dashboard – Policy Secure

Figure 295 shows the dashboard for Connect Secure.

Figure303: Dashboard – Connect Secure

Selecting a Data Timeframe

To select a data timeframe:

1. Select **System > Status > Activity > Dashboard**.

2. Select one of the following periods from the Timeframe list box:

   - Last 24 Hours– (Default) Refers to the last 24 hours from the current hour.

   - Last 7 Days– Refers to current day and the previous last 6 days.

   - Last 30 Days– Refers to current day and the previous last 29 days.

     NOTE: Access records are kept for 30 days. Older records are removed and not included in dashboard charts and reports.

Figure 296 shows the dashboard for a timeframe of 30 days.

Figure304: Dashboard Showing a 30-Day Timeframe

Figure 297 shows the dashboard for a timeframe of 7 days.

Figure305: Dashboard Showing a 7-Day Timeframe

Refreshing Data

To refresh data:

1. Select **System > Status > Activity > Dashboard**.

2. Select one of the following refresh rates from the Refresh list box:

   - Disabled

   - 5 Minutes

   - 10 Minutes

   - 30 Minutes

   - 60 Minutes

Figure 298 shows the dashboard with a refresh rate of 5 minutes.

Figure306: Dashboard Showing a 5-Minute Refresh Rate

Drilling Down to Detailed Reports

To drill down to view detailed reports:

1. Select **System > Status > Activity > Dashboard**.

2. Click the search icon to display the corresponding tabular report with predefined search filters.

Figure 299 shows the detailed authentication report. The Authentication Results filter is set to Success.

Figure307: Detailed Authentication Report

Using the User Summary Report

This topic describes the user summary report. It includes the following information:

- About the User Summary Report
- Displaying the User Summary Report
- Applying Data Filters
- Sorting Records
- Drilling Down to the Single User Report
- Exporting User Summary Report

About the User Summary Report

The user summary report displays user statistics such as realm, username, last login time, last login IP, successful login, and so on for each user based on the user activity in the selected time range.

Displaying the User Summary Report

To display the user summary report , select **System > Reports > User Summary**.

Figure 300 shows the user summary report for Policy Secure.

Figure308: User Summary Report – Policy Secure

Figure 301 shows the user summary report for Connect Secure.

Figure309: User Summary Report – Connect Secure

Table 191 describes the columns on the user summary report.

Table192: User Summary Report Columns

Column	Description
Username	Specifies the name of the user.
Realm	Specifies the realm.
Last Login Time	Specifies the last time the user logged in.
Last Login IP	Specifies the last IP that the user logged in with.
Login Success	Specifies the number of successful logins.
Login Failure	Specifies the number of failed logins.
Compliant Sessions	Specifies the number of compliant sessions.
Non Compliant Sessions	Specifies the number of non compliant sessions.
Remediated Sessions	Specifies the number of remediated sessions.
Total Session Length	Specifies the total length of the sessions.
Average Session Length	Specifies the average length of the sessions.

Applying Data Filters

To apply a data filter:

1. Select **System > Reports > User Summary**.

2. Select one of the following periods from the Date Range list box:

   - Last 24 Hours– (Default) Refers to the last 24 hours from the current hour.

   - Last 7 Days– Refers to current day and the previous last 6 days.

   - Last 30 Days– Refers to current day and the previous last 29 days.

3. Enter search criteria in one or more of the following attribute columns:

   - Username

   - Realm

4. Click **Apply Filter**.

Figure 302 shows the user summary report filtered by username for Policy Secure. The user summary report filtered by username for Connect Secure is similar.

Figure310: Filter User Summary Report

Sorting Records

The data source determines the default sort order of the data rows in the report. Typically, data appears randomly, so sorting is an important task in creating a useful report. You can sort single data columns or multiple columns.

To sort the user summary report:

1. Select **System > Reports > User Summary**.

2. Select one of the following columns from the user summary report table, and click either the ascending or descending order icon.

   - Username

   - Realm

   - Last Login Time

   NOTE: The username column is sorted in ascending order by default.

   Figure 303 shows the user summary report sorted by username for Policy Secure. The user summary report sorted by username for Connect Secure is similar.

Figure311: Sort User Summary Report–Policy Secure

Drilling Down to the Single User Report

To drill down to a single user report:

1. Select **System > Reports > User Summary**.

2. Click the username to view the single user report.

Figure 304 shows the single user report displayed for Connect Secure. The single user report for Policy Secure is similar.

Figure312: Detailed Single User Report– Connect Secure

Exporting User Summary Report

To export device summary report:

1. Select **System > Reports > User Summary**.

2. Select a Download Report option.

   • CSV– Exports the report in CSV format.

   • Tab Delimited– Exports the report in tab-delimited format.

Figure 305 shows the export user summary report for Policy Secure. The export user summary report for Connect Secure is similar.

Figure313: Export User Summary Report–Policy Secure

Using the Device Summary Report

This topic describes the device summary report. It includes the following information:

• About the Device Summary Report

• Displaying the Device Summary Report

• Applying Data Filters

• Sorting Records

• Exporting Device Summary Report

About the Device Summary Report

The device summary report displays device information such as device detail, MAC address, last login time, last login IP, login successful, and so on for each user based on device activity in the selected time range.

Displaying the Device Summary Report

To display the device summary report:

1. Select **System > Reports > Device Summary**.

2. Select one of the following periods from the Date Range list box:

    - Last 24 Hours– (Default) Refers to the last 24 hours from the current hour.

    - Last 7 Days– Refers to current day and the previous last 6 days.

    - Last 30 Days– Refers to current day and the previous last 29 days.

3. Enter search criteria in one or more of the following columns:

    - Last Login Username

    - MAC Address

4. Click **Apply Filter**.

Figure 306 shows the device summary report for Policy Secure.

Figure314: Device Summary Report – Policy Secure

Figure 307 shows the device summary report for Connect Secure.

Figure315: Device Summary Report –Connect Secure

Table 192 describes the columns on the device summary report.

Table193: Device Summary Report Columns

Column	Description
Device ID	Specifies a unique identifier to identify the endpoint. Click the device ID icon to view a single device report.
MAC Address	Specifies the MAC address of the device. Click the MAC address to view a single device report.
Last Login Time	Specifies the last time the device was logged in.
Last Login IP	Specifies the last IP that the device logged in with.
Last Login Username	Specifies the username that the user logged in with.
Login Success	Specifies the number of successfully logins.
Login Failure	Specifies the number of failed logins.
Compliant Sessions	Specifies the number of compliant sessions.
Non-Compliant Sessions	Specifies the number of non-compliant sessions.
Remediated Sessions	Specifies the number of remediated sessions.
Total Session Length	Specifies the total session length.
Average Session Length	Specifies the average session length.

NOTE: If a device has more than one MAC address in a session, then the value appearing in the MAC Address column will be multiple instead of the actual MAC addresses. Note that the value multiple is not hyperlinked.

Applying Data Filters

To apply a data filter:

1. Select **System > Reports > Device Summary**.

2. Select one of the following periods from the Date Range list box:

   • Last 24 Hours– (Default) Refers to the last 24 hours from the current hour.

   • Last 7 Days– Refers to current day and the previous last 6 days.

- Last 30 Days– Refers to current day and the previous last 29 days.

3. Enter search criteria in one or more of the following columns:

- Last Login Username

- Mac Address

4. Click **Apply Filter**.

   Figure 308 shows the device summary report for Policy Secure. The device summary report for Connect Secure is similar.

Figure316: Filter Device Summary Report– Policy Secure

Sorting Records

The data source determines the default sort order of the data rows in the report. Typically, data appears randomly, so sorting is an important task in creating a useful report. You can sort single data columns or multiple columns.

To sort the device summary report:

1. Select **System > Reports > Device Summary**.

2. Select any one of the following column and click either the ascending or descending order icon.

- Last Login Time

- Last Login Username

   NOTE:
   - You can sort the column in either ascending order or descending order.

Figure 309 shows the device summary report sorted by last login time for Policy Secure. The device summary report sorted by login time for Connect Secure is similar.

Figure317: Sort Records in Device Summary Report

Exporting Device Summary Report

To export device summary report:

1. Select **System > Reports > Device Summary**.

2. Select a Download Report option.

   - CSV– Exports the report in CSV format.

   - Tab Delimited– Exports the report in tab-delimited format.

   shows the export device summary report for Policy Secure. The export device summary report for Connect Secure is similar.

Figure318: Export Device Summary Report

Using the Single Device Report

This topic describes the single device report. It includes the following information:

- About the Single Device Activities Report

- Displaying the Single Device Activities Report

- Applying Data Filters

- Sorting Records

- Exporting Single Device Activities Report

About the Single Device Activities Report

The single device activities report displays the device activity information such as username, realm, login time, logout time, device detail, MAC address, authentication mechanism, authentication result, compliance, IP address, role and so on for each device.

Displaying the Single Device Activities Report

To display the single device activities report , select **System > Reports > Single Device Activities**.

Figure 311 shows the single device activities report for Policy Secure.

Figure319: Single Device Activities Report – Policy Secure

Figure 312 shows the single device activities report for Connect Secure.

Figure320: Single Device Activities Report–Connect Secure

Table 193 describes the columns on the single device report.

Table194: Single Device Report Columns

Column	Description
Username	Specifies the name of the user.
Realm	Specifies the realm.
Login Time	Specifies the time the user logged in.
Logout Time	Specifies the time the user logged out.

Duration	Specifies the total duration of the user session.
Device ID	Specifies a unique identifier used to identify an end point. Click the device ID icon to view a single device report.
MAC Address	Specifies the MAC address of the device.
Auth Mechanism	Specifies the authentication mechanism: 802.1x, Layer 3, MAC address. It applies to Policy Secure only.
Auth Result	Specifies the authentication result.
Compliance	Specifies the Host Checker posture assessment results: Compliant, Not Compliant, Not Assessed, or Remediated.
IP Address	Specifies the IP that the user logged in with.
Role	Specifies the role of the user.
Device Detail	Displays the URL that is used for connecting to the MDM server.

Applying Data Filters

To apply a data filter:

1. Select **System > Reports > Single Device Activities**.

2. Select one of the following periods from the **Filter by: Date Range** list box:

   - Last 24 Hours– (Default) Refers to the last 24 hours from the current hour.

   - Last 7 Days– Refers to current day and the previous last 6 days.

   - Last 30 Days– Refers to current day and the previous last 29 days.

3. Enter search criteria in one or more of the following columns:

   - Compliance Results

   - MAC Address

   - Device ID

   - Authentication Mechanism. It applies only to Policy Secure.

4. Click **Apply Filter**.

   Figure 313 shows the single device activities report for Policy Secure. The single device activities report for Connect Secure is similar.

Figure321: Filter Single Device Activities Report

Sorting Records

The data source determines the default sort order of the data rows in the report. Typically, data appears randomly, so sorting is an important task in creating a useful report. You can sort single data column.

To sort the single device activities report:

1. Select **System > Reports > Single Device Activities**.

2. Select **Login Time** column and click either the ascending or descending order icon.

> NOTE:
> - You can sort the column in either ascending order or descending order.

Figure 314 shows the single device activities report sorted by last login time for Policy Secure. The single device activities report sorted by last login time for Connect Secure is similar.

Figure322: Sort Records in Single Device Activities Report

Exporting Single Device Activities Report

To export single device activities report:

1. Select **System > Reports > Single Device Activities**.

2. Select a Download Report option.

   - CSV– Exports the report in CSV format.

   - Tab Delimited– Exports the report in tab-delimited format.

   Figure 315 shows the single device activities report for Policy Secure. The single device activities report for Connect Secure is similar.

Figure323: Export Single Device Activities Report

Using the Authentication Report

This topic describes the authentication report. It includes the following information:

- About the Authentication Report
- Displaying the Authentication Report
- Applying Data Filters
- Sorting Records
- Exporting Authentication Report

About the Authentication Report

The authentication report displays the authentication result for each user based on the device activity in the selected time range.

Displaying the Authentication Report

To display the authentication report, select **System > Reports > Authentication**.

Figure 316 shows the authentication report for Policy Secure.

Figure324: Authentication Report – Policy Secure

Figure 317 shows the authentication report for Connect Secure.

Figure325: Authentication Report – Connect Secure

Table 194 describes the columns on the authentication report.

Table195: Authentication Report Columns

Column	Description
Username	Specifies the name of the user.
Realm	Specifies the realm.
Login Time	Specifies the time the user logged in.
Auth Mechanism	Specifies the authentication mechanism: 802.1x, Layer 3, MAC address. It applies only to Policy Secure.
Auth Result	Specifies the authentication result.
Failure Reason	Specifies the host checker failure reason.
Device ID	Specifies a unique identifier used to identify an end point. Click the device ID icon to view a single device report.
Role	Specifies the user role.
Device OS	Specifies the operating system of the device.

Applying Data Filters

To apply a data filter:

1. Select **System > Reports > Authentication**.

2. Select one of the following periods from the Filter by: Date Range list box:

   - Last 24 Hours– (Default) Refers to the last 24 hours from the current hour.

   - Last 7 Days– Refers to current day and the previous last 6 days.

   - Last 30 Days– Refers to current day and the previous last 29 days.

3. Enter search criteria in one or more of the following columns:

   - Authentication Results

   - Username

   - Realm

4. Click **Apply Filter**.

   Figure 318 shows the authentication report for Policy Secure. The authentication report for Connect Secure is similar.

Figure326: Filter Authentication Report

Sorting Records

The data source determines the default sort order of the data rows in the report. Typically, data appears randomly, so sorting is an important task in creating a useful report. You can sort single data column.

To sort the authentication report:

1. Select **System > Reports > Authentication**.

2. Select **Login Time** column and click either the ascending or descending order icon.

   Figure 319 shows the authentication report sorted by last login time for Policy Secure. The authentication report sorted by last login time for Connect Secure is similar.

Figure327: Sort Records in Authentication Report

Exporting Authentication Report

To export a authentication report:

1. Select **System > Reports > Authentication**.

2. Select a Download Report option.

   - CSV– Exports the report in CSV format.

   - Tab Delimited– Exports the report in tab-delimited format.

   Figure 320 shows the authentication report displayed for Policy Secure. The authentication report for Connect Secure is similar.

Figure328: Export Authentication Report

Using the Compliance Report

This topic describes the compliance report. It includes the following information:

- About the Compliance Report

- Displaying the Compliance Report

- Applying Data Filters

- Sorting Records

- Exporting Compliance Report

About the Compliance Report

The compliance report displays compliance status such as compliant, not compliant, remediated, not assessed information for each user based on the device activity in the selected time range.

Displaying the Compliance Report

To display the compliance report , select **System > Reports > Compliance**.

Figure 321 shows the compliance report for Policy Secure.

Figure329: Compliance Report – Policy Secure

Figure 322 shows the compliance report for Connect Secure.

Figure330: Compliance Report –Connect Secure

Table 195 describes the different columns on the compliance report.

Table196: Compliance Report Columns

Column	Description
Username	Specifies the name of the user.
Realm	Specifies the realm.

Device ID	Specifies a unique identifier used to identify an end point. Click the device ID icon to view a single device report.
MAC Address	Specifies the MAC address of the device.
Session Compliance	Specifies the Host Checker posture assessment results: Compliant, Not Compliant, Not Assessed, or Remediated.
Initial Host Check Time	Specifies the initial host check time.
Initial Host Check Details	Specifies the host check result.

The posture assessment chart is also a part of compliance report. It is displayed based on Initial Host Checker evaluation details (Login time).

Table 196 lists the type and the failure reasons for Host Checker.

Table197: Host Checker Failure Reasons– Posture Assessment Chart

Type	Failure Reason
Antivirus	Anti-virus not installed
	Anti-virus not running
	Anti-virus not up to date
	Anti-virus scan time check failed
Firewall	Firewall not installed
	Firewall not running
Antimalware	Anti-malware not installed
Antispyware	Anti-spyware not installed
	Anti-spyware not running
OS Checks	Unsupported OS
Port	Restricted ports open
	Required ports not open
Process	Detected restricted processes
	Required processes not detected
File	Detected restricted files
	Required files missing
Registry	Incorrect registry settings
NetBIOS	Detected restricted NetBIOS names
	Required NetBIOS names not found
MAC Address	Detected restricted MAC address
	Required MAC address not present
Machine Certificate	Certificate missing
Patch Assessment	Patches missing
Cache Cleaner	Cache cleaner failed
SVW	NA (Not considered for reporting)

SVW sub policy (.SVWActive)	Connected from non-SVW
Remote IMV	Remote IMV failure
EES	Enhanced Endpoint Security failed (no longer supported)
3rd party	NA (Not considered for reporting)
3rd party sub policy	3rd party sub policy failed
Rooting Detection	Detected rooted devices
Jail Breaking Detection	Detected jail broken devices
Mobile Security Suite (MSS)	Mobile security suite not active
3rd party NHC Check	Generic failure
Statement of Health	Generic failure
Connection Control	Generic failure

Applying Data Filters

To apply a data filter:

1. Select **System > Reports > Compliance**.

2. Select one of the following periods from the Filter by: Date Range list box:

   - Last 24 Hours– (Default) Refers to the last 24 hours from the current hour.

   - Last 7 Days– Refers to current day and the previous last 6 days.

   - Last 30 Days– Refers to current day and the previous last 29 days.

3. Enter search criteria in one or more of the following columns:

   - Compliance Results

   - Username

   - Realm

   - MAC Address

4. Click **Apply Filter**.

   Figure 323 shows the compliance report for Policy Secure. The compliance report for Connect Secure is similar.

Figure331: Filter Compliance Report

Sorting Records

The data source determines the default sort order of the data rows in the report. Typically, data appears randomly, so sorting is an important task in creating a useful report. You can sort single data column.

To sort the compliance:

1. Select **System > Reports > Compliance**.

2. Select **Initial Host Check Time** or **Username** column and click either the ascending or descending order icon.

   Figure 324 shows the compliance report sorted by last login time for Policy Secure. The compliance report sorted by last login time for Connect Secure is similar.

Figure332: Sort Records in Compliance Report

Exporting Compliance Report

To export a compliance report:

1. Select **System > Reports > Compliance**.

2. Select a Download Report option.

   - CSV– Exports the report in CSV format.

   - Tab Delimited– Exports the report in tab-delimited format.

   Figure 325 shows the export compliance report displayed for Policy Secure. The export compliance report for Connect Secure is similar.

Figure333: Export Compliance Report

Related Documentation

- Dashboard and Report Overview

- Using the Dashboard

- Using the User Summary Report

- *Using the Single-User Report*

- Using the Device Summary Report

Troubleshooting a Top Roles Chart from the Dashboard

Problem	**Description:**
	Environment:
	Symptoms: The same role for a selected time period appears multiple times in a top user roles report generated from the dashboard.
Diagnosis	The same role can appear multiple times when the role was deleted but created again using the same name.

CHAPTER 44 Pulse One Integration

- Overview
- Pulse One Configuration
- Pulse Workspace Handlers

Overview

PCS appliance can be integrated with the Pulse Workspace console server to auto-provision workspace based on user's group membership and to enable seamless active sync email access for mobile clients. Once this integration is in place, the mobile devices that are managed by Pulse Workspace will get seamless mail access from Enterprise mail server without requiring the users to configure their mail clients.

To configure Pulse Workspace command handlers to auto-provisioning workspace or to enable seamless active sync email access for mobile clients, do the following:
1. Register PCS with Pulse Workspace
2. Maintain Notification Channel
3. Renew Credentials
4. Configure User Role (For seamless Active Sync support)
5. Configure LDAP Authentication Servers to use for Group Lookup (For User's group membership based auto-provisioning)

Register PCS with Pulse Workspace

PCS has to be registered with Pulse Workspace before it can be used for seamless mail access for Pulse Workspace configured mobile devices. On successful registration, Pulse Workspace sends PCS the following information:

Table198: Registration Information

Registration Information	Description
Hawk Credentials	All communication from PCS to Pulse Workspace are authenticated using the HAWK. Pulse Workspace sends this information in the registration response. The response consists of: 1. Key 2. Key Identifier 3. Message Authentication Code Generation Algorithm
Device Identification Information	Each PCS device is uniquely identified in Pulse Workspace. This identification information is sent to PCS in the registration response to be used in all communications.
Notifications Channel URL	To receive any unsolicited notification from Pulse Workspace, PCS creates and maintains a websocket channel with Pulse Workspace. The endpoint URL on the Pulse Workspace for this channel is sent as part of the registration response.
Base API URL	On receiving any unsolicited notification on the websocket, PCS sends a REST request to Pulse Workspace to fetch additional information. The base URL for these REST APIs is sent by Pulse Workspace in the registration response.

Maintain Notification Channel

PCS creates a websocket channel with the Pulse Workspace server. Pulse Workspace sends notification to PCS over this channel. This channel is teared down by the Pulse Workspace once in 24 hours and PCS needs to reconnect to Pulse Workspace on this event. Also, when the HAWK credentials become invalid, the websocket channel is teared down.

PCS keeps the websocket channel up all the time and also takes corrective measures whenever there is a disruption on this channel.

Renew Credentials

HAWK credentials sent by Pulse Workspace are valid for 7 days. After this time, the credentials need to be renewed. When the credentials are in renew state, the notification channel will fail and any communication from PCS to Pulse Workspace cannot be authenticated. The existing credentials can only be used to request the new credentials.

HAWK credentials expire after 30 days. Once the credentials expire, PCS needs to be reconfigured and reregistered using a new registration code. This results into new device identification information and new HAWK credentials.

Configure User Role (For seamless Active Sync support)

Configure the User role that will be used for creating the device records on PCS for Pulse Workspace devices. On creation of a workspace, Pulse Workspace requests PCS to create a device record so that the mobile device which maps to that workspace can access email using PCS as activesync proxy. This requires PCS to know which role should be used for creating the device records. PCS administrator needs to configure this information using the admin UI.

Configure LDAP Authentication Servers to use for Group Lookup (For User's group membership based auto-provisioning)

Configure the LDAP Authentication server that will be used for handling group validation and user's group membership related requests on PCS for Pulse Workspace Server. PCS administrator needs to configure this information using the admin UI.

Pulse One Configuration

This section covers the configuration required on PCS to enable it to register with the Pulse Workspace console server.

Figure334: Pulse One Settings

Table199: Pulse One Configuration Details

Field	Description
Registration URL	This is the URL to which PCS sends the registration request. The format of the URL is https://<PWS API Host Name>/api/v1/register. The Pulse Workspace API Host name is displayed to the administrator when he/she creates an entry for this appliance on the Pulse Workspace console server.
Registration Code	This is the code that PCS sends to Pulse Workspace in the registration request. This code is generated and displayed to the administrator when he/she creates an entry for this appliance on the Pulse Workspace console server.
Credential Renegotiation Interval	This is the time in days after which PCS automatically does renegotiation of HAWK credentials with Pulse Workspace.
Credentials Exchange time	This is the time at which the last successful credential exchange took place.
Hashing Algorithm	This is the algorithm used for generating the MAC for HAWK authentication. Currently the only supported value is HS256 which is HMAC using SHA-256.
Client Device ID	This is the unique identification information of the PCS device on the Pulse Workspace server. This information is received in the registration response.
Notification URL	This is the URL at which the websocket endpoint is present at the Pulse Workspace server. This information is received in the registration response.
Registration Status	Reports current status of registration. Gray – not yet registered Yellow – registration in progress Green – registered successfully RED – registration failed/renew credentials/credentials expired
Notification Channel Status	Reports curent status of notification channel. Gray – not yet connected/connection not required Yellow – connection in progress Green – connected RED – connection failed
Save Changes	Saves the configuration and triggers registration, if required.
Clear configuration	Clears all the configuration and disconnects the notification channel.
Renegotiate credentials	Triggers renegotiation of credentials.

> **NOTE:**
> - Hawk is an HTTP authentication scheme providing a method for making authenticated HTTP requests with partial cryptographic verification of the request, covering the HTTP method, request URI, and host.

> **NOTE:**
> - To back up and restore Pulse One configuration, administrator should use the binary export/import of system configuration.

Pulse Workspace Handlers

This section covers the configuration of the command handlers that handle the messages received on the notification channel.

Figure335: Pulse Workspace Handlers

Active Sync Handler Configuration

This section covers the configuration of the activesync command handlers that create/delete the device records in PCS when Pulse Workspace sends a notification.

Table200: Active Sync Handler Configuration

Field	Description
Device Role	This is the role assigned to the device records created by PCS for the Pulse Workspace registered devices.
Clear Active sync Device Records	This option would delete all the device records pushed from Pulse Workspace Console Server.

> **NOTE:**
> - Administrator should ensure that secure email feature is enabled for this user role.
> - Use "Clear Active sync Device Records" option only if:
>   - This PCS is no longer the active sync provider for Pulse Workspace Server.

o To troubleshoot Device Record sync-up related issues, clear all Pulse Workspace Onboarded Device Records and recreate only the valid Device Records during next active sync Device Record sync-up. Device Record sync-up can happen if there is any new workspace created or existing workspace state is modified or due to periodic sync up initiated by the Pulse Workspace server for every one hour.

Group Lookup Handler Configuration

This section covers the configuration of group lookup command handlers that validate the group existence and also fetches the user's group membership from the configured backend LDAP server when Pulse Workspace sends a notification.

Table201: Group Lookup Handler Configuration

Field	Description
Available Auth Servers	All the configured LDAP Server will be listed under this.
Selected Auth Servers	Select the LDAP authentication server to handle the Group lookup requests.

NOTE:

- Only one authentication server per domain should be selected.
- This functionality is supported only with 'Active Directory' type LDAP server.
- To back up and restore Pulse One command handler configuration, administrator should use the binary export/import of user configuration

PART 6 Appendix

- Customizable Admin and End-User UIs

- SA4500 and SA6500 Series Appliances

- Using the SA4500 and SA6500 FIPS Appliances (Hardware FIPS)

- FIPS Level 1 Support (Software FIPS)

- Compression

- Localization

- Handheld Devices and PDAs

- Custom Expressions and System Variables

CHAPTER 45 Customizable Admin and End-User UIs

-

Customizable Admin and End-User UIs

The system enables you to customize a variety of elements in both the admin console and the end-user interface. This section contains information about which elements you can customize and where you can find the appropriate configuration options.se

You can customize the look and feel of the following user interface elements in the admin console:

- **Sign-in pages (default and custom)**—You can customize the page that administrators see when they sign into the admin console using settings in the Authentication > Signing In > Sign-in Pages page. Using settings in this page, you can create welcome messages, sign out messages and other instructions; control page headers; customize select error messages; and create a link to a custom help page within the default system sign-in page. Or, you can upload your own custom sign-in page.

- **UI look and feel**—You can customize the header, background color, and logo displayed in the admin console using settings in the Administrators > Admin Roles > Select Role > General > UI Options page. You can also use settings in this page to enable or disable the "fly out" hierarchical menus that appear when you mouse over one of the menus in the left panel of the admin console.

- **System utilization graphs**—You can choose which system utilization graphs to display on the opening page of the admin console using settings in the System > Status > Overview page. You can also use settings in this page to fine-tune the look and data within each of the graphs.

- **Show auto-allow options**—You can show or hide the auto-allow option from yourself or other administrators who create new bookmarks for roles using settings in the Maintenance > System > Options page.

- **User role views**—You can use customization options on the Users > User Roles page to quickly view the settings that are associated with a specific role or set of roles.

- **User realm views**—You can use customization options on the Users > User Realms page to quickly view the settings that are associated with a specific user realm or set of user realms.

- **Resource policy views**—You can limit which resource policies to display on any given resource policy page based on user roles. For instance, you can configure the Users > Resource Policies > Web page of the admin console to only display those resource policies

---

that are assigned to the "Sales" user role. You can customize these using settings in the Users > Resource Policies > Select Policy Type page of the admin console.

- **Web resource policy views**—You can limit which Web resource policy configuration pages to display using settings in Users > Resource Policies > Web > Policy Type of the admin console.

- **Administrator roles**—You can delegate select responsibilities to other administrators using settings in the Administrators > Admin Roles section of the admin console. In doing so, you can restrict the visibility of certain options and capabilities to those other administrators.

Customizable End-User Interface Elements Overview

The system enables you to customize the look and feel of the following elements in the end-user interface:

- **Sign-in pages (default and custom)**—You can customize the page that users see when they sign into the admin console using settings in the Authentication > Signing In > Sign-in Pages page. Using settings in this page, you can create welcome messages, sign out messages and other instructions; control page headers; customize select error messages; and create a link to a custom help page within the default system sign-in page. Or, you can upload your own custom sign-in page.

- **UI look and feel**—You can customize the header, background color, and logo displayed in the admin console using settings in the Users > User Roles > Select Role > General > UI Options page. You can also use settings in this page to specify the first page the users see after they sign in, the order in which to display bookmarks, the help system to display to users, and various toolbar settings.

- **Default messages and UI look and feel**—You can specify what the default look and feel should be for all user roles using settings in Users > User Roles > [Default Options] pages of the admin console. You can also use settings in these pages to define the default errors that users see when they try to access a blocked site, SSO fails, or SSL is disabled.

CHAPTER 46 SA4500 and SA6500 Series Appliances

SA4500 and SA6500

The SA4500 and SA6500 (SA 4500/6500) are next-generation appliances featuring a number of notable hardware features.

Standard Hardware

The SA4500 and SA6500 chassis features the following hardware components:

- Console port—You use the console port to initially set up the SA4500 and SA6500 before you fully integrate it as the secure gateway to your internal network. You can also use the console port to perform certain configuration and clustering tasks after the system begins operating as the secure gateway.

- Bonding ports—By default, on the SA6500 only, the device uses bonding of the multiple ports to provide failover protection. Bonding two ports automatically shifts traffic to the secondary port when the primary port fails.

  The SA6500 appliance bonds ports as follows:

  - Internal port = Port 0+Port 1

  - External port = Port 2+Port 3

Internal Ports   External Ports

The system indicates in a message on the System > Network > Overview page of the administrator admin console whether or not the failover functionality is enabled.

Bonding ports cannot span separate networks (multi-homed).

- Management port—The SA6500's management port:

  - Enables seamless integration into a dedicated Management Network.

  - Provides continuously available management access.

  - Enables you to perform management activities without impacting user traffic.

  - Allows you to separate administrative access from user access between the system and Enterprise devices on the internal network.

  You can configure the Management port information and advanced settings via the admin console, just as you would configure the internal port.

- SFP ports—4-port Small Form-factor Pluggable (SFP) ports are available as an optional feature for link redundancy to internal switches.

- Status LEDs—Three device status LEDs are located on the left-side of the front panel to display power, hard disk access and fault status.

- Ethernet Port LEDs—The Ethernet port LEDs show the status of each Ethernet port.

The appliance supports up to four node active/active clusters or 2 node active/passive.

SA6500 Field-Replaceable Units

The SA6500 chassis features three types of field-replaceable units (FRUs) that you can add or replace. The FRUs are "hot-swappable," meaning you do not have to first shut down the SA6500 before adding or replacing any of the FRUs. The SA4500 has a "cold-swappable" power supply.

For safety information, refer to the *Pulse Secure Products Safety Guide* available on the Pulse Secure Support site.

- **Hard disks**—The SA6500 ships with one hard disk, however, you can add an optional second hard disk to the SA6500 chassis to offer component redundancy and help minimize the system down time. When a second (redundant) hard disk is installed, it maintains an exact copy of the software image and configuration information on the working hard disk. Therefore, if the working hard disk fails, the redundant hard disk immediately assumes

responsibility for all system operations. This function is referred to as the Redundant Array of Independent Disks (RAID) mirroring process.

> **NOTE:** The SA6500 hard disk modules are hot-swappable. You must make sure that the device finishes booting and is operating correctly before removing, replacing, or upgrading a hard disk module. After you insert a new hard disk module, you must wait until the RAID mirroring process is completely finished—which takes approximately 40 minutes—before rebooting or turning off the device.

- **Power supplies**—The SA6500 ships with one AC power supply installed in the back of the chassis. You can add an optional second power supply to support redundancy and load-sharing features. In addition, if you need to replace one of the power supplies, you can "swap" the faulty power supply for a replacement while the optional second power supply assumes responsibility for the entire power load, thus avoiding a situation where you have to power off the device before replacing the removable unit.

- **Cooling fans**—The SA6500 ships with two cooling fans installed in the back of the chassis. If you need to replace one of the cooling fans, you can "swap" the faulty fan for a replacement during operation in a matter of moments. You can purchase additional cooling fans from your vendor when you order your SA6500, or you can purchase them in the future to replace faulty or failed cooling fans, as necessary, in the future.

Related Documentation

- Device Status LED Behavior

- Ethernet Port LED Behavior

- Replacing the Cooling Fans

- Replacing a Hard Drive

- Replacing IOC Modules

- Replacing a Power Supply

Device Status LED Behavior

Startup takes approximately one minute to complete. If you want to turn the device off and on again, we recommend you wait a few seconds between shutting it down and powering it back up.

There are three device status LEDs located on the left-side of the front panel:

- Power

- Hard disk access

- Fault

Table 197 lists the name, color, status, and description of each device status LED.

Table202: Device Status LEDs

Name	Color	State	Description
POWER	Green	Off	Device is not receiving power
		On Steady	Device is receiving power
HARD DISK ACCESS	Yellow	Off	Hard disk is idle
		Blinking	Hard disk is being accessed
FAULT	Red	Off	Device is operating normally
		Slow blinking	Power supply fault
		Fast blinking	Fan failure
		Solid	Thermal failure

Related Documentation

- SA4500 and SA6500

- Ethernet Port LED Behavior

- Replacing the Cooling Fans

- Replacing a Hard Drive

- Replacing IOC Modules

- Replacing a Power Supply

Ethernet Port LED Behavior

Table 198 show the status of each Ethernet port.

Table203: 4-Port Copper Gigabit Ethernet LEDs (available on IC4500 and IC6500)

LED	Color and State	Description
Link/Activity	Green	Link
	Blinking green	Activity
Link Speed	Off	10 Mbps
	Green	100 Mbps
	Yellow	1 Gbps

Related Documentation

- SA4500 and SA6500

- Device Status LED Behavior

- Replacing the Cooling Fans

- Replacing a Hard Drive

- Replacing IOC Modules

- Replacing a Power Supply

Replacing the Cooling Fans

The SA6500 ships with two cooling fans installed in the back of the chassis. If you need to replace one of the cooling fans, you can "hot-swap" the faulty fan for a replacement during operation in a matter of moments. You can purchase additional cooling fans from your authorized Juniper reseller, or you can purchase them in the future to replace faulty or failed cooling fans, as necessary.

To remove and install a cooling fan module:

1. To release the cooling fan module, do one of the following:

   - Press and slide the release trigger toward the center of the cooling fan module

   - Loosen the thumbscrews

2. Grasp the cooling fan module and carefully pull it out.

   ⚠️ CAUTION: Once you remove the cooling fan module, it is important that you replace it with a replacement cooling fan. The second fan is required for proper air flow across the chassis's internal components; it is not a redundant fan.

3. Line the a cooling fan module up with an empty cooling fan port on the back of the chassis.

4. Slowly slide the module into the chassis until it clicks into place.

5. If your cooling fan is equipped with thumb screws, tighten the screws.

Related Documentation

- SA4500 and SA6500

- Replacing a Hard Drive

- Replacing IOC Modules

- Replacing a Power Supply

Replacing a Hard Drive

The SA6500 ships with two standard hard drives to offer component redundancy and help minimize down time. The second (redundant) hard disk maintains an exact copy of the software image and configuration information on the working hard disk. Therefore, if the working hard disk fails, the redundant hard disk immediately assumes responsibility for all operations. This function is referred to as the Redundant Array of Independent Disks (RAID) mirroring process.

ⓘ NOTE: The hard disk modules are hot-swappable. Once a new hard disk module is inserted, you should wait until the RAID mirroring process has completed before rebooting or turning off the appliance.

To remove and install a hard drive:

1. On the hard drive module, press the blue handle release trigger in and to the right to release the insertion and removal handle.

2. Grasp the handle and pull the hard drive module straight out of the chassis.

   Once you have removed the hard drive module, be sure to replace it with a replacement hard drive.

3. With the insertion and removal handle on the hard drive module in the released/out position, line the hard drive module up with an empty hard drive port on the front of the chassis.

4. Carefully slide the hard drive module into the chassis until it is clicks into place.

Retract the handle by swinging it back across the face of the hard drive until it is completely flush with the face of the hard drive module.

Related Documentation

- SA4500 and SA6500

- Replacing the Cooling Fans

- Replacing IOC Modules

- Replacing a Power Supply

Replacing IOC Modules

This section contains information about removing and installing IOC Modules (IOMs) in the SA6500.

⚠️ CAUTION: Power off the device before removing or installing IOMs. IOMs are not hot-swappable.

Removing a Blank IOM Faceplate

To maintain proper airflow through the device, leave blank faceplates in place over slots that do not contain IOMs. Do not remove a blank faceplate unless you are installing an IOM in the empty slot.

To remove a blank faceplate:

1. Unplug the power cord.

2. Loosen the thumbscrews on each side of the faceplate.

3. Grasp the thumbscrews and pull to remove the faceplate.

Installing an IOM

1. Unplug the power cord.

2. Line the IOM up with an empty port on the front of the chassis.

3. Carefully slide the IOM in until it seats firmly in the device.

4. Tighten the screws on each side of the IOM faceplate.

5. Insert the appropriate cables into the cable connectors on the IOM.

6. If necessary, arrange the cables to prevent them from dislodging or developing stress points:

- Secure the cable so that it is not supporting its own weight as it hangs to the floor.

- Place excess cable out of the way in a neatly coiled loop.

- Use fasteners to maintain the shape of cable loops.

7. Insert the power cord into the AC power receptacle.

Removing an IOM

To remove an IOM:

1. Unplug the power cord.

2. Disconnect the cables from the IOM.

3. If necessary, arrange the cables to prevent them from dislodging or developing stress points.

4. Loosen the thumb screws on each side of the IOM faceplate.

5. Grasp the thumbscrews and pull to remove the IOM.

If you are not reinstalling an IOM into the empty slot, install a blank IOM faceplate over the empty slot to maintain proper airflow.

Related Documentation

- SA4500 and SA6500

- Replacing a Hard Drive

- Replacing a Hard Drive

- Replacing a Power Supply

Replacing a Power Supply

Removing and Installing an AC Power Supply
The Pulse Secure appliance ships with one AC power supply installed in the back of the chassis. You can add an optional second power supply to support redundancy and load-sharing features. In addition, if you need to replace one of the power supplies, you can "hot-swap" the faulty power supply for a replacement while the optional second power supply assumes responsibility for the entire power load, thus avoiding a situation where you have to power off the device before replacing the removable unit.

To remove and install an AC power supply module:

1. Press the release trigger in and to the right to release the module.

2. Grasp the insertion and removal handle and pull the power supply module straight out of the chassis.

   Once you have removed the supply module, be sure to replace it with a replacement power supply or the "dummy" power supply port cover installed in your chassis at the time of shipping.

3. Line the new power supply module up with an empty power supply port on the back of the chassis.

4. Slowly slide the power supply module into the chassis until it clicks into place.

Removing and Installing a DC Power Supply

To remove and install a DC power supply module:

1. Unplug the power cord.

2. Disconnect the DC supply wires from the lugs on the DC power supply.

3. Press the release trigger in and to the right to release the module.

4. Grasp the power supply module and pull it straight out of the chassis.

5. Slowly slide the new module into the chassis until it clicks into place.

6. Connect the DC supply wires to the module using the lugs. Be sure to attach the ground wire.

7. Attach the power cord

Related Documentation

- SA4500 and SA6500

- Replacing the Cooling Fans

- Replacing a Hard Drive

- Replacing IOC Modules

CHAPTER 47 Using the SA4500 and SA6500 FIPS Appliances (Hardware FIPS)

FIPS Overview

The Pulse Secure SA4500FIPS and SA6500FIPS is a standard SA4500 or SA6500 appliance equipped with a FIPS-compliant crypto card. The tamper-proof hardware security module installed on a FIPS system is certified to meet the FIPS 140-2 level 3 security benchmark.

The configuration process for FIPS administrators is almost exactly the same as for the non-FIPS administrators, requiring only minor configuration changes during the initialization, clustering, and certificate generation processes. In the few cases where administration tasks are different, this guide includes the appropriate instructions for both FIPS and non-FIPS administrators. For end users, a FIPS system is exactly the same as a non-FIPS system.

The FIPS-compliant crypto card is a host bus adapter card that combines IPsec and SSL cryptographic acceleration with Hardware Security Module (HSM) features. This combination of a dedicated HSM, advanced cryptographic security and secure key management meet the security and performance needs for any service.

This card has two main roles: a security officer and a user role. The FIPS-compliant crypto card replaces the need for administrator cards with the concept of a security officer who is responsible for key and password management. The security officer credential protects the keystore from being exported and imported onto another FIPS-compliant crypto card.

User roles perform cryptographic operations such as accessing keying material within the keystore as well as performing bulk encryption operations.

The security officer credentials, user credentials, and RSA private keys are stored in the HSM encrypted keystore located on the hard disk. You are prompted to provide these credentials whenever any operation requires them. Credentials are not automatically retrieved from the HSM keystore.

Keystores are stored on the disk and are encrypted with a master key. The master key is stored in the crypto card firmware and can be backed up by a security officer using a restore password. This restore password can then be used to restore the master key onto the same or different FIPS-compliant crypto cards allowing the keystore to be shared across a cluster, for example.

Related Documentation

- Name and Password Restrictions
- Initializing a Keystore
- Reinitializing the Keystore
- Joining a Cluster
- Importing Device Certificates
- Changing the Security Officer Password
- Changing the Web User Password
- Resetting the HSM Card In Case Of An Error
- Upgrading the HSM Firmware
- Binary Importing and Exporting of the Keystore

Name and Password Restrictions

Security officer names and usernames must adhere to the requirements shown in Table 199.

Table204: Security Officer Name and Username Requirements

Security Officer Name and Username Requirement	Description
Minimum Length	At least one character

Maximum Length	63 characters
Valid Characters	Alphanumeric, underscore (_), dash (-) and period (.)
First Character	Must be alphabetic

Passwords must be at least six characters and no more than 63 characters. Three characters must be alphabetic and one character must be non-alphabetic.

Related Documentation

- FIPS Overview

Initializing a Keystore

When the FIPS appliance is powered on from a factory-reset or when its configuration is reset, the serial console requires the initialization of a keystore and a self-signed device certificate. The steps for initialization are:

- During the boot process, the current release's HSM firmware is installed on the FIPS-compliant crypto card HSM.

- You are prompted to create a new keystore. As part of the new keystore creation, you must provide the following data:

  - The security officer name and password. Save these credentials as they are required for such tasks as creating new restore passwords and for changing the security officer password.

  - The keystore restore or HSM master key backup password. Every time you export the system configuration, save the current restore password for the archived keystore.

  - Web username and password for running cryptographic operations using keys stored in the HSM's keystore.

- The self-signed certificate creation proceeds as normal except that the HSM is used to generate a secure RSA private key which is stored in the HSM's database.

Related Documentation

- FIPS Overview

Reinitializing the Keystore

If there is a change in the security policy of the deployment that requires the creation of new RSA key pairs and corresponding certificates, you will need to reinitialize the keystore. You can reinitialize the keystore from either a standalone node or from a cluster.

To reinitialize the keystore from a standalone node:

1. Reboot the standalone node.

   During the boot process, you are prompted to re-initialize the keystore.

2. Press y to delete the current keystore and server certificates.

   (i) NOTE: If you do not press y within 10 seconds, the appliance will proceed to boot normally.

To reinitialize the keystore from a cluster:

1. Reboot a node within the cluster.

   During the boot process, you are prompted to re-initialize the keystore.

2. Press y to delete the current keystore and server certificates. A new keystore is initialized.

   > **NOTE:** If you do not press y within 10 seconds, the appliance will proceed to boot normally.

3. On the node that you rebooted, open the cluster status page in the admin console and wait for all nodes to exit from the "Transitioning" state.

4. For all other nodes in the cluster, connect to the serial console and enter 9 to select FIPS Options and then 1 to select Complete import of keystore and server certificates.

5. Enter the restore password when prompted.

Related Documentation

- FIPS Overview

Joining a Cluster

Joining a cluster involves using both the admin console and serial console. To join a cluster:

1. If you have not already done so, define and initialize a cluster

   If you are currently running stand alone appliances that you want to cluster, we recommend that before you create a cluster, you first configure system and user settings on one machine. After doing so, use the same machine to create the cluster. This machine joins the cluster as part of the creation process. When other devices join the cluster, this machine propagates its configuration to the new cluster member.

2. Before you can add an appliance to a cluster, you need to make its identity known to the cluster.

3. Join the appliance to the cluster through the admin console or through the serial console.

   - When joining a node to a cluster using the serial console, you are prompted for the cluster keystore's restore password. If the restore password fails, enter 9 to select FIPS Option and then enter 1 to select Complete import of keystore and server certificates.

   When a cluster is created on a node, the node's keystore becomes the cluster's keystore. Any node joining the cluster must import the cluster's keystore. You need the current keystore restore password to do this.

4. When you see the message confirming that the machine has joined the cluster, click the System > Clustering > Cluster Status tab in the admin console of any active cluster member.

5. When all nodes have exited from the "Transitioning" state, connect to the serial console of each node that has a non-CL license and enter 9 to select FIPS Options and then 1 to select Complete import of keystore and server certificates.

6. Enter the cluster keystore restore password.

Related Documentation

- FIPS Overview

Importing Device Certificates

To import a device certificate, generate a CSR from the appliance and then import its corresponding certificate after it is validated by a CA. Each CSR request generates a new RSA key pair.

> **ⓘ** NOTE: Device certificates without a CSR request from the appliance cannot be imported.

> **ⓘ** NOTE: The SA Series FIPS device is said to be in a disassociated state when the key store state in the cache and on disk are different. As a security measure, you cannot create or delete a CSR when the appliance is in a disassociated state. The options are grayed-out. To resolve a disassociated state, connect to the serial console and reload the FIPS keystore database (Option 9 > Sub-option 1).

Related Documentation

- FIPS Overview

Changing the Security Officer Password

Occasionally you may want to change the security officer password. In a cluster, you can perform this operation from any node. The new security officer password is updated to the other nodes automatically.

> **ⓘ** NOTE: Changing the security officer password restarts the web server.

To change the security officer password:

1. Connect to the serial console of the FIPS appliance you want to reset.

2. Enter 9 to select FIPS Option.

3. Enter 2 to select Change security officer password.

4. Enter the existing security officer password.

5. Enter the new password.

6. Re-enter the new password when prompted to confirm.

> **ⓘ** NOTE: The SA Series FIPS device is said to be in a disassociated state when the key store state in the cache and on disk are different. As a security measure, you cannot change the security officer password when the appliance is in a disassociated state. The option is disabled. To resolve a disassociated state, connect to the serial console and reload the FIPS keystore database (Option 9 > Sub-option 1).

Related Documentation

- FIPS Overview

Changing the Web User Password

The web username and password are used to securely store the RSA private keys in the HSM's encrypted database. These credentials are used by to carry out RSA operations. The keys will never be available for use outside the HSM. You can later change the web password but not the web username.

In a cluster, you can perform this operation from any node. The new password is updated to the other nodes automatically.

i   NOTE: Changing the web user password restarts the web server.

To change the web password:

1. Connect to the serial console of the FIPS appliance you want to reset.

2. Enter 9 to select FIPS Option.

3. Enter 3 to select Change web user password.

4. Enter the existing web user password.

5. Enter the new password.

i   NOTE: The SA Series FIPS device is said to be in a disassociated state when the key store state in the cache and on disk are different. As a security measure, you cannot change the web user password when the appliance is in a disassociated state. The option is disabled. To resolve a disassociated state, connect to the serial console and reload the FIPS keystore database (Option 9 > Sub-option 1).

Related Documentation

- FIPS Overview

Resetting the HSM Card In Case Of An Error

If the FIPS card LEDs indicates an error or fault, try resetting the HSM card prior to rebooting your appliance.

To reset the HSM card:

1. Connect to the serial console of the FIPS appliance you want to reset.

2. Enter 9 to select FIPS Option.

3. Enter 5 to select Reset the HSM.

4. Observe the LEDS on the FIPS card. If they do not eventually turn green, reboot your appliance.

Related Documentation

- FIPS Overview

Upgrading the HSM Firmware

Some system software upgrades may also require firmware updates. Typically, firmware upgrades occur during the boot process. After the system software updates, the serial console prompts you for the keystore restore password before upgrading the HSM's firmware. If you do not remember the password, you have the option of upgrading the firmware at a later date using the serial console. Note that the web server may not function properly if the firmware upgrade is required and is not updated.

To upgrade the firmware using the serial console:

1. Click **System > Clustering > Cluster Status** tab in the admin console and wait for the node to be in the "FIPS disassociated" state.

2. Open a serial console and enter 9 to select the FIPS option.

3. Enter 6 to select Load Firmware.

> **(i)** NOTE: The SA Series FIPS device is said to be in a disassociated state when the key store state in the cache and on disk are different. As a security measure, you cannot load firmware when the appliance is in a disassociated state. The option is disabled. To resolve a disassociated state, connect to the serial console and reload the FIPS keystore database (Option 9 > Sub-option 1).

Related Documentation

- FIPS Overview

Binary Importing and Exporting of the Keystore

Select **Maintenance > Import/Export** from the admin console to import and export the keystore. You can do this from a standalone node or from a node within a cluster. The keystore is exported as part of the system settings configuration file. Safely store the restore password associated with the archived keystore as you will need it for various FIPS operations. If you forget the restore password you can create a new one from the serial console and then re-export the configuration.

To import the keystore, select the **Import Key Store and Device Certificate(s)** check box and import your configuration. After the import process has completed, open a serial console for that FIPS appliance and enter 9 for FIPS Options and then 1 to select Complete import of keystore and server certificates. If the keystore is different from the one installed on the HSM you will be prompted for the keystore's restore password.

> **(i)** NOTE: If you reboot the FIPS appliance without performing the serial console step above, you are prompted to import the keystore during the boot process. Enter y to import the keystore. If you do not enter y within five seconds, the FIPS appliance continues to boot normally. If this occurs, perform the serial console step after the FIPS appliance completes its boot process.

If the FIPS appliance is in a cluster, go to each node within the cluster and perform the serial console step above to complete the keystore import process.

Related Documentation

- FIPS Overview

FIPS Device Status LED Behavior

Table 200 lists the three device status LEDs located on the FIPS card:

- S (Status)

- F (FIPS)

- I (INIT)

Table205: Status LED

LED	Color and State	Description
STATUS	Off	Bootstrap firmware is executing
	Blinking green	IDLE, OPERATIONAL, or FAILSAFE state
	Green	POST or DISABLED state (driver not attached)
	Blinking red	Error occurred during boot process
	Red	HALTED (fatal error) state or when a low-level hardware initialization failure occurred
FIPS	Off	Operating in non-FIPS mode
	Green	Operating in FIPS mode
	Blinking yellow	Zeroize jumper is present
INIT	Off	Board is not initialized
	Green	Board initialized by security officer
	Yellow	POST, DIAGNOSTIC or FAILSAFE (firmware not upgraded) state
	Blinking yellow	Running diagnostics

Related Documentation

- FIPS Overview

SA FIPS

FIPS, or Federal Information Processing Standards, are National Institute of Standards and Technology regulations for handling keys and encrypting data. Pulse Secure FIPS is a standard SA4000, SA4500, SA6000, or SA6500 NetScreen Instant Virtual Extranet equipped with a FIPS-certified cryptographic module. The tamper-proof hardware security module installed on a FIPS device is certified to meet the FIPS 140-2 level 3 security benchmark. The module handles private cryptographic key management and SSL handshakes, simultaneously, ensuring FIPS compliance and off-loading CPU-intensive public key infrastructure (PKI) tasks from the device to a dedicated module.

The configuration process for FIPS administrators is almost exactly the same as for the non-FIPS administrators, requiring only minor configuration changes during the initialization, clustering, and certificate generation processes. In the few cases where administration tasks

are different, this guide includes the appropriate instructions for both FIPS and non-FIPS administrators. For end users, FIPS is exactly the same as a standard system.

Related Documentation

- SA FIPS Execution

- Creating Administrator Cards

- Creating a New Security World

- Recovering an Archived Security World

SA FIPS Execution

When you first install a FIPS system, the system serial console walks you through the process of creating a security world through the serial console. A security world is a key management system used by FIPS consisting of the following elements:

- Cryptographic module—The cryptographic module (also sometimes called the hardware security module, or HSM) included with FIPS devices includes hardware and firmware installed directly on the appliance. A security world may contain a single cryptographic module (standard environment) or multiple modules (clustered environment). However, a single FIPS device is always equipped with a single cryptographic module.

- Security world key—A security world key is a unique Triple DES encrypted key that protects all other application keys within a security world. As required by the Federal Information Processing Standards, you cannot import this key into a security world—you must directly create it from a cryptographic module. In a clustered environment, all of the modules within the security world share the same security world key.

- Smart cards—A smart card is a removable key device that looks like a credit card. A smart card authenticates users, allowing them access to various data and processes controlled by the cryptographic hardware module. During the initialization process, you must insert one of your smart cards into the reader (built-in or external, depending upon which device model you own). As part of the initialization process, the smart card is transformed into an administrator card that allows the card holder access to the security world.

- Encrypted data—Encrypted host data in a FIPS environment includes keys and other data required to share information in a secure manner.

These elements interlock to create a comprehensive security world. When you start the appliance, it confirms that the security world is valid and that the cryptographic module is in operational mode before starting normal operations.

You can set the cryptographic module into operational mode using a hardware switch on the outside of the module. The switch's settings include:

- I—Initialization mode. Use this setting when initializing the cryptographic module with a new security world or when adding a module to an existing security world in a cluster. Note that once you set the switch to I and begin initialization, you must complete the process. Otherwise, your security world is only partially initialized, making it unusable.

- O—Operational mode. Use this setting to place the cryptographic module into operational mode after initialization. Note that you must set the switch to O before the module powers up

in order to alert the unit that you want to begin day-to-day processing. Otherwise, the module prompts you through the serial console to join the existing security world or initialize a new one.

- M—Maintenance mode. In future releases, this setting will be used to upgrade the firmware on the cryptographic module. (Not yet supported.)

Related Documentation

- SA FIPS
- Creating Administrator Cards
- Creating a New Security World
- Recovering an Archived Security World

Creating Administrator Cards

When you receive your FIPS product, you receive 6 smart cards as part of the package. A smart card is a removable key device that you must use in order to gain access to some of the critical data and processes controlled by the cryptographic module. The system first requires you to use one of your smart cards while initializing the cryptographic module through the serial console. During this process, the system creates a security world and transforms the smart card into an administrator card that gives the holder access only to that security world.

Once the module is initialized, you do not need the administrator card for normal system operations. However, you are required to use the administrator card whenever you want to add another FIPS device to a cluster, reinitialize a module with a new or different security world or replace administrator cards.

As a rule-of-thumb, any FIPS operation that you must execute through the serial console requires an administrator card.

> NOTE: Whenever you change your security world, you must determine how to handle your existing administrator cards. Your choices include:
>
> - Reset your existing administrator cards to the new security world.
> - Use administrator cards that are pre-initialized to the new security world and leave your existing administrator cards unchanged. Note that if you choose this option, however, you cannot use the old, unchanged cards to access the new security world.

Administrator Card Precautions

Since administrator cards are so critical to FIPS operations and the security of the keys within your security world, we strongly recommend that you take the following precautions:

- Create multiple administrator cards—You cannot replace an administrator card unless you have another valid card and the pass phrase for that card; the cryptographic module does not store administrator card recovery data. Therefore, we strongly recommend that you create at least one administrator card for standard administrative operations and another for backup purposes. Otherwise, you run the risk of losing your only administrator card and subsequently losing access to your security world and all the data it stores. You can only

create a set of administrator cards, all at once. You cannot add additional cards to an existing set.

- Store a backup administrator card in a secure location—Always keep your backup administrator card(s) in a secure location separate from the card you use for standard administrative operations to ensure that you do not lose all of your administrator cards to the same event (such as a fire or theft).

- Overwrite all remaining administrator cards if one gets lost—If you lose or damage an administrator card, immediately create a new security world and overwrite all remaining cards from the old security world. Otherwise, an attacker with an old administrator card may be able to access old host data stored on a backup tape or another host. With the old host data and an old card, the attacker may then be able to re-create your keys.

- Protect the administrator card's pass phrase—For maximum security, you should never write down your pass phrase, tell it to untrusted users, or use a pass phrase that is easy to guess. Protecting your pass phrase adds an extra level of security to your operations.

- Only use your administrator card with known, trusted sources—Always obtain smart cards from a trusted source, never insert a smart card into an untrusted smart card reader, and never insert untrusted smart cards into your smart reader.

Related Documentation

- SA FIPS

- Creating a New Security World

- Recovering an Archived Security World

Deploying a Cluster in a Connect Secure FIPS Environment

In addition to sharing state, user profile, user session, and monitoring state data, the members of a FIPS cluster also share security world data. All cluster members share the same private key and are accessible using the same administrator cards. Since changing a security world requires physical access to a cryptographic module, however, FIPS cluster members cannot share all of their data using the standard synchronization process. Instead, to create a FIPS cluster, you must:

- Create a cluster of FIPS devices through the admin console—As with a standard Connect Secure cluster, each cluster node in a FIPS cluster is initialized using system state data from the specified cluster member, overwriting all existing data on the node machine.

- Manually update the security world on each of the machines—After creating a cluster, you must initialize each cluster node with the specified member's security world using an administrator card that is pre-initialized to the security world and the serial console. Prior to joining a cluster, each node is in its own security world. As a consequence, after a node joins the cluster, the administrator card from the joining node will be invalid. Only the administrator card set from the cluster will be valid.

Similarly, if you want to modify an existing security world on a cluster, you must individually update each cluster member's cryptographic module using an administrator card and the serial console.

The basic process for creating a cluster follows these high-level steps:

1. Initialize one device from the serial console, creating administrator cards.

2. Create the cluster from this device's admin console.

3. Add nodes to the cluster from this device's admin console.

4. Reboot the joining node from the serial console.

5. When prompted, supply the cluster details, including the current node's IP address, netmask, and domain.

6. When prompted, insert an administrator card from the cluster's set of cards. The node's administrator card, if any, will become invalid as the node joins the security world of the cluster.

To initialize a FIPS cluster member's security world via the serial console:

1. Insert an administrator card that is pre-initialized with the active cluster member's security world into the smart card slot with the contacts facing up.

> NOTE: If you have already performed the procedures required to configure the FIPS appliance, as described in the Quick Start Guide, you might be able to skip this step.

2. Switch the cryptographic module's mode switch to I (initialization mode) if it is not already in that position.

3. Connect to the machine's serial console.

4. Cycle the power to reboot the machine and watch its serial console. After the system software starts, you will see a message that the machine is about to boot as a standalone device and to hit Tab for clustering options. Press the Tab key as soon as you see this option.

> NOTE: The interval to press the Tab key is five seconds. If the machine begins to boot in standalone mode, wait for it to finish and then reboot again.

5. Enter the number 2 to join the existing cluster or 1 to continue as a standalone device.

6. Enter the initialization information as prompted, including:

- Cluster name

- Cluster password

- IP address of a node in the cluster

- IP address of the node you are adding

- Netmask

- Gateway IP address

> NOTE: After you initialize members of a FIPS cluster with the same security world, you may disable and re-enable the cluster through the admin console. You are no longer required to use the serial console once the cluster members are all members of the same security world.

7. Select 1 to continue joining the cluster.

8. After the FIPS appliance initializes the card, switch the cryptographic module's mode switch to O (operational mode).

Related Documentation

- Using the Serial Port

Creating a New Security World

You cannot begin using a Connect Secure FIPS device until you create a security world on it. However, in some case you may need to overwrite that security world with a new one. For example, if you lose an administrator card, we recommend that you create a brand new security world to prevent an untrusted source from finding the card and accessing your security world. You may also need to create a new security world if you cannot remember your original administrator cards pass phrases.

In order to create a new security world, you must have physical access to:

- The cryptographic module(s) that belong to the security world.

- A smart card reader (if you use an older model device that does not contain a built-in card reader).

- One or more unformatted smart cards or administrator cards containing data that you can safely overwrite.

  i NOTE: Your old administrator cards will not work with the new security world until you reformat them with the new security world's data. Also note that once you set the switch to I and begin initialization, you must complete the process. Otherwise, your security world is only partially initialized, making it unusable.

  WARNING: You must obtain one or more new device certificates from your CA whenever you create a new security world.

Creating a Security World on a Standalone Device

To create a new security world on a standalone device:

1. Insert an unformatted smart card or an administrator card containing data that you can safely overwrite into the card slot with the card contacts facing up.

2. Set the mode switch on the cryptographic module to I (initialization mode).

3. Access the serial console and reboot the device. After the device reboots, you are prompted on the serial console with the following question: **Do you want to use the currently installed security world (y/n)?**

4. Perform one of the following:

   - If you want to create a new security world, then:

     1. Enter n and press Enter.

2. You are asked to confirm this choice with the prompt "Are you sure you want to delete your existing Security World (including server certificates) (y/n)?". If you choose to continue enter y and press Enter.

3. Enter the number of administrator cards you want to create and press Enter.

4. Enter y and press Enter to confirm the number of cards you want to create.

- If you want to use the currently installed security world, then:

  1. Enter y and press Enter.

  2. Proceed to the next numbered step in this procedure.

3. Reset the cryptographic module's mode switch to O (operational mode).

4. Add the common name and company name when prompted. The system uses the existing self-signed certificate temporarily.

5. Create a new device certificate that shares the new security world's private key.

> WARNING: You must obtain one or more new server certificates from your CA whenever you create a new security world.

Creating a Security World in a Clustered Environment

To create a new security world in a clustered environment:

1. Sign in to the admin console of a cluster node. To access a node's admin console, enter its internal IP address followed by "/admin" in a browser. For example:

   https://x.x.x.x/admin

2. On the System > Clustering > Status tab, select the check box for all nodes other than the current node in the Cluster Members column and then click **Disable**.

3. Initialize the cluster member with a security world. If this is the first node in the cluster, create a new security world.

4. Return to the node's System > Clustering > Status tab, select the check box next to disabled nodes in the Cluster Members column, and then click **Enable**.

5. Wait for all the cluster members to go into an "Enabled" state.

6. Set the mode switch on the cryptographic modules of cluster members that were earlier disabled to I (initialization mode).

7. Reboot each of these nodes from the serial console.

8. After a node joins the security world, reset its cryptographic module's mode switch to O (operational mode).

Replacing Administrator Cards

You can replace an administrator card by selecting the **Replace Administrator Card Set** option from the serial console. You cannot increase the number of administrator cards in an existing set. If you want to do this, you have to create a new security world which replaces all of the existing cards in a set and allow you to create a set with a larger or smaller number of cards.

NOTE: Replacing administrator cards restarts services on your standalone device or cluster.

If you need to replace administrator cards for a security world, you must have physical access to:

- A cryptographic module that belongs to the security world.

- A smart card reader (if you use an older model device that does not contain a built-in card reader).

- An administrator card that is pre-initialized with the security world.

- An unformatted smart card or administrator card containing data that you can safely overwrite.

- The same number of unformatted smart cards or administrator cards as in the original set containing data that you can safely overwrite.

NOTE: If you need to replace administrator cards, you must replace the same number of cards that you first initialized for the security world. You cannot replace a subset of the cards.

NOTE: If you require additional smart cards, please contact your reseller.

To replace all administrator cards or to create a larger number of cards for a security world:

1. Create a new security world.

2. Choose **Replace Administrator Card Set** from the list of configuration tasks.

3. Enter the pass phrase for the security world.

4. When prompted, insert an unformatted smart card or an administrator card whose data you can safely overwrite into the smart card reader with the contacts facing up.

5. Enter the additional initialization information for which you are prompted.

6. Repeat steps 4 and 5 for as many cards as you want to create.

7. Store at least one of the administrator cards in a secure location.

Related Documentation

- Recovering an Archived Security World

Recovering an Archived Security World

In rare cases, you may need to recover your system using an archived security world. The archived security world may be an older version of the security world that already exists on your system or the same version. In order to recover your system, you must have access to the system configuration file (by default, system.cfg) that holds the archived security world and its corresponding certificate.

In addition, if you are overwriting your security world with a different security world, you must have physical access to:

- All of the cryptographic modules that belong to the security world.

- A smart card reader (if you use an older model device that does not contain a built-in card reader).

- An administrator card that is pre-initialized with the security world and administrator passphrase that you want to import.

Importing a Security World Into a Standalone Device

To import an existing security world into a standalone device:

1. Import the system configuration file that contains the archived security world and its corresponding certificate into the device, and then initialize the security world if necessary. If the configuration file contains an archive of:

   - The same security world that was already present on the machine, no further configuration is required.

   - A different security world than was already present on the machine, you must initialize the new security world.

     > NOTE: If you import a configuration file containing a different security world, note that your existing administrator cards will not work with the imported security world until you reformat them with the new security world's data. Also note that once you set the switch to I and begin initialization, you must complete the process. Otherwise, your security world is only partially initialized, making it unusable.

2. Insert an administrator card that is pre-initialized with the imported security world into the smart card reader slot with the contacts facing up.

3. Set the mode switch on the cryptographic module to I (initialization mode).

4. Access the device's serial console and reboot the device.

5. Reset the cryptographic module's mode switch to O (operational mode) when prompted.

Importing a Security World Into a Cluster

To import an existing security world into a cluster:

1. Sign in to the admin console of a cluster node. To access a node's admin console, enter its internal IP address followed by "/admin" in a browser. For example:

   https://x.x.x.x/admin

2. On the System > Clustering > Status tab, select the check box for all nodes other than the current node in the Cluster Members column and then click **Disable**.

3. Import an archived security world in to the cluster member.

4. When the installation process completes, return to the node's System > Clustering > Status tab, select the check box next to the disabled nodes in the Cluster Members column, and then click **Enable**.

5. Wait for all the cluster members to go into the "Enabled" state.

6. Set the mode switch on the cryptographic modules of cluster members' that were earlier disabled to I (initialization mode).

7. Reboot each of these nodes from the serial console.

8. After a node joins the security world, reset its cryptographic module's mode switch to O (operational mode).

Related Documentation

- Creating a New Security World

CHAPTER 48 FIPS Level 1 Support (Software FIPS)

Understanding Pulse Secure FIPS Level 1 Support

What Is FIPS?

Federal Information Processing Standard (FIPS) are a set of standards that define security requirements for products that implement cryptographic modules used to secure sensitive but unclassified information. The most recent standards are defined in the FIPS Publication 140-2.

The FIPS documents define, among other things, security levels for computer and networking equipment. U.S. Federal Government departments, and other organizations, use FIPS to evaluate the cryptographic capabilities of the equipment they consider for purchase. Cryptographic modules are validated against separate areas of the FIPS specification. An overall certification level is assigned based on the minimum level achieved in any area. Although primarily aimed at environments requiring strict security, FIPS levels are increasingly enforced as qualifying criteria for all U.S. Federal Government contracts. Security-conscious private enterprises might also use FIPS levels as an equipment evaluation benchmark. FIPS levels also serve as a customer-neutral description of vendor requirements. Vendors can engineer security products to FIPS levels and extend the applicability and eligibility of these products across a broad customer base, thereby eliminating exhaustive and time-consuming customer-by-customer product qualification procedures.

What Is FIPS Level 1 Support?

Pulse Secure offers FIPS level 1 support for both Connect Secure and Policy Secure. Both services use a 140-2 level 1 certified cryptographic module to comply with FIPS. When FIPS level 1 support is enabled applications, such as browsers, accessing the web server must support Transport Layer Security (TLS), the latest version of Secure Socket Layer (SSL). If the platform features hardware acceleration, then for SSL processing SSL hardware acceleration is disabled as hardware acceleration does not comply with FIPS validation. Only FIPS approved algorithms are used when in FIPS level 1 support is enabled.

NOTE: You cannot run FIPS level 1 support on a hardware FIPS platform such as the SA6500 FIPS SSL VPN Appliance. For more information on the hardware FIPS

platform, see the *Connect Secure Administration Guide* and the *Policy Secure Administration Guide.*

For more information about the Pulse Secure Cryptography Module, see the security policy and the validation certificate. For a complete list of validated FIPS 140-1 and FIPS 140-2 cryptography modules, see http://csrc.nist.gov/groups/STM/cmvp/documents/140-1/140val-all.htm#2012.

Related Documentation

- FIPS Supported Platforms

- Enabling FIPS Level 1 Support

- *Pulse Secure FIPS Mode Overview for Pulse Policy Secure*

- *Pulse Secure FIPS Mode for Pulse Connect Secure Overview*

FIPS Supported Platforms

The following platforms support FIPS level 1:

- Pulse Secure Gateway MAG2600

- Pulse Secure Gateway MAG4610

- Pulse Secure Gateway MAG6610

- Pulse Secure Gateway MAG6611

- Pulse Secure Gateway MAG-SM160

- Pulse Secure Gateway MAG-SM360

- Connect Secure and Policy Secure virtual appliances

Related Documentation

- Understanding Pulse Secure FIPS Level 1 Support

- Enabling FIPS Level 1 Support

Enabling FIPS Level 1 Support

Once you enable FIPS level 1 support, your browser is restricted to specific custom cipher strengths. A list of supported ciphers is shown during the enabling process.

When you enable FIPS level 1 support, the following events occur on the system:

- The Web server restarts and turns on FIPS level 1 support. The Web server now allows only TLSv1.0, TLSv1.1 and TLSv1.2 protocols that include FIPS approved cryptographic algorithms which include Suite B cipher suites.

  > NOTE: Once FIPS level 1 support is enabled, new client sessions will use FIPS if the client supports FIPS. Existing client sessions may not be using FIPS. To ensure FIPS capable clients are in FIPS level 1 support, all client sessions should be terminated after the FIPS level 1 support is enabled. Administrators can use the **System > Status > Active Users** page to terminate client sessions.

- If the platform features hardware acceleration, when FIPS level 1 support is enabled SSL processing does not utilize the hardware acceleration. IPSec hardware acceleration is not affected.

- The **FIPS compliant Network Connect** option is enabled automatically on user roles. However, if this option is disabled manually for a role, and FIPS level 1 support is disabled and then reenabled, the option remains disabled for the role.

ℹ️   NOTE: You cannot enable FIPS level 1 support on a Pulse Secure FIPS platform, such as the SA6500 FIPS SSL VPN Appliance.

The following event logs are generated for FIPS level 1 support:

- SYS30966 when the web server turns FIPS level 1 support on.

- ADM30965 when the administrator turns FIPS level 1 support on or off.

- ERR30967 when the web server fails to turn on FIPS level 1 support.

To enable FIPS level 1 support:

1. Select **System > Configuration > Security > SSL Options**.

2. Under SSL FIPS Mode option, select **Turn on FIPS mode**. See Figure 326.

Figure336: Enabling FIPS Level 1 Support

Once you turn on FIPS level 1 support, the following changes are made:

- Under Allowed SSL and TLS Version, the **Accept only TLS** option is selected.

- Under Allowed Encryption Strength, the Custom SSL Cipher Selection is set. See Figure 326. Only FIPS approved algorithms are selected. All other options under this section are disabled. See "Supported Cipher Suites when FIPS Level 1 Support is Enabled and Disabled"

- Under Encryption Strength, the **Do not allow connections from browsers that only accept weaker ciphers** option is selected. You cannot disable this selection.

3. Click **Save Changes**.

A window shows the supported custom ciphers. See Figure 327 for an example. Your output may be different.

Figure337: Example of Allowed Encryption Strengths

4. Click **Change Allowed Encryption Strength**.

FIPS Level 1 support is now enabled on the device. If your browser does not support any of the listed ciphers, you will not be able to log in to the device.

When enabling FIPS mode on Policy Secure, only the following protocols are FIPS compliant:

- EAP-TTLS

- EAP-PEAP

- EAP-TLS

A warning is displayed if non-FIPS protocols are configured on an Policy Secure FIPS enabled device. However, these protocols are not disabled as they may be required for other use cases such as for the MAC authentication bridge.

Entries are made in the Events logs (see Figure 328) and Admin Access logs (see Figure 329) to show that FIPS level 1 support is enabled.

Figure338: Events Log Entries for FIPS Level 1

Figure339: Admin Access Logs for FIPS Level 1 Encryption Strength Changes

Related Documentation

- Understanding Pulse Secure FIPS Level 1 Support

- FIPS Supported Platforms

- Supported Cipher Suites when FIPS Level 1 Support is Enabled and Disabled

Turning Off FIPS Level 1 Support from the Serial Console

Problem

Description: If you have FIPS level 1 support enabled and your browser does not support the required cipher suites, you cannot access the device. If this happens to an administrator account, you can no longer administer or configure the system.

Solution

You can turn off FIPS level 1 support and reset the encryption strength from the device's serial console. After choosing that option, SSL options are reset to **Accept only SSL V3 and TLS** and to **Accept only 128-bit and greater encryption strength**.

Open a serial console to your device and select option **8. Turn off FIPS Mode and reset allowed encryption strength for SSL**.

Turning Off FIPS Level 1 and Resetting Encryption Strength From the Serial Console

```
Please choose the operation to perform:
    1. Network Settings and Tools
    2. Create admin username and password
    3. Display log/status
    4. System Operations
    5. Toggle password protection for the console (Off)
    6. Create a Super Admin session.
    7. System Maintenance
    8. Turn off FIPS Mode and reset allowed encryption strength for SSL
Choice: 8
```

ⓘ　　　　NOTE: Once you turn off FIPS level 1 support, option 8 is relabeled "Reset allowed encryption strength for SSL."

Related Documentation

- Enabling FIPS Level 1 Support

- Installing a Self-Signed Certificate From the Serial Console

Installing a Self-Signed Certificate From the Serial Console

Problem　　**Description:** An administrator can be locked out of the system if their browser does not support the certificate assigned to the network port. For example, if your system has an ECC certificate assigned to the internal port and your browser does not support ECC certificates you cannot log in to the device using the internal port.

Solution　　You can use the serial console to create and install a self-signed RSA certificate onto the internal port to allow access. Once you connect to the serial console, select option **4. System Operations** followed by Option **7. Install self-signed certificate**. It may take a few minutes for the 2048-bit key size self-signed certificate to be created and installed on your device. Once the certificate is installed, you can now log in to the device.

Creating and Installing an RSA Certificate From the Serial Console

```
Please choose the operation to perform:
    1. Network Settings and Tools
    2. Create admin username and password
    3. Display log/status
    4. System Operations
    5. Toggle password protection for the console (Off)
    6. Create a Super Admin session.
    7. System Maintenance
    8. Turn off FIPS Mode and reset allowed encryption strength for SSL
Choice: 4

Please choose the operation to perform:
    1. Reboot this Pulse Connect Secure
    2. Shutdown this Pulse Connect Secure
    3. Restart services at this Pulse Connect Secure
    4. Rollback this Pulse Connect Secure
    5. Factory reset this Pulse Connect Secure
    6. Clear all configuration data at this Pulse Connect Secure
    7. Install self-signed certificate
Choice: 7

Are you sure you want to install a newly-created RSA self-signed certificate on
the internal port? (y/n) y

Please provide information to create a self-signed Web server
digital certificate.
    Common name (example: secure.company.com): myname.mycompany.com
    Organization name (example: Company Inc.): MyCompany Inc.

Please enter some random characters to augment the system's
random key generator.  We recommend that you enter approximately
```

```
thirty characters.

Random text (hit enter when done):abcdef1234567


Creating self-signed digital certificate - this may take several minutes...
The self-signed digital certificate was successfully created.
```

Related Documentation

- Enabling FIPS Level 1 Support

Supported Cipher Suites when FIPS Level 1 Support is Enabled and Disabled

The tables in this topic list the cipher suites that are supported by the web server when the FIPS level 1 support is disabled and enabled.

Supported Cipher Suites When FIPS Level 1 Support is Disabled

When the FIPS level 1 support is disabled and when SSL hardware acceleration is not present or is disabled, the web server gives preference to cipher suites that use RC4 for bulk encryption. When SSL hardware acceleration is present and enabled, the web server gives preference to AES128, AES256 and 3DES over RC4, in that order.

Table206: Supported Cipher Suites With FIPS Level 1 Support Off, Hardware Acceleration Enabled and RSA Server Certificates In Use

Cipher Suite	Protocol
TLS_RSA_WITH_AES_128_CBC_SHA256	TLSv1.2
TLS_RSA_WITH_AES_128_CBC_SHA	SSLv3 and later
TLS_RSA_WITH_AES_256_CBC_SHA256	TLSv1.2
TLS_RSA_WITH_AES_256_CBC_SHA	SSLv3 and later
TLS_RSA_WITH_3DES_EDE_CBC_SHA	SSLv3 and later
TLS_RSA_WITH_RC4_128_SHA	SSLv3 and later
TLS_RSA_WITH_AES_128_GCM_SHA256	TLSv1.2
TLS_RSA_WITH_AES_256_GCM_SHA384	TLSv1.2
TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256	TLSv1.2
TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA	SSLv3 and later
TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA384	TLSv1.2
TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA	SSLv3 and later
TLS_ECDHE_RSA_WITH_3DES_EDE_CBC_SHA	SSLv3 and later
TLS_ECDHE_RSA_WITH_RC4_128_SHA	SSLv3 and later
TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256	TLSv1.2
TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384	TLSv1.2
TLS_RSA_WITH_RC4_128_MD5	SSLv3 and later

Table207: Supported Cipher Suites With FIPS Level 1 Support Off, Hardware Acceleration Enabled and ECC Server Certificates In Use

Cipher Suite	Protocol
TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256	TLSv1.2
TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA256	TLSv1.2
TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA	SSLv3 and later
TLS_ECDH_RSA_WITH_AES_128_GCM_SHA256	TLSv1.2
TLS_ECDH_ECDSA_WITH_AES_128_GCM_SHA256	TLSv1.2
TLS_ECDH_RSA_WITH_AES_128_CBC_SHA256	TLSv1.2
TLS_ECDH_ECDSA_WITH_AES_128_CBC_SHA256	TLSv1.2
TLS_ECDH_RSA_WITH_AES_128_CBC_SHA	SSLv3 and later
TLS_ECDH_ECDSA_WITH_AES_128_CBC_SHA	SSLv3 and later
TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384	TLSv1.2
TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA384	TLSv1.2
TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA	SSLv3 and later
TLS_ECDH_RSA_WITH_AES_256_GCM_SHA384	TLSv1.2
TLS_ECDH_ECDSA_WITH_AES_256_GCM_SHA384	TLSv1.2
TLS_ECDH_RSA_WITH_AES_256_CBC_SHA384	TLSv1.2
TLS_ECDH_ECDSA_WITH_AES_256_CBC_SHA384	TLSv1.2
TLS_ECDH_RSA_WITH_AES_256_CBC_SHA	SSLv3 and later
TLS_ECDH_ECDSA_WITH_AES_256_CBC_SHA	SSLv3 and later
TLS_ECDHE_ECDSA_WITH_3DES_EDE_CBC_SHA	SSLv3 and later
TLS_ECDH_RSA_WITH_3DES_EDE_CBC_SHA	SSLv3 and later
TLS_ECDH_ECDSA_WITH_3DES_EDE_CBC_SHA	SSLv3 and later
TLS_ECDHE_ECDSA_WITH_RC4_128_SHA	SSLv3 and later
TLS_ECDH_RSA_WITH_RC4_128_SHA	SSLv3 and later
TLS_ECDH_ECDSA_WITH_RC4_128_SHA	SSLv3 and later

Table208: Supported Cipher Suites With FIPS Level 1 Support Off, Hardware Acceleration Disabled and RSA Server Certificates In Use

Cipher Suite	Protocol
TLS_RSA_WITH_RC4_128_SHA	SSLv3 and later
TLS_RSA_WITH_AES_128_CBC_SHA256	TLSv1.2
TLS_RSA_WITH_AES_128_CBC_SHA	SSLv3 and later
TLS_RSA_WITH_AES_256_CBC_SHA256	TLSv1.2
TLS_RSA_WITH_AES_256_CBC_SHA	SSLv3 and later
TLS_RSA_WITH_3DES_EDE_CBC_SHA	SSLv3 and later
TLS_RSA_WITH_AES_128_GCM_SHA256	TLSv1.2
TLS_RSA_WITH_AES_256_GCM_SHA384	TLSv1.2
TLS_ECDHE_RSA_WITH_RC4_128_SHA	SSLv3 and later

TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256	TLSv1.2
TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA	SSLv3 and later
TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA384	TLSv1.2
TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA	SSLv3 and later
TLS_ECDHE_RSA_WITH_3DES_EDE_CBC_SHA	SSLv3 and later
TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256	TLSv1.2
TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384	TLSv1.2
TLS_RSA_WITH_RC4_128_MD5	SSLv3 and later

Table209: Supported Cipher Suites With FIPS Level 1 Support Off, Hardware Acceleration Disabled and ECC Server Certificates In Use

Cipher Suite	Protocol
TLS_ECDHE_ECDSA_WITH_RC4_128_SHA	SSLv3 and later
TLS_ECDH_RSA_WITH_RC4_128_SHA	SSLv3 and later
TLS_ECDH_ECDSA_WITH_RC4_128_SHA	SSLv3 and later
TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256	TLSv1.2
TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA256	TLSv1.2
TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA	SSLv3 and later
TLS_ECDH_RSA_WITH_AES_128_GCM_SHA256	TLSv1.2
TLS_ECDH_ECDSA_WITH_AES_128_GCM_SHA256	TLSv1.2
TLS_ECDH_RSA_WITH_AES_128_CBC_SHA256	TLSv1.2
TLS_ECDH_ECDSA_WITH_AES_128_CBC_SHA256	TLSv1.2
TLS_ECDH_RSA_WITH_AES_128_CBC_SHA	SSLv3 and later
TLS_ECDH_ECDSA_WITH_AES_128_CBC_SHA	SSLv3 and later
TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384	TLSv1.2
TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA384	TLSv1.2
TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA	SSLv3 and later
TLS_ECDH_RSA_WITH_AES_256_GCM_SHA384	TLSv1.2
TLS_ECDH_ECDSA_WITH_AES_256_GCM_SHA384	TLSv1.2
TLS_ECDH_RSA_WITH_AES_256_CBC_SHA384	TLSv1.2
TLS_ECDH_ECDSA_WITH_AES_256_CBC_SHA384	TLSv1.2
TLS_ECDH_RSA_WITH_AES_256_CBC_SHA	SSLv3 and later
TLS_ECDH_ECDSA_WITH_AES_256_CBC_SHA	SSLv3 and later
TLS_ECDHE_ECDSA_WITH_3DES_EDE_CBC_SHA	SSLv3 and later
TLS_ECDH_RSA_WITH_3DES_EDE_CBC_SHA	SSLv3 and later
TLS_ECDH_ECDSA_WITH_3DES_EDE_CBC_SHA	SSLv3 and later

Supported Cipher Suites When FIPS Level 1 Support is Enabled

When FIPS level 1 support is enabled, only TLSv1.0, v1.1, v1.2 and AES256, 3DES and AES128 are allowed. The order of the cipher suites is not dependent on the SSL hardware acceleration module since hardware acceleration is not used when FIPS level 1 support is enabled.

When FIPS level 1 support is enabled, the following settings are automatically configured:

- In the SSL Options window:

  - Under Allowed SSL and TLS Version, the **Accept only TLS** option is selected. All other options under this section are disabled.

  - Under Allowed Encryption Strength, the **Custom SSL Cipher Selection** option is selected. Only FIPS approved algorithms are selected. All other options under this section are disabled.

  - Under Encryption Strength Option, the **Do not allow connections from browsers that only accept weaker ciphers** option is selected.

- SSL hardware acceleration is disabled. IPsec hardware acceleration is not affected by the FIPS level 1 support being enabled.

In Table 205, the first four cipher suites are given preference due to the requirements in RFC 6460. The first two cipher suites meeting the requirement for Suite B Profile for TLS 1.2. The next two meeting the requirement for Suite B Transitional Profile for TLS 1.0 and 1.1.

Table210: Supported Cipher Suites With FIPS Level 1 Support On and ECC Server Certificates In Use

Cipher Suite	Protocol
TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256	TLSv1.2
TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384	TLSv1.2
TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA	TLSv1.0 and later
TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA	TLSv1.0 and later
TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA384	TLSv1.2
TLS_ECDH_RSA_WITH_AES_256_GCM_SHA384	TLSv1.2
TLS_ECDH_ECDSA_WITH_AES_256_GCM_SHA384	TLSv1.2
TLS_ECDH_RSA_WITH_AES_256_CBC_SHA384	TLSv1.2
TLS_ECDH_ECDSA_WITH_AES_256_CBC_SHA384	TLSv1.2
TLS_ECDH_RSA_WITH_AES_256_CBC_SHA	TLSv1.0 and later
TLS_ECDH_ECDSA_WITH_AES_256_CBC_SHA	TLSv1.0 and later
TLS_ECDHE_ECDSA_WITH_3DES_EDE_CBC_SHA	TLSv1.0 and later
TLS_ECDH_RSA_WITH_3DES_EDE_CBC_SHA	TLSv1.0 and later
TLS_ECDH_ECDSA_WITH_3DES_EDE_CBC_SHA	TLSv1.0 and later
TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA256	TLSv1.2
TLS_ECDH_RSA_WITH_AES_128_GCM_SHA256	TLSv1.2
TLS_ECDH_ECDSA_WITH_AES_128_GCM_SHA256	TLSv1.2
TLS_ECDH_RSA_WITH_AES_128_CBC_SHA256	TLSv1.2

TLS_ECDH_ECDSA_WITH_AES_128_CBC_SHA256	TLSv1.2
TLS_ECDH_RSA_WITH_AES_128_CBC_SHA	TLSv1.0 and later
TLS_ECDH_ECDSA_WITH_AES_128_CBC_SHA	TLSv1.0 and later

Table211: Supported Cipher Suites With FIPS Level 1 Support On and RSA Server Certificates In Use

Cipher Suite	Protocol
TLS_RSA_WITH_AES_256_CBC_SHA256	TLSv1.2
TLS_RSA_WITH_AES_256_CBC_SHA	TLSv1.0 and later
TLS_RSA_WITH_3DES_EDE_CBC_SHA	TLSv1.0 and later
TLS_RSA_WITH_AES_128_CBC_SHA256	TLSv1.2
TLS_RSA_WITH_AES_128_CBC_SHA	TLSv1.0 and later
TLS_RSA_WITH_AES_256_GCM_SHA384	TLSv1.2
TLS_RSA_WITH_AES_128_GCM_SHA256	TLSv1.2
TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA384	TLSv1.2
TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA	TLSv1.0 and later
TLS_ECDHE_RSA_WITH_3DES_EDE_CBC_SHA	TLSv1.0 and later
TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256	TLSv1.2
TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA	TLSv1.0 and later
TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384	TLSv1.2
TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256	TLSv1.2

Related Documentation

- Enabling FIPS Level 1 Support

- Example: Assigning an ECC P-256 Certificate to an External Virtual Port and Giving Preference to Suite B Ciphers

CHAPTER 49 Compression

- About Compression
- Enabling System-Level Compression

About Compression

The system improves performance by compressing common types of Web and file data such as HTML files, Word documents, and images.

The system determines whether it should compress the data accessed by users by using the following process:

1. The system verifies that the accessed data is a compressible type. Compressing many common data types such as such as HTML files, and Word documents is supported.

2. If the user is accessing Web data, the system verifies that the user's browser supports compression of the selected data type.

   The system determines compression supportability based on the browser's user-agent and the accept-encoding header. It supports the compression of all of the standard Web data types if it determines that the user-agent is compatible with Mozilla 5, Internet Explorer 5, or Internet Explorer 6. The system supports only compressing HTML data, however, if it determines that the browser's user-agent is only compatible with Mozilla 4.

3. The system verifies that compression is enabled at the system level. You can enable system-level compression through the Maintenance > System > Options page of the admin console.

4. The system verifies that compression resource policies or autopolicies are enabled for the selected data type and comes with resource policies that compress data. You may enable these policies or create your own through the following pages of the admin console:

   - Users > Resource Policies > Web > Compression.
   - Users > Resource Policies > Files > Compression.

   You may also create resource profile compression autopolices through the Users > Resource Profiles > Web > Web Applications/Pages page of the admin console.

If all of these conditions are met, the system runs the appropriate resource policy either compresses or does not compress the data accessed by the user based on the configured action.

If all of these conditions are not met, the system does not run the appropriate resource policy and no resource policy items appear in the log files.

The system comes pre-equipped with three resource policies that compress Web and file data. If you are upgrading from a pre-4.2 version of the system software and you previously had

compression enabled, these policies are enabled. Otherwise, if you previously had compression disabled, these policies are disabled.

The Web and file resource policies created during the upgrade process specify that the system should compress all supported types of Web and File data, including types that were not compressed by previous versions of the appliance. All data types that were not compressed by previous product versions are marked with an asterisk (*) in the supported data types list below.

The system supports compressing the following types of Web and file data:

- text/plain (.txt)

- text/ascii (.txt)*

- text/html (.html, .htm)

- text/css (.css)

- text/rtf (.rtf)

- text/javascript (.js)

- text/xml (.xml)*

- application/x-javascript (.js)

- application/msword (.doc)

- application/ms-word (.doc)*

- application/vnd.ms-word (.doc)*

- application/msexcel (.xls)*

- application/ms-excel (.xls)*

- application/x-excel (.xls)*

- application/vnd.ms-excel (.xls)*

- application/ms-powerpoint (.ppt)*

- application/vnd.ms-powerpoint (.ppt)*

> NOTE: The data types denoted by an asterisk * were not compressed by pre-4.2 versions of the system software.
>
> Also note that the system does not compress files that you upload—only files that you download from the system.

Additionally, the system supports compressing the following types of files:

- text/html (.html, .htm)

- application/x-javascript (.js)

- text/javascript (.js)

- text/css (.css)

- application/perl (.cgi)

Related Documentation

- Enabling System-Level Compression

Enabling System-Level Compression

To enable system-level compression:

1. Select **Maintenance > System > Options**.

2. Select the **Enable gzip compression** check box to reduce the amount of data sent to browsers that support HTTP compression. Note that after you enable this option, you must also configure Web and file resource policies specifying which types of data the system should compress.

3. Click **Save Changes**.

Related Documentation

- About Compression

CHAPTER 50 Localization

About Multi-Language Support for Connect Secure

The system provides multi-language support for file encoding, end-user interface display, and customized sign-in and system pages. It supports the following languages:

- English (US)
- Chinese (Simplified)
- Chinese (Traditional)
- French
- German
- Japanese
- Korean
- Spanish

Related Documentation

Encoding Files for Multi-Language Support

Character encoding is a mapping of characters and symbols used in written language into a binary format used by computers. Character encoding affects how you store and transmit data. The encoding option in Users > Resource Policies > Files > Encoding allows you to specify the encoding to use when communicating with Windows and NFS file shares. The encoding option does not affect the end-user language environment.

To specify the internationalization encoding for system traffic:

1. In the admin console, choose **Users > Resource Policies > Files > Encoding**.

2. Select the appropriate option:

   - Western European (ISO-8859-1) (default) (Includes English, French, German, Spanish)

- Simplified Chinese (CP936)

- Simplified Chinese (GB2312)

- Traditional Chinese (CP950)

- Traditional Chinese (Big5)

- Japanese (Shift-JIS)

- Korean

3. Click **Save Changes**.

Related Documentation

- About Multi-Language Support for Connect Secure

- Localizing the User Interface

- Localizing Custom Sign-In and System Pages

Localizing the User Interface

The system provides a means to display the end-user interface in one of the supported languages. Combining this feature with (custom) sign-in and system pages and a localized operating system provides a fully localized user experience.

When you specify a language, the system displays the user interface, including all menu items, dialogs generated by the system, and the help file in the chosen language for all users regardless of which realm they sign in to.

To configure localization options:

1. In the admin console, choose **Maintenance > System > Options**.

2. Use the End-user Localization drop-down list to specify the language in which to display the end-user interface (optional). If you do not specify a language, the end-user interface displays based on the settings of the browser.

3. Click **Save Changes**.

Related Documentation

- About Multi-Language Support for Connect Secure

- Encoding Files for Multi-Language Support

- Localizing Custom Sign-In and System Pages

Localizing Custom Sign-In and System Pages

The system provides several zip files that contain different sets of sample template files for various pages that may appear during the sign-in process. Use these template files along with the template toolkit language to create localized custom sign-in and system pages for your end users.

Editing the default sign-in page using text in the language of your choice is a quick way to provide your users with a localized sign-in page.

Related Documentation

- *Custom Sign-In Pages Solutions Guide*

- About Multi-Language Support for Connect Secure

- Encoding Files for Multi-Language Support

- Localizing the User Interface

CHAPTER 51 Handheld Devices and PDAs

Handheld Devices and PDAs

In addition to allowing users to access the system from standard workstations and kiosks, the system also allows end users access from connected PDAs, handhelds and smart phones such as i-mode and Pocket PC. When a user connects from a PDA or handheld device, the system determines which pages and functionality to display based on settings in the System > Configuration > Client Types page of the admin console. By default, settings in this page specify that when accessing the system using a(n):

- **i-mode device**—The system displays compact HTML (cHMTL) pages without tables, images, JavaScript, Java, or frames to the user. Depending on which features you enable through the admin console, the end user may browse the Web, link to Web bookmarks, single sign-on to other applications, and edit their preferences (including clearing their cache and editing their system/LDAP password). The system allows i-mode users to access supported features using access keys on their phone's keypad as well as through standard browse-and-select navigation.

- **Pocket PC device**—The system displays mobile HTML pages with tables, images, JavaScript and frames, but does not process Java. Depending on which features you enable through the admin console, the end user may access Mobile Notes and OWA e-mail applications, browse the Web, link to Web bookmarks, single sign-on to other applications, and edit their preferences (including clearing their cache and editing their system/LDAP password).

PDA and handheld users cannot access the admin console or most of the system's advanced options, including file browsing, VPN Tunneling, Pulse Collaboration, Telnet/SSH, Email Client, Host Checker, and Cache Cleaner, since PDA and handheld devices do not generally support the ActiveX, Java, or JavaScript controls on which these features depend.

Also note that i-mode users cannot access cookie-based options, including session cookies and SiteMinder authentication and authorization, since most i-mode browsers do not support HTTP cookies. The system rewrites hyperlinks to include the session ID in the URL instead of using cookies. The system reads the session ID when the user accesses the URL.

> **ⓘ** NOTE: In order to improve the response time, the following icons are not displayed when accessing the home page: help, sign out, open bookmark in new page, and WSAM.

Related Documentation

- Task Summary: Configuring Connect Secure for PDAs and Handhelds

- Defining Client Types

- Enabling WSAM on PDAs

- Enabling ActiveSync For Handheld Devices

Task Summary: Configuring Connect Secure for PDAs and Handhelds

To properly configure the system to work with PDAs and handheld devices, you must:

1. **Enable access at the system level**—If you want to support browsers other than the defaults provided with the system, you must enter the user agent strings of the PDA and handheld operating systems that you want to support in the System > Configuration > Client Types tab. For a complete list of supported PDA and handheld browsers, see the Supported Platforms document posted on the Support web site.

2. **Evaluate your user roles and resource policies**—Depending on which Connect Secure features you have enabled, you may need to either modify your existing roles and resource policies for PDA and handheld users or create new ones. Note that:

   - Mobile device users cannot access roles or policies that require Host Checker or Cache Cleaner since handheld devices do not generally support the ActiveX, Java, or JavaScript controls on which these features depend. You can disable these options through the following tabs:

     - Users > User Roles > *Role* > General > Restrictions

     - Resource Policies > Web > Access > Web ACL> *Policy* > Detailed Rules

   - Mobile device users may have trouble reading long role names on their small screens. If you require users to pick from a list of roles when they sign in, you may want to shorten role names in the Users > User Roles > Role > General > Overview tab.

   - Mobile device users may have trouble reading long bookmark names on their small screens. You can edit Web bookmarks in the following tabs:

     - Users > Resource Profiles > Web Application Resource Profiles > *Profile* > Bookmarks

     - Users > User Roles > *Role* > Web > Bookmarks

     - Resource Policies > Web > Access > Web ACL> *Policy* > General

   - Although advanced features such as file browsing are not supported for PDAs and handhelds, you do not need to disable them in the roles and resource policies used by mobile device users. The system simply does not display these options to mobile device users.

3. **Evaluate your authentication and authorization servers**—The system supports all of the same authentication and authorization servers for PDA and handheld users as standard

users, except the eTrust SiteMinder policy server. SiteMinder is dependent on cookies, which are not supported with i-mode browsers.

4. **Evaluate your realms**—Depending on which system features you have enabled, you may need to either modify your existing realms for PDA and handheld users or create new ones. Note that:

- Mobile device users cannot access the system when they try to sign into a realm that requires Host Checker or Cache Cleaner since handheld devices do not generally support the ActiveX, Java, or JavaScript controls on which these features depend. You can disable these options through sub-tabs in the System > Configuration > Security page.

- Mobile device users cannot authenticate against an eTrust SiteMinder server. You can choose a different authentication server for the realm in the Users > User Realms > *Realm* > General tab.

- Mobile device users may have trouble reading long realm names on their small screens. If you require users to pick from a list of realms when they sign in, you may want to shorten realm names in the Users > User Realms > *Realm* > General tab.

5. **Evaluate your sign-in policy to use**—If you want to use a different sign-in page for Pocket PC users, you can define it in the Authentication > Signing In > Sign-in Pages tab and then create a sign-in policy that references the page using options in the Authentication > Signing In > Sign-in Policies tab. Or, you can create a custom sign-in page using the Pocket PC template files that are available in sample.zip.

6. **Specify allowed encryption strength**—Different types of devices allow different encryption strengths. You should specify the encryption strength in Connect Secure to match the requirement of your devices. For example, mobile phones often only accept 40-bit encryption. Review your end-users' device requirements and specify the allowed encryption strength on the System > Configuration > Security tab.

Related Documentation

- Defining Client Types
- Enabling WSAM on PDAs
- Enabling ActiveSync For Handheld Devices

Defining Client Types

The Client Types tab allows you to specify the types of systems your users may sign in from and the type of HTML pages to display when they do. In addition, client types are used to identify the operating system shown on the Device Management page for devices that use ActiveSync to synchronize e-mail with a Microsoft Exchange server. The user agent string used to identify a device during login may be different from the one in the ActiveSync message. For example, in the list of default user agent strings, *Apple-iPhone* and *Apple-iPad* are used only in ActiveSync messages.

To manage the client types:

1. In the admin console, choose **System > Configuration > Client Types**.

2. In the **User-agent string pattern** text box, enter the user agent string for the operating system(s) that you want to support. You can specify all or part of the string. For example,

you can use the default *DoCoMo* string to apply to all DoCoMo operating systems, or you can create a string such as DoCoMo/1.0/P502i/c10 to apply to a single type of DoCoMo operating system. You can use the * and ? wildcard characters in the string. Note that user agent strings on the system are case-insensitive.

If a device operating system shown on the Device Management page is **Other**, the ActiveSync message for the device has a user-agent string that is not defined here. To add the missing user-agent string:

1. Select **System > Log/Monitoring > User Access > Log**.

2. Search the User Access Log using the filter **id='AUT31094' && user=***'username'*. The AUT31094 is the ActiveSync log message ID, and you can select **System > Status > Devices** to get the device's username from the Device Management page. The log message looks like the following:

   Device record created for user jsmith@asglab.onmicrosoft.com to obtain Authorization Only access. (activesync_id=SAMSUNG1355815045478007_AM, user-agent=SAMSUNG-SAMSUNG-SGH-I997/100.202)

3. Copy the **user-agent=** value from the log message to the **User-agent string pattern** text box.

4. Select the client type (see Step 3) and click **Add**.

5. Select the type of HTML to display to users who sign in from the operating system specified in the previous step. Options include:

   - **Standard HTML**—The system displays all standard HTML functions, including tables, full-size graphics, ActiveX components, JavaScript, Java, frames, and cookies. Ideal for standard browsers, such as Firefox, Mozilla, and Internet Explorer.

   - **Compact HTML (iMode)**—The system displays small-screen HTML-compatible pages. This mode does not support cookies or the rendering of tables, graphics, ActiveX components, JavaScript, Java, VB script, or frames. (The only difference between this option and the Smart Phone HTML Basic option is the user interface.) Ideal for iMode browsers.

     NOTE: Form Post SSO is not supported on iMode appliances.

   - **Mobile HTML (Pocket PC)**—The system displays small-screen HTML-compatible pages that may contain tables, small graphics, JavaScript, frames, and cookies, but this mode does not facilitate the rendering of java applets or ActiveX components. Ideal for Pocket PC browsers.

   - **Smart Phone HTML Advanced**—The system displays small-screen HTML-compatible pages that may contain tables, small graphics, frames, cookies, and some JavaScript, but this mode does not facilitate the rendering of java applets, ActiveX components, or VB scripts. Ideal for Treo and Blazer browsers.

   - **Smart Phone HTML Basic**—The system displays small-screen HTML-compatible pages. This mode does not support cookies or the rendering of tables, graphics, ActiveX components, JavaScript, Java, VB script, or frames. (The only difference between this option and the Compact HTML option is the user interface.) Ideal for Opera browsers on Symbian.

> NOTE: The system rewrites hyperlinks to include the session ID in the URL instead of using cookies.

- **Mobile Safari, Android, Symbian, iPad**—The Mobile Safari (iPhone/iPod Touch), Android, and Symbian selections have Basic, Advanced, and Full HTML options.

6. Specify the order that you want to evaluate the user agents. The system applies the first rule in the list that matches the user's system. For example, you may create the following user agent string/HTML type mappings in the following order:

   1. User Agent String: *DoCoMo* Maps to: Compact HTML

   2. User Agent String: DoCoMo/1.0/P502i/c10 Maps to: Mobile HTML

   If a user signs in from the operating system specified in the second line, the system will display compact HTML pages to him, not the more robust mobile HTML, since his user agent string matches the first item in the list.

   To order mappings in the list, select the check box next to an item and then use the up and down arrows to move it to the correct place in the list.

3. Select the **Enable password masking for Compact HTML** check box if you want to mask passwords entered in iMode and other devices that use compact HTML. (Devices that do not use compact HTML mask passwords regardless of whether or not you select this check box.) Note that if your iMode users' passwords contain non-numeric characters, you must disable password masking because iMode devices only allow numeric data in standard password fields. If you disable masking, passwords are still transmitted securely, but are not concealed on the user's display.

4. Click **Save Changes**.

Related Documentation

- Task Summary: Configuring Connect Secure for PDAs and Handhelds

- Enabling WSAM on PDAs

- Enabling ActiveSync For Handheld Devices

Enabling WSAM on PDAs

When defining client/server applications to secure through Windows Secure Application Manager (WSAM) on PDA devices, you should define PDA-specific applications through the Users > User Roles > Select Role > SAM > Applications > Add Application page.

Listed below are some PDA-specific executable files that you might want to enable for PSA devices:

- **tmail.exe**—Specifies the Pocket Outlook application

  The system supports the following modes through Pocket Outlook:

  - S-IMAP/S-POP and S-SMTP

  - ActiveSync—If the supported PDAs to which you are providing Pocket Outlook access are using ActiveSync, you must ensure that the IP address of the Exchange Server appears in the list of destination hosts defined within the user role. Direct Push, a feature built into

Exchange Server 2007, is supported however you must set HTTPServerTimeout to 20 minutes or less.

- **mstsc40.exe**—Specifies the Windows Terminal Services application

- **iexplore.exe**—Specifies the Pocket Internet Explorer application

> NOTE: When using an existing WSAM role configuration originally set up for Windows PC users to provide secure access PDA users, ensure that the list of destination hosts defined within the user role is no larger than 1500 bytes. Very large lists of destination hosts can lock up the WSAM launcher on PDA devices due to memory buffer constraints.

For Windows Phone 5 users, WSAM adds log files to the \Program Files\Pulse Secure\WSAM\Log directory.

Related Documentation

- Task Summary: Configuring Connect Secure for PDAs and Handhelds

- Defining Client Types

- Enabling ActiveSync For Handheld Devices

Enabling ActiveSync For Handheld Devices

Using ActiveSync, you can synchronize data between a Windows-based desktop computer and handheld devices. Connect Secure can be used as a reverse proxy to allow users to synchronize their data without installing an additional client application, such as WSAM, on their handheld devices. More than 1000 concurrent connections is supported on an SA 6500.

Please note the following:

- Supports Windows Phone 5.0, 6.0 and 8.0 only.

- Supports Exchange Server 2003, 2007, 2010, 2013.

- ActiveSync does not use up concurrent user licenses, even when configured with certificate authentication.

- Both NTLM & Basic Auth on the Exchange server are supported.

- Both HTTP and HTTPS between Connect Secure and an Exchange server are supported.

- If Connect Secure is used for OWA & ActiveSync, the hostnames for OWA access and ActiveSync must be different.

- Direct Push is supported with ActiveSync, however you must set HTTPServerTimeout to 20 minutes or less. Direct Push is a feature built into Exchange Server 2007.

- ActiveSync does not work through a back-end web proxy.

- VIP sourcing settings are ignored for ActiveSync sessions. ActiveSync traffic from Connect Secure to a backend server is always sent with the Internal Port's source IP address.

To configure the system as a reverse proxy for use with ActiveSync:

1. In the admin console, choose **Authentication > Signing In > Sign-in Policies**.

2. To create a new authorization only access policy, click **New URL** and select **authorization only access**. Or, to edit an existing policy, click a URL in the Virtual Hostname column.

3. In the **Virtual Hostname** field, enter the name that maps to the system IP address. The name must be unique among all virtual hostnames used in pass-through proxy's hostname mode. The hostname is used to access the Exchange server entered in the Backend URL field. Do not include the protocol (for example, http:) in this field.

   For example, if the virtual hostname is myapp.ivehostname.com, and the backend URL is http://www.xyz.com:8080/, a request to https://myapp.ivehostname.com/test1 by the system is converted to a request to http://www.xyz.com:8080/test1. The response of the converted request is sent to the original requesting web browser.

4. In the **Backend URL** field, enter the URL for the Exchange server. You must specify the protocol, hostname and port of the server. For example, http://www.mydomain.com:8080/*.

   When requests match the hostname in the Virtual Hostname field, the request is transformed to the URL specified in the Backend URL field. The client is directed to the backend URL unaware of the redirect.

5. Enter a Description for this policy (optional).

6. Select the server name or **No Authorization** from the Authorization Server drop-down menu. If you select a server, ensure that the front-end server provides the SMSESSION cookie otherwise you will receive an error.

7. Select a user role from the Role Option drop-down menu.

   Only the following user role options are applicable for Autosync.

   - HTTP Connection Timeout (Users > User Roles > *RoleName* > Web > Options > View advanced options)

   - Allow browsing untrusted SSL web sites (Users > User Roles > *RoleName* > Web > Options > View advanced options)

   - Source IP restrictions (Users > User Roles > *RoleName* > General > Restrictions)

   - Browser restrictions (Users > User Roles > *RoleName* > General > Restrictions)

   Ensure the user role you select has an associated Web Access policy.

8. Select the **Allow ActiveSync Traffic only** option to perform a basic of validation of the HTTP header to ensure the request is consistent with ActiveSync protocol. If you select this option only ActiveSync protocol requests can be processed. If validation fails, a message is created in the user's event log. If you do not select this option, both ActiveSync and non-ActiveSync requests are processed.

9. Click **Save Changes**.

The System Status Overview page displays the number of current active concurrent connections and a histogram of the active concurrent connections (Authorization Only Access Active Connections plot in the Concurrent SSL Connections graph).

Related Documentation

- Task Summary: Configuring Connect Secure for PDAs and Handhelds

- Defining Client Types

- Enabling WSAM on PDAs

CHAPTER 52 Custom Expressions and System Variables

- Using Custom Expressions in Rule Configuration

Using Custom Expressions in Rule Configuration

This topic describes custom expressions. It is intended for advanced users. It includes the following information:

- Custom Expressions
- Custom Expression Elements
- Wildcard Matching
- Using Multivalued Attributes
- Specifying Multivalued Attributes in a Bookmark Name
- Distinguished Name Variables
- System Variables
- Custom Variables and Macros
- append
- daysdiff
- regmatch
- Specifying Fetch Attributes in a Realm
- Specifying the homeDirectory Attribute for LDAP

Custom Expressions

Many system rules, such as role mapping rules or resource policy rules, support custom expressions. A custom expression is a combination of variables that the system evaluates as a Boolean object. The expression returns true, false, or error.

You can write custom expressions in the following formats. Note that elements of these formats are described in greater detail in the table that follows:

- *variable comparisonOperator variable*
- *variable comparisonOperator simpleValue*
- *variable comparisonOperator* (*simpleValue*)
- *variable comparisonOperator* (OR *Values*)
- *variable comparisonOperator* (AND *Values*)

- *variable comparisonOperator* (*time* TO *time*)

- *variable comparisonOperator* (*day* TO *day*)

- isEmtpy (*variable*)

- isUnknown (*variable*)

- (*customExpr*)

- NOT *customExpr*

- ! *customExpr*

- *customExpr* OR *customExpr*

- *customExpr* || *customExpr*

- *customExpr* AND *customExpr*

- *customExpr* && *customExpr*

     NOTE: The custom expression should be less than 64K.

Custom Expression Elements

Table212: Custom Expression Elements

Element	Description
variable	Represents a system variable. A variable name is a dot-separated string, and each component can contain characters from the set [a-z A-Z 0-9_] but cannot start with a digit [0-9]. Variable names are case-insensitive. For system variables that you may use in role mapping rules and resource policies.
	When writing a custom expression in a log query field, you need to use system log variables. These variables are described in the **Filter Variables Dictionary** on the Filter page (**System > Log/Monitoring > Events | User Access | Admin Access > Filters** > *Select Filter* tab).
	Quoting syntax for variables:
	The system supports a quoting syntax for custom expression variables that allows you to use any character except '.' (period) in a user attribute name. To escape characters in an attribute name, quote some or all of the variable name using { } (curly-braces). For example, these expressions are equivalent:
	• userAttr.{Login-Name} = 'xyz'
	• userAttr.Login{-}Name = 'xyz'
	• {userAttr.Login-Name} = 'xyz'
	• userA{ttr.L}{ogin-}Name = 'xyz'

- Escape characters supported within quotes:
- ****—Escape a backslash (\).
- **\{**—Escape a left curly brace ({).
- **\}**—Escape a right curly brace (}).
- **\hh**—Escape a hexadecimal value where hh is two characters from [0-9A-Fa-f].

Examples:

- userAttr.{Tree Frog} = 'kermit'
- userAttr.{Tree\20Frog} = 'kermit'

Notes:

- There is no limit to the number of quotes you can use in a variable name.
- You can use the quoting syntax with any variable, not just userAttr.* variables.
- You need to use curly-brace quotes only when writing custom expressions.

comparisonOperator	One of the following: • **=**—Equal to. Use with strings, numbers, and DNs. • **!=**—Not equal to. Use with strings, numbers, and DNs. • **<**—Less than. Use with numbers. • **<=**—Less than or equal to. Use with numbers. • **>**—Greater than. Use with numbers. • **>=**—Greater than or equal to. Use with numbers.
simpleValue	One of the following: • *string* — quoted string that may contain wildcards. • *IP Address*—a.b.c.d • *subnet*—a.b.c.d/subnetBitCount or a.b.c.d/netmask • *number*—Positive or negative integer • *day*—SUN MON TUE WED THU FRI SAT Notes about strings: • A string may contain all characters except <nl> (newline) and <cr> (carriage return). • Strings can be any length. • String comparisons are case-insensitive. • Strings can be quoted with single- or double-quotes. A quoted string may contain wildcards, including star(*), question mark (?), and square brackets ([ ]). • *variable comparisonOperator variable* comparisons are evaluated without wildcard matching. • Use a backslash to escape these characters: single-quote (') — \' double-quote (") — \" backslash (\) — \\ hexadecimal — \hh [0-9a-fA-F] Note about *day*: Day and time comparisons are evaluated in the system's time zone. Day range (*day* TO *day*) calculations start with the first day and step forward until the second day is reached. In time range (*time* TO *time*) calculations, the first value must be earlier than the second value. Only time variables can be compared to day and time values. The time variables are: time.* and loginTime.*.

time	Time of day in one of the following formats:
	• *HH:MM* — 24-hour
	• *HH:MMam* — 12-hour
	• *HH:MMpm* — 12-hour
	• *H:MM* — 24-hour
	• *H:MMam* — 12-hour
	• *H:MMpm* — 12-hour
	Day and time comparisons are evaluated in the system's time zone. Day range (day TO day) calculations start with the first day and step forward until the second day is reached. In time range (time TO time) calculations, the first value must be earlier than the second value. Only time variables can be compared to day and time values. The time variables are: time.* and loginTime.*.
OR *Value*	String containing one or more OR comparisons:
	Examples:
	variable comparisonOperator (*number* OR *number* ...)
	variable comparisonOperator (*string* OR *string* ...)
AND *Value*	String containing one or more AND comparisons.
	Examples:
	variable comparisonOperator (*number* AND *number* ...)
	variable comparisonOperator (*string* AND *string* ...)
isEmpty	Function that takes a single variable name (*variable*) argument and returns a boolean value. isEmpty() is true if the variable is unknown or has a zero-length value, zero-length strings, and empty lists.
	Example: isEmpty(userAttr.terminationDate)
isUnknown	Function that takes a single variable name (*variable*) argument and returns a boolean value. isUnknown() is true if the variable is not defined. User attributes (userAttr.* *variables*) are unknown if the attribute is not defined in LDAP or if the attribute lookup failed (such as if the LDAP server is down).
	Example: isUnknown(userAttr.bonusProgram)
NOT, !	Logical negation *comparisonOperator*. The negated expression evaluates to true if the *customExpr* is false and evaluates to false if the *customExpr* is true. The operators NOT, AND, and OR are evaluated from highest to lowest precedence in this order: NOT (from right), AND (from left), OR (from left).
OR, ||	Logical operator OR or ||, which are equivalent. The operators NOT, AND, and OR are evaluated from highest to lowest precedence in this order: NOT (from right), AND (from left), OR (from left).
AND, &&	Logical AND or &&, which are equivalent. The operators NOT, AND, and OR are evaluated from highest to lowest precedence in this order: NOT (from right), AND (from left), OR (from left).
customExpr	Expression written in the Custom Expression Syntax (see above).

Wildcard Matching

In a quoted string, supported wildcards include:

• star (*)—A star matches any sequence of zero or more characters.

• question mark (?)—A question mark matches any single character.

- square brackets ([ ])—Square brackets match one character from a range of possible characters specified between the brackets. Two characters separated by a dash (-) match the two characters in the specified range and the lexically intervening characters. For example, 'dept[0-9]' matches strings "dept0", "dept1", and up to "dept9".

To escape wildcard characters, place them inside square brackets. For example, the expression ' userAttr.x = " *value* [*]" ' evaluates to true if attribute x is exactly "*value**".

Using Multivalued Attributes

Multivalued attributes—attributes that contain two or more values—provide you with a convenient method for defining resources that expand into multiple individual bookmarks on the users' bookmarks page.

For example, assume that the user's LDAP directory contains the multivalued attribute HomeShares: \\Srv1\Sales;\\Srv2\Marketing. When you configure the Windows File share resource definition using the HomeShares multivalued attribute, \\<userAttr.HomeShares>, the user sees two bookmarks:

- \\Srv1\Sales

- \\Srv2\Marketing

Now let's assume the user's LDAP directory contains a second multivalued attribute defined as HomeFolders: Folder1;Folder2;Folder3. When you configure the Windows File share resource using both of the multivalued attributes, \\<userAttr.HomeShares>\<userAttr.HomeFolders>, the user sees the following six bookmarks:

- \\Srv1\Sales\Folder1

- \\Srv1\Sales\Folder2

- \\Srv1\Sales\Folder3

- \\Srv2\Marketing\Folder1

- \\Srv2\Marketing\Folder2

- \\Srv2\Marketing\Folder3

The only exception to this functionality is when the variable includes an explicit separator string. In this case, only one bookmark containing multiple resources displays on the users' bookmark page.

You specify the separator string in the variable definition using the syntax sep='string' where string equals the separator you want to use. For example, to specify a semi-colon as the separator, use the syntax <variable.Attr sep=';'>.

Use the following syntax for multivalued attributes handling. Note that <variable> refers to a session variable such as <userAttr.name> or <CertAttr.name>:

- <variable[Index]>—You specify indexes in a variety of ways. If, for example, the total number of values for a given index is 5, and you want to specify the entire range of values you use <variable[ALL]>. If you want to specify only the fourth value, you use <variable[4]>.

- <variable> is the same as <variable[ALL]>.

- <variable> is the same as <variable[ALL]>.

- <variable sep='str'> and <variable[All] sep='str'> — These variable definitions always refer to a single string value with all the tokens expanded out with separator strings between the values.

i   NOTE: Variable names cannot contain spaces.

Specifying Multivalued Attributes in a Bookmark Name

Another common case of using multivalued attributes occurs when you include a variable in a bookmark name and in a URL or file server/share field.

For example, again assume that the user's LDAP directory contains the multivalued attribute HomeShares: \\Srv1\Sales;\\Srv2\Marketing. When you configure the Windows File share resource definition using the HomeShares multivalued attribute, \\<userAttr.HomeShares>, and you use the same attribute in the bookmark name field, <userAttr.HomeShares>, the system creates two bookmarks:

- Srv1\Sales bookmark pointing to \\Srv1\Sales

- Srv2\Marketing bookmark pointing to \\Srv2\Marketing

This does not create a situation in which you end up with the following set of conditions:

- Srv1\Sales bookmark pointing to \\Srv1\Sales

- Srv1\Marketing bookmark pointing to \\Srv1\Marketing (error)

- Srv2\Sales bookmark pointing to \\Srv1\Sales (error)

- Srv2\Marketing bookmark pointing to \\Srv2\Marketing

Distinguished Name Variables

You can compare a distinguished name (DN) to another DN or to a string, but the system ignores wildcards, white space, and case. Note, however, that the system takes the order of DN keys into consideration.

When the system compares an expression to a DN to a string, it converts the string to a distinguished name before evaluating the expression. If the system cannot convert the string due to bad syntax, the comparison fails. The DN variables are:

- userDN

- certDN

- certIssuerDN

The system also supports DN suffix comparisons using the **matchDNSuffix** function. For example:

```
matchDNSuffix( certDn, "dc=danastreet,dc=net")
```

Within the parenthesis, the first parameter is the " full" DN and the second is the suffix DN. You can use a variable or string for each parameter. Note that this first parameter should have more keys than the second (suffix parameter). Otherwise, if they are equal, it is the same as *<firstparam>* = *<secondparam>*. If the second parameter has more keys, **matchDNsuffix** returns false.

System Variables

Table 208 lists and defines system variables, gives an example for each system variable, and provides a guide as to where you may use system variables.

Table213: System Variables and Examples

Variable	Description	Usage	Examples
authMethod	Type of authentication method used to authenticates a user.	role mapping rules, resource policy rules	authMethod = 'ACE Server'
cacheCleanerStatus	The status of Cache Cleaner. Possible values: 1 - if it is running 0 - if otherwise		cacheCleanerStatus = 1 cacheCleanerStatus = 0
certAttr.<cert-attr>	Attributes from a client-side certificate. Examples of certAttr attributes include: • C - country • CN - common name • description - description • e-mailAddress - e-mail address • GN - given name • initials - initials • L - locality name • O - organization • OU - organizational unit • SN - surname • serialNumber- serial number • ST - state or province • title - title • UI - unique identifier Use this variable to check that the user's client has a client-side certificate with the value(s) specified.	• role mapping rules • resource policy rules • SSO parameter fields • LDAP configuration	certAttr.OU = 'Retail Products Group'

certAttr.altName.<Alt-attr>	Subject alternative name value from a client-side certificate where <Alt-attr> may be: - Email - EmailId - EmailDomain - DNS - registeredId - ipAddress - UPN - UPNid - UPNDomain - fascn - fascnAC - fascnSC - fascnCN - fascnCS - fascnICI - fascnPI - fascnOC - fascnOI - fascnPOA - fascnLRC	- role mapping rules - resource policy rules - SSO parameter fields - LDAP configuration	- certAttr.altName.email = "joe@company.com" - certAttr.altName.ipAddress = 10.10.83.2
certAttr.serialNumber	Client certificate serial number. Note that all characters other than [0-9 a-f A-F] are stripped out of a string before comparison with certAttr.SN. Wildcards are not supported.	- role mapping rules - resource policy rules - SSO parameter fields - LDAP configuration	- certAttr.SerialNumber = userAttr.certSerial - certAttr.SerialNumber = "6f:05:45:ab"
certDN	Client certificate subject DN. Wildcards are not permitted.	role mapping rules, resource policy rules	- certDN = 'cn=John Harding,ou=eng,c=Company' - certDN = userDN (match the certificate subject DN with the LDAP user DN) - certDN = userAttr.x509SubjectName - certDN = ('cn=John Harding,ou=eng,c=Company' or 'cn=Julia Yount,ou=eng,c=Company')

certDN.<subject-attr>	Any variable from the client certificate subject DN, where subject-attr is the name of the RDN key. Use to test the various subject DN attributes in a standard x.509 certificate.	• role mapping rules • resource policy rules • SSO parameter fields • LDAP configuration	• certDN.OU = 'company' • certDN.E = 'joe@company.com' • certDN.ST = 'CA'
certDNText	Client certificate user DN stored as a string. Only string comparisons to this value are allowed.	• role mapping rules • resource policy rules • SSO parameter fields	certDNText = 'cn=John Harding,ou=eng,c=Company'
certIssuerDN	Client certificate-issuer subject DN. This variable works like a standard DN attribute such as CertDN. Wildcards are not permitted.	• role mapping rules • resource policy rules • SSO parameter fields	• certIssuerDN = 'cn=John Harding,ou=eng,c=Company' • certIssuerDN = userAttr.x509Issuer • certIssuerDN = ('ou=eng,c=Company' or 'ou=operations,c=Company')
certIssuerDN.<issuer-attr>	Any variable from the client certificate-issuer subject DN, where issuer-attr is the name of the RDN key.	• role mapping rules • resource policy rules • SSO parameter fields	• certIssuerDN.OU = 'company' • certIssuerDN.ST = 'CA'
certIssuerDNText	Client certificate-issuer subject DN stored as a string. Only string comparisons to this value are allowed.	• role mapping rules • resource policy rules • SSO parameter fields	certIssuerDNText = 'cn=John Harding,ou=eng,c=Company'
defaultNTDomain	Contains the Domain value set in the authentication server configuration when you use AD/NT authentication.	• role mapping rules • resource policy rules • SSO parameter fields	defaultNTDomain=" CORP"

group.<group-name>	User's group membership as provided by the realm authentication or directory server.	• role mapping rules • resource policy rules Only those groups evaluated for role mapping rules are available in the detailed rules (conditions) in the resource policies. We recommend that you use the **groups** variable instead of **group.<group-name>**, which is supported only for backwards compatibility.	• group.preferredPartner • group.goldPartner or group.silverPartner • group.employees and time.month = 9 Combination examples: Allow all partners with active status from Monday to Friday but preferred partners Monday through Saturday: ((group.partners and time = (Mon to Fri)) or (group.preferredPartners and time = (Mon to Sat))) and userAttr.partnerStatus = 'active' NOTE: Spaces are not supported, such as, group.sales managers
groups	List of groups as provided by the realm authentication or directory server. NOTE: You can enter any characters in the groupname, although wildcard characters are not supported.	• role mapping rules • resource policy rules • SSO parameter fields	groups=('sales managers')
hostCheckerPolicy	Host Checker polices that the client has met.	• role mapping rules • resource policy rules • SSO parameter fields	hostCheckerPolicy = ('Norton' and 'Sygate') and cacheCleanerStatus = 1hostCheckerPolicy = ('Norton' and 'Sygate')
loginHost	Hostname or IP address that the browser uses to contact the Pulse Secure client service.	• role mapping rules • resource policy rules • SSO parameter fields • LDAP configuration	loginHost = 10.10.10.10

loginTime	The time of day at which the user submits his credentials. The time is based on system time. NOTE: When using this variable in an SSO parameter field, the variable returns the UNIX string time.	• role mapping rules • resource policy rules • SSO parameter fields	• loginTime = (8:00am) • loginTime= (Mon to Fri)
loginTime.day	The day of month on which the user submits his credentials, where day is 1-31. The time is based on the system time. You cannot use the TO operator with this variable.	• role mapping rules • resource policy rules	loginTime.day = 3
loginTime.dayOfWeek	The day of the week on which the user submits his credentials, where dayOfWeek is in the range [0-6] where 0 = Sunday. The system does not support the **TO** operator with time.dayOfWeek expressions if you use numbers instead of strings. In other words, " loginTime.dayOfWeek = (2 TO 6)**"** does not work, but " loginTime.dayOfWeek = (mon to fri)**"** does work.	• role mapping rules • resource policy rules	• loginTime.dayOfWeek = (0 OR 6) • loginTime.dayOfWeek = (mon TO fri) • loginTime.dayOfWeek = (1) • loginTime.dayOfWeek = 5
loginTime.dayOfYear	The numeric day of the year on which the user submits his credentials, where dayOfYear can be set to [0-365]. You cannot use the TO operator with this variable.	• role mapping rules • resource policy rules	loginTime.dayOfYear = 100
loginTime.month	The month in which the user submits his credentials, where month can be set to [1-12] where 1 = January. You cannot use the TO operator with this variable.	• role mapping rules • resource policy rules	loginTime.month >= 4 AND loginTime.month <=9
loginTime.year	The year in which the user submits his credentials, where year can be set to [1900-2999]. You cannot use the TO operator with this variable.	• role mapping rules • resource policy rules	loginTime.year = 2005
loginURL	URL of the page that the user accessed to sign in. The system gets this value from the Administrator URLs|User URLs column on the **Authentication > Signing In > Sign-in Policies** page of the admin console.	• role mapping rules • resource policy rules • SSO parameter fields • LDAP configuration	loginURL = */admin

networkIf	The network interface on which the user request is received. Possible values: internal, external	• role mapping rules • resource policy rules • SSO parameter fields	sourceIp = 192.168.1.0/24 and networkIf = internal
ntdomain	The NetBIOS NT domain used in NT4 and Active Directory authentication.	• role mapping rules • SSO parameter fields	ntdomain = jnpr
ntuser	The NT username used in Active Directory authentication	• role mapping rules • SSO parameter fields	ntuser = jdoe
password password[1] password[2]	The password entered by the user for the primary authentication server (password and password[1]) or the secondary authentication server (password[2]).	• role mapping rules • resource policy rules • SSO parameter fields	**password = A1defo2z**
realm	The name of the authentication realm to which the user is signed in.	• role mapping rules • resource policy rules • SSO parameter fields	Realm = ('GoldPartners' or 'SilverPartners') NOTE: AND condition will always fail as a user is only allowed to sign in to a single realm in a session.
role	List of all the user roles for the session. In SSO, if you want to send all the roles to back-end applications, use <role sep = ";"> - where sep is the separator string for multiple values. The system supports all separators except " and >.	• resource policy rules • SSO parameter fields	• Role = ('sales' or 'engineering') • Role = ('Sales' AND 'Support')

sourceIP	The IP address of the machine on which the user authenticates. You can specify the netmask using the bit number or in the netmask format: '255.255.0.0'. Note that you can evaluate the sourceIP expression against a string variable such as an LDAP attribute.	• role mapping rules • resource policy rules • SSO parameter fields	• sourceIP = 192.168.10.20 • sourceIP = 192.168.1.0/24 and networkIf internal • userAttr.dept = ('eng' or 'it') and sourceIP = 10.11.0.0/16 • sourceIP = 192.168.10.0/24 (Class C) is the same as: sourceIP = 192.168.10.0/255.255.255.0 • sourceIP=userAttr.sourceip
time	The time of day at which the role mapping rule or resource policy rule is evaluated. The time of the day can be in 12-hour or 24-hour format.	• role mapping rules • resource policy rules	• time = (9:00am to 5:00pm) • time = (09:00 to 17:00) • time = (Mon to Fri) Combination examples: Allow executive managers and their assistants access from Monday to Friday: userAttr.employeeType = ('*manager*' or '*assistant*') and group.executiveStaff and time = (Mon to Fri)
time.day	The day of month on which the user submits his credentials to, where day is 1-31. The time is based on the system time.	• role mapping rules • resource policy rules	loginTime.day = 3
time.dayOfWeek	The day of the week on which the role mapping rule or resource policy rule is evaluated, where dayOfWeek is in the range [0-6] where 0 = Sunday.	• role mapping rules • resource policy rules	• loginTime.dayOfWeek = (0 OR 6) • loginTime.dayOfWeek = (1 to 5) • loginTime.dayOfWeek = 5
time.dayOfYear	The day of the year on which the role mapping rule or resource policy rule is evaluated. Possible values include: 1-365.	• role mapping rules • resource policy rules	time.dayOfYear = 100
time.month	The month in which the role mapping rule or resource policy rule is evaluated. Possible values include: 1-12	• role mapping rules • resource policy rules	• time.month >= 9 and time.month <= 12 and time.year = 2004 • group.employees and time.month = 9

time.year	The year in which the role mapping rule or resource policy rule is evaluated, where year can be set to [1900-2999].	• role mapping rules • resource policy rules	time.year = 2005
user user@primary_auth_server_name user@secondary_auth_server_name	Pulse Secure client username for the user's primary authentication server (user and user@primary_auth_server_name) or secondary authentication server (user@secondary_auth_server_name). Use when authenticating against an Active Directory server, domain and username. primary_auth_server_name is the name of the primary auth server. If there are spaces or special characters in the name, it can be enclosed in curly brackets. For example user@{My Primary Auth Server} secondary_auth_server_name is the name of the secondary auth server. If there are spaces or special characters in the name, it can be enclosed in curly brackets. For example user@{My Secondary Auth Server} NOTE: When including a domain as part of a username, you must include two slashes between the domain and user. For example: user='yourcompany.net\\joeuser'.	• role mapping rules • resource policy rules • SSO parameter fields	• user = 'steve' • user = 'domain\\steve'
username username@primary_auth_server_name username@secondary_auth_server_name	Pulse Secure client system username for the user's primary authentication server (username and username@primary_auth_server_name) or secondary authentication server (username@secondary_auth_server_name). If the user is signing in to a certificate authentication server, then the user's Pulse Secure client system username is the same as CertDN.cn. primary_auth_server_name is the name of the primary auth server. If there are spaces or special characters in the name, it can be enclosed in curly brackets. For example user@{My Primary Auth Server} secondary_auth_server_name is the name of the secondary auth server. If there are spaces or special characters in the name, it can be enclosed in curly brackets. For example user@{My Secondary Auth Server}	• role mapping rules • resource policy rules • SSO parameter fields	• username = 'steve' and time = mon • username = 'steve' • username = 'steve*' • username = ('steve' or '*jankowski')

userAgent	The browser's user agent string.	• role mapping rules • resource policy rules • SSO parameter fields	The browser's user agent string.
userAttr.<auth-attr>	User attributes retrieved from an LDAP, RADIUS, or SiteMinder authentication or directory server.	• role mapping rules • resource policy rules • SSO parameter fields	• userAttr.building = ('HQ*' or 'MtView[1-3]') • userAttr.dept = ('sales' and 'eng') • userAttr.dept = ('eng' or 'it' or 'custsupport') • userAttr.division = 'sales' • userAttr.employeeType != 'contractor' • userAttr.salaryGrade > 10 • userAttr.salesConfirmed >= userAttr.salesQuota *Negative examples:* • userAttr.company != "Acme Inc" or not group.contractors • not (user = 'guest' or group.demo) *Combination examples:* Allow executive managers and their assistants access from Monday to Friday: userAttr.employeeType = ('*manager*' or '*assistant*') and group.executiveStaff and time = (Mon to Fri) Allow all partners with active status from Monday to Friday but preferred partners Monday through Saturday: ((group.partners and time = (Mon to Fri)) or (group.preferredPartners and time = (Mon to Sat))) and userAttr.partnerStatus = 'active'

userDN	The user DN from an LDAP server (not applicable to Active Directory auth server with ldap group lookup). If the user is authenticated by the LDAP server, then this DN is from the authentication server; otherwise, the DN comes from the realm's Directory/Attribute server.	• role mapping rules • resource policy rules	• userDN = 'cn=John Harding,ou=eng,c=Company' • userDN = certDN
userDN.<user-attr>	Any variable from the user DN, where user-attr is the name of the RDN key.	• role mapping rules • resource policy rules • SSO parameter fields	Any variable from the user DN, where user-attr is the name of the RDN key.
userDNText	User DN stored as a string. Only string comparisons to this value are allowed.	• role mapping rules • resource policy rules • SSO parameter fields	userDNText = 'cn=John Harding,ou=eng,c=Company'

Custom Variables and Macros

Custom variables, like system variables, are name-value pair tags that you can use when defining role mapping rules, resource policy rules and SSO parameter fields.

Custom variables are created in the Server Catalog (for example, Authentication > Auth Server > *Name* > Settings) by using a predefined macro on a system variable. Available macros are:

- REGMATCH – Matches a regular expression pattern against a string text.

- APPEND – Appends a text string to another text string.

- DAYSDIFF – Calculates the difference between two dates.

> NOTE: These macros are located under Variable Operators in the Variables tab of the Server Catalog window.

A custom variable name is a dot-separated string. Each component can contain characters from the set [a-z A-Z 0-9 _] but cannot start with a digit [0-9]. Custom variable names are case-insensitive.

Custom variables are referenced as **customVar.<variableName>** . For example, if you create a custom variable with the name **check-prefix**, you reference this custom variable as **customVar.check-prefix**.

append

Syntax   APPEND (attr, TextString)

APPEND (attr, attr2)

Description Append a text string to an attribute or append an attribute to another attribute and store the resulting string in the custom variable.

Options **attr**—System variable of type string.

  TextString—Quoted ASCII string.

  attr2—System variable of type string.

Output Fields Returns a String value. If no match is found, returns an empty string.

  If the system variable is multivalued, the custom variable is also multivalued and uses the same order as the system variable.

Sample Output

```
APPEND (userName, "@juniper.net")
```

In this example, the string "@juniper.net" is appended to the userName value.

daysdiff

Syntax DAYSDIFF (attr, timeformat)

Description Calculates the number of days between the attribute and the current time.

Options **attr**—System variable of type string.

  timeformat—Output time format. Valid values are: UTC, TIMET, MMDDYYYY

Output Fields Returns an Integer value.

Sample Output

```
DAYSDIFF ( certAttr.validUpto, UTC)
```

In this example, calculate the difference in days between the current time and the value of certAttr.validUpto and express the time in UTC (Coordinated Universal Time).

regmatch

Syntax REGMATCH (attr, regex, groupingNumber)

Description Match the regular expression pattern against an attribute and store the result in the custom variable.

Options **attr**—System variable of type string.

  regex—Quoted string containing the regular expression to be applied to the attr option.

  groupingNumber—The group value to assign to the custom variable.

Additional Information The regular expression supports the Perl Compatible Regular Expressions (PCRE) syntax. A grouping (capture buffer) in the regex pattern can also be used to define a custom variable.

Output Fields Returns a String value. If no match is found, returns an empty string.

If the system variable is multivalued, the custom variable is also multivalued and uses the same order as the system variable.

Sample Output

```
REGMATCH (mailId, "^(.*)@juniper.net$", 1)
```

In this example, a mailId of myName@pulsesecure.net creates a custom variable with value "myName".

Specifying Fetch Attributes in a Realm

To facilitate the support for various parameterized settings in user roles and resource policies, you have the ability to specify additional fetch attributes. The system stores the fetch attributes when users log in so that you can use them in parameterized role or resource policy definitions.

The system pulls all the attributes that are currently stored in the Sever Catalog for the user's authentication or authorization LDAP server. So, make sure to add the LDAP user attributes that are used in role or resource policy definitions in the LDAP Server Catalog first.

When a user logs in, the system retrieves user attributes that are referenced in the role mapping rules plus all of the additional attributes referenced in the Server Catalog and stores all these values. Note that this should not incur a significant performance overhead because all the user attributes are retrieved in one single LDAP query.

> NOTE: When you substitute variables, such as in IP/Netmasks or hostnames, the values in the session are appropriately converted into the data type that is required by the particular application definition.

Specifying the homeDirectory Attribute for LDAP

You can create a bookmark that automatically maps to a user's LDAP home directory. You can accomplish this using the LDAP attribute homeDirectory. You need to configure a realm that specifies the LDAP server instance as its auth server, and you need to configure role-mapping rules and a bookmark that points to the LDAP homeDirectory attribute.

Related Documentation

- *User Access Management Framework Feature Guide*
- *User Access Management Framework Feature Guide*

PART 7 Index

- Index

Index

Symbols

A

I

J