

# StoreVault S500 Performance



---

## overview

The S500 is a multi-protocol system capable of running the CIFS, NFS, and iSCSI protocols simultaneously. It serves data using two interfaces:

1. **Ethernet.** Two embedded Broadcom NICs are included standard, and may be used for any TCP/IP protocol (CIFS, NFS and iSCSI).
2. **Fibre Channel.** The Fibre Channel Protocol is supported using an optional dual-port QLogic target card with limited QLogic switch and HBA components available from the StoreVault Division.

Many end users and resellers are interested in a level of precision that can be frustratingly difficult to provide. We have learned that results are entirely dependent on variables that can be out of our control, like network design, application requirements, disk drive layouts and the ultimate expectations of the user. It makes more sense to give an honest answer, and say that we don't know precisely what will happen for you, and to provide context and examples. This document includes benchmark testing from our own labs with detailed configuration parameters, as well as data collected from several real-live end user sites.

---

## StoreVault S500 test configuration

The performance information below was generated using the following specifications:

### 1. Benchmark StoreVault System Configuration

The S500 system was a standard, factory-configured unit. Specifications are as follows:

- Standard shipping options:
  - 1GB of memory
  - Data ONTAP StoreVault Edition
  - iSCSI license
- User-definable options:
  - 12 250GB SATA II drives
  - CIFS license
  - NFS license
  - Snapshots turned on, using the factory default settings of every night at midnight and 4 times daily.

Fibre Channel Protocol was not available at the time of testing (we had not yet qualified the QLogic components in the kit).

### 2. Benchmark Network Configuration

The testbed network included:

- Standard Gigabit Ethernet network connections to Intel-based servers through a NETGEAR unmanaged 1GBps switch.
- Jumbo Frames were NOT enabled, since some networks may not have this feature available

### 3. Tests Completed

The StoreVault Division ran the following tests internally to produce these results:

- SIO, a NetApp internal test that runs on UNIX and measures read and write performance in manner similar to IOMeter. This is very useful when comparing storage system performance from Windows and UNIX servers.
- IOMeter 2004.07.30 to test disk subsystem performance. We're not big fans of this test since network storage is often impacted by LAN factors, but people always ask us about it.

We also had Lionbridge Labs, a third-party testing facility, run the following tests and have included them in our report:

- NetBench 7.0.3
- IOMeter 2004.07.30

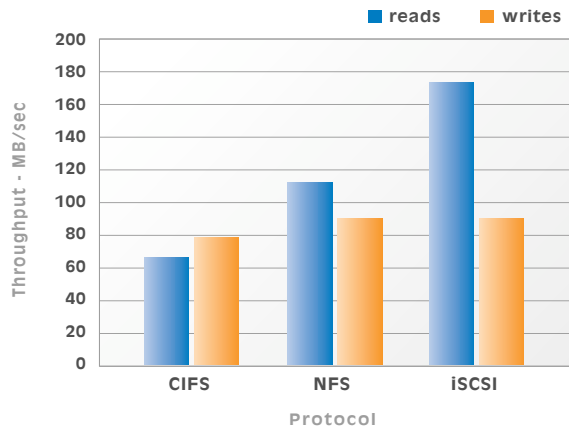
Oracle performed some baseline benchmarking to ensure that the StoreVault was capable of handling database transactions.

## StoreVault S500 test results

Overall, we found that the S500 is fastest when writing blocks than when writing using file protocols, and can read data about the same for all protocols. We used NetBench, SiO and Iometer to produce this chart.

**Figure 1 - S500 Performance by Protocol\***


\*sequential reads with 32KB block size



## CIFS results

Windows file sharing clients behaved as shown below. These are Lionbridge Labs results.

**Figure 2 - S500 CIFS Performance**



The S500 maintains its peak performance of just over 70 MB/sec even as large numbers of Windows clients are simultaneously reading and writing data. We believe that more disks in a StoreVault contribute to higher throughput, and recommend that systems be filled for best performance. We also noted that our default configuration includes 1GB of RAM, and that more RAM may impact the speed at which the system peaks. Further testing and analysis remains to be done.

## iSCSI results

See Figure 1 for Iometer results when the StoreVault is an iSCSI target.

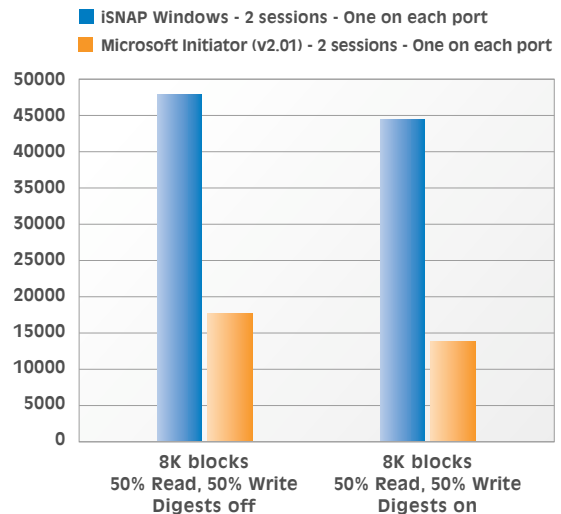
Performance scales fairly linearly for the StoreVault as block size increases. Tests were performed with traffic across both embedded NICs.

A note on iSCSI TOE cards. TOE cards are designed to improve the performance of a host server by reducing the amount of CPU overhead required by network processing. Actual results depend entirely on the setup and configuration of the host server, the applications it is running and the network configuration. In general, IOPS improve and system CPU workload reduces on the host when an iSCSI TOE card is added.

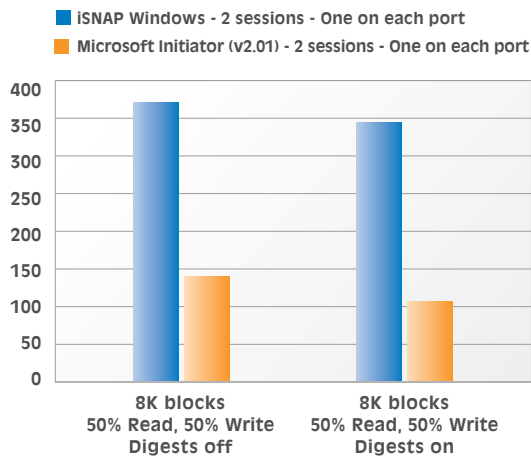
Since most host servers will not have two processors dedicated to network traffic, the results displayed below are highly case specific. The most useful visual point in this testing is that there is a difference when a TOE card is added; the amount of the difference is variable.

Silverback Systems compared their iSNAP 2110 hardware accelerator to the Microsoft software iSCSI initiator. The test was configured so that each of the two iSCSI HBAs was deployed in a single-processor 2.4 GHz Xeon server with 512 MB of RAM running Windows 2003 SP1 Server, while a dual-processor 2.8 GHz Xeon server equipped with the iSNAP 2110 HBA running in RAMdisk mode acted as the iSCSI target. The following Iometer results were achieved:

## iSCSI TOE Card IOPS



## iSCSI TOE Card MBPS



## NFS results

See Figure 1 for our own SiO NFS results (a very similar test to IOMeter but runs on UNIX).

## Oracle OAST results

Oracle used their standard Oracle Automated Stress Test (OAST) to measure Average Transactions Per Minute (ATPM) for Windows and Linux over iSCSI and Linux over NFS. The Automated Storage Manager (ASM) feature of Oracle 10g provides a database administrator with a simple storage management interface that is consistent across all server and storage platforms.

## Testbed configuration:

### Software

- Oracle 10gR2 version 10.2.0.1
- Linux RedHat REL4 Update 4, 32 bit version OR Windows Server 2003 32 bit version, EE edition SP1
- iSCSI initiator – RPM OR Microsoft iSCSI Software Initiator version 2.01
- SVM 2.0 on Windows
- DataONTAP SVE 7.0.3S4
- ASMLIB 2.0

### Host

- 2 CPU 64Bit EM64T chips with hyperthreading
- 2 Gigs of memory
- One 76 GB internal SCSI drive
- Three internal Ethernet cards in server. 2 GE Intel PRO 1000 cards bonded using Linux channel binding with failover and load balancing enabled. No changes to Frame Size/MTU
- 1 Gb Switch

### StoreVault S500

- ONTAP version: 7.0.3S4
- Total disks: 12
- Size of each disk: 250G
- RAID DP and a hot spare:

## Test Results:

After a 24 hour run, the testbed reported the following results:

NFS – Linux	6385.21 ATPM
ASM – Linux	7023.16 ATPM
ASM – Windows	7033.29 ATPM

## real-life anecdotes

All the benchmarking in the world won't tell predict what will happen in real life. In hopes that this provides some context, several end users have graciously allowed us to provide some detail on their networks and their results.

### Scenario 1: Mellano & Co

(see the success story on [www.storevault.com](http://www.storevault.com) for more information)

- Switch: Dell PowerEdge Gigabit Ethernet
- Applications: Multiple VMware virtual machines on a single StoreVault LUN
- Throughput (measured by the Dell switch utilities): ~430MB/s over iSCSI
- Comments: "Our StoreVault holds more data and is faster than our old servers, even when getting small amounts of data from multiple locations, a typical task in a VMware environment." – Danny Covell

### Scenario 2: Virginia Military Institute

(see the success story on [www.storevault.com](http://www.storevault.com) for more information)

- Host: Sun Fire x4600 x64 server with 4x AMD Opteron Model 885 processors, 8x 2GB DDR-400 memory, 2x 73GB SAS drives, 4 on-board NICs, 4 dual port NICs
- Switch: Two Dell iSCSI switches
- Applications: VMware for virtualizing application servers, and NAS
- Comments: "I honestly didn't expect iSCSI to perform as well as it does. I read a lot of articles and spoke with quite a few people with differing views on iSCSI. Some said we would love it, others didn't recommend touching it – or at least not using it to support high disk i/o applications like Windows operating systems or SQL servers. There was also some debate over whether we needed to use TOE cards or just a basic NIC with the Microsoft iSCSI initiator. We had been told that the iSCSI overhead would slow everything else running on the server down but it hasn't. We're using Intel Gigabit NICS with the software initiator without any trouble or degradation in performance of the host. From a disk performance perspective, I can't say that I've noticed a difference between the StoreVault and the EMC SAN that we replaced – neither have I heard anything negative about slow response from users." – Jef McReery

### Scenario 3: iBiquity

(see the success story on [www.storevault.com](http://www.storevault.com) for more information)

- Clients:
- Switch: Gigabit
- Comments: "I have a 40GB Exchange store which is around 20GB most of the time. I had it on local disks and then moved it to an iSCSI LUN on the StoreVault, keeping the logs local (as Microsoft recommends for SAN and Exchange). Performance got faster when I put the stores on iSCSI than when it ran

locally. Microsoft is clear that when you “split up the OS / store” to different devices you will get better performance and it was true for me. The entire store move from start to finish took me about 20 minutes.” – Desmond Fuller

**Scenario 4: Quest, Inc.**

- Clients: Exchange and SQL servers, an Intel Xeon 2.8Ghz and a dual 700Mhz. Both servers run RAID 5 arrays with 10k SCSI drives. The newer server has 3 drives and the older has 6 drives.
- Switch: Gigabit
- “Just some rough info but moving from DAS to the StoreVault cut 1 minute off of the 8GB SQL backup when using the faster server. The StoreVault CPU never got over 50%.” - Jon Eakins

**Scenario 5: A financial analytics company which develops powerful, real-time decision support applications for participants in the financial markets and declined to be named**

- Clients: NFS Debian (Sarge Fast99 2.8.18-2-amd64) on dual Intel 5150 2.66GHz with 8GB RAM.
- Local disk: RAID 1 array with 7200rpm SATAII drives.
- Switch: Gigabit
- Comments: “We have an internal test called FacTask, a mix of jobs that are typical of what we plan to do with the system. The task does a lot of file i/o, mainly opening and closing some small to medium files. The point of test was to throw 8 threads of file requests to the StoreVault over NFS and see what happened.” – IT Administrator

Task	Local Disk	StoreVault S500
	<b>MBps</b>	
Read 60GB file	50.35	45.82
Write 60GB file	19.09	51.38
Copy 60GB file between dirs	23.62	40.85
	<b>Seconds</b>	
gzip and untar a 60GB file containing 116182 files		4099.27
1 copy FacTask	32.3	39.4
2 copies FacTask	33.3	41.1
3 copies FacTask	39.7	45.2
4 copies FacTask	45.1	48.9
20 repeats of 1 copy of FacTask on one node		778.1
20 repeats of 1 copy of FacTask on one each of two nodes		827.6
5 repeats of 4 copies of FacTask on one node		247.7
5 repeats of 4 copies of FacTask on each of two nodes		318.1

**scalability recommendations**

The table below indicates the recommended limits that we are aware of for various tasks.

Variable	Recommended Limit
Number of CIFS Shares	20,000
Number of NFS Exports	10,000
Number of snapshots per volume <sup>1</sup>	255
Number of LUNs	64
Max LUN size	2TB
Number of snapshots per LUN <sup>1</sup>	255
Number of network connections	10,000
Number of open files	200,000
Number of Microsoft Exchange Mailboxes*	1000

\* Tested by an independent third party software vendor using the Microsoft Exchange testing benchmark, MMB-3. For more info on this benchmark refer to <http://www.microsoft.com/exchange/evaluation/performance/mmb3.mspx>

<sup>1</sup> A note on snapshots. The StoreVault supports up to 255 snapshots per volume, so the question that really matters is “How many volumes will I have?”

A StoreVault creates one volume for CIFS shares and another for NFS exports depending on the system configuration. One volume is also created for every LUN, whether iSCSI or Fibre Channel. Therefore, the theoretical maximum numbers of snapshots supported are as follows:

- A total of 510 for NAS data (255\*2 volumes)
- A total of 16320 for SAN data (255\*64 volumes)

Most recommended configurations will never experience this maximum and, obviously, careful thought should be invested in the amount of capacity reserved for this many snapshots should they be implemented!



Network Appliance, Inc.  
495 East Java Drive  
Sunnyvale, CA 94089



Your StoreVault Reseller  
[www.altaware.com](http://www.altaware.com)

[sales@altaware.com](mailto:sales@altaware.com)

(866) 833-4070

For more information on StoreVault, a NetApp division, go to [www.storevault.com](http://www.storevault.com).

© 2006 Network Appliance, Inc. All rights reserved. Specifications subject to change without notice. NetApp and the Network Appliance logo are registered trademarks and Network Appliance, Data ONTAP, Snapshot, SnapVault, and The evolution of storage are trademarks of Network Appliance, Inc. in the U.S. and other countries. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such.