



Exchange Server 2007 Performance Characteristics Using NetApp iSCSI Storage Systems

Brad Garvey, Network Appliance, Inc.

March 2007 | TR-3565



Table of Contents

1. Introduction	3
1.1 Purpose and Scope	3
2. Exchange Server Performance	3
2.1 Exchange 2003 Workload Characteristics	3
2.2 Exchange 2007 Workload Characteristics	4
3. Microsoft Exchange Server 2007 on NetApp iSCSI/IP SAN Storage Systems	4
3.1 Improved Performance and Manageability	4
4. Configuration Overview	5
4.1 Test Environment	5
4.2 Exchange Storage Group and Storage Layout	6
5. Test Results: Exchange 2003 and Exchange 2007 Performance Comparison	6
5.1 Memory Utilization	6
5.2 Database Cache	7
5.3 Read/Write Comparison	8
5.4 Read:Write Ratio per User	8
5.4 Effects of I/O Reduction on Total iSCSI IOPS	9
6. Online Versus Cached Mode Performance Characteristics	11
6.1 LoadGen Tests	11
6.2 LoadGen Configuration	11
6.3 LoadGen Results	12
7.0 Summary	14

1. Introduction

Corporate messaging has become mission critical in nearly all organizations. This change in priority for messaging has imposed additional constraints on information technology departments. Not only do these departments have to deal with requests for increased mailbox sizes and additional users, they need to ensure high availability of their messaging infrastructure while providing longer online hours, which in turn reduces available maintenance windows. Enterprises are struggling to keep up with the complexity and costs of these deployments. Microsoft® Exchange Server 2007 provides increased availability options and performance improvements that address some of these messaging issues. However, as user mailboxes become larger, Exchange Server performance must be balanced with storage capacity to achieve acceptable cost/performance metrics. This paper examines some of the changes in Exchange Server 2007 and compares them to Exchange Server 2003. In particular, the use of iSCSI as a storage interconnect is examined with a view to determining the role it can play in Exchange Server storage availability and performance.

1.1 Purpose and Scope

This document has two main purposes:

- To provide metrics illustrating the workload changes in Exchange Server 2007 versus Exchange Server 2003
- To demonstrate how NetApp IP SAN solutions provide high-performance and cost-effective solutions for Exchange Server 2007 without sacrificing flexible infrastructure and data management options that are critical to Exchange Server deployments

2. Exchange Server Performance

2.1 Exchange 2003 Workload Characteristics

Exchange Servers generate I/O to the storage subsystem each time information is read from or written to disks. The amount of I/O that the Exchange Server generates and the amount of I/O the storage subsystem can sustain directly affects Exchange performance and usability.

Exchange Server generates I/O to the storage subsystem through a number of different processes. The majority of the I/O transactions are composed of reads and writes to database and transaction log files. Several factors contribute to the amount of I/O placed on the storage subsystem by Exchange Servers.

- Number of Exchange users
- Average mailbox size
- Number of storage groups
- Outlook client online versus cached mode

As these items increase, the number of I/Os to the storage subsystem also increases. Due to limited application memory available in the Windows™ 32-bit architecture, Exchange Server 2003 has a limited amount of database cache. This results in a limited number of read hits from database cache, which in turn causes more read I/Os to the storage subsystem. For more information on Exchange Server 2003 workloads, refer to [Optimizing Storage for Exchange Server 2003](#) on the Microsoft Web site.

The use of Outlook client online mode also causes an increase in I/Os because additional process requirements are placed on the server. When using cache mode, operations such as searches, indexing, and repeated access to e-mail are processed on the client, thus reducing the amount of I/O placed in the

Exchange Server. When running in online mode, these processes are pushed to the server, causing additional I/O on the Exchange Server.

2.2 Exchange 2007 Workload Characteristics

Exchange Server 2007 provides significant performance improvements and enables larger mailbox sizes and more users per server. This is accomplished by reducing the total amount of I/O generated by Exchange workloads. The three main factors responsible for the reduction in I/O in Exchange Server 2007 are:

- Additional application memory available
- Increased database page size
- Reduction of database reads

Additional application memory available

In Exchange Server 2003, available memory is limited by the 32-bit architecture. Exchange Server 2007 is based on 64-bit (x64) architecture. This change enables Exchange Server 2007 to access larger amounts of memory which in turn increases the amount of database cache available to the Exchange information store. Exchange Server 2007 database cache is capable of going beyond the previous 900MB limit to potentially multiple gigabytes, depending on the amount of server system memory. This allows for a great reduction in read I/O, which in turn causes a decrease in IOPS per user. Additional database cache results in a decrease in total IOPS per mailbox.

Increased database page size

The database page size for Exchange Server 2007 has increased from 4kb to 8kb, which enables larger amounts of data to be read from or written to disk. This decreases the number of times the Exchange Server has to go to the disk for information.

Reduction of database reads

In previous versions of Exchange, a read:write ratio of 2:1 or 3:1 was typical. With a 64-bit architecture and a larger page size, Exchange Server 2007 is able to decrease the read:write ratio to approximately 1:1. The reduction in reads helps decrease the overall IOPS.

3. Microsoft Exchange Server 2007 on NetApp iSCSI/IP SAN Storage Systems

NetApp storage solutions for Exchange Server 2007 offer superior performance and scalability with the ability to reconfigure and expand storage, and without the need for costly maintenance downtime.

By enabling storage consolidation and centralized administration, NetApp IP SAN solutions enable organizations to improve Exchange Server data accessibility and availability and provide greater application uptime than direct-attached storage, at a much lower total cost of ownership.

3.1 Improved Performance and Manageability

NetApp storage systems make it possible to optimize storage utilization by partitioning NetApp storage arrays into logical units called FlexVol™ volumes, which make it possible to manage data without the need

to assign physical drives to volumes. These volumes all share performance benefits from a larger pool of drives called an *aggregate*, resulting in the following benefits for Microsoft Exchange Server 2007 environments:

- Using FlexVol volumes, the capacity and performance bandwidth of a large collection of fast drives can be available to the volume. Even small FlexVol volumes benefit from the number of drives in the aggregate.
- All FlexVol volumes can be managed independently and maintain individual Snapshot™ copies and schedules, providing greater flexibility in recovery options.

4. Configuration Overview

To establish a comparison between Exchange Server 2003 and Exchange Server 2007 performance, a baseline test was first run with JetStress 2004, then with JetStress 2007. When all tests were complete, the results were compared to establish differences between Exchange Server 2003 and Exchange Server 2007 performance characteristics.

To determine the I/O reduction between online and cached mode in Exchange Server 2007, Microsoft LoadGen was used to simulate Outlook 2007 client traffic in both online and cached modes.

JetStress is Microsoft's recommended tool for simulating Exchange Server-like workloads to test the performance and stability of disk subsystems by simulating Exchange disk I/O load. JetStress 2004 was used to simulate Exchange 2003 performance characteristics. JetStress 2007 was used to simulate Exchange 2007 performance characteristics. For additional information on Microsoft JetStress, visit <http://www.microsoft.com/technet/prodtechnol/exchange/2007/downloads/tools.msp>.

Microsoft Exchange Load Generator (LoadGen) as a simulation tool used to measure the impact of MAPI clients on Exchange Servers. LoadGen simulates MAPI workloads over a period of time to measure an Exchange Servers responds to e-mail loads. LoadGen is a useful tool for administrators who are sizing servers and validating a deployment plan. Specifically, LoadGen helps determine whether servers can handle the load they are to carry.

4.1 Test Environment

This section details server and storage configurations used in the test environment.

Storage

- NetApp FAS3050C storage system, composed of two FAS3050 storage controllers configured as an active/active cluster
- Data ONTAP® 7.1
- A total of 12 disk shelves (DS14 MK2) were used (6 per storage controller)
- 15K 144GB drives (see Table 1 for total number of disks used in each configuration)

Servers

- 1 Hewlett Packard Proliant DL 385 with 2 dual-core AMD Opteron processors and 8GB RAM
- Windows Server 2003 Enterprise Edition (32-bit) service pack 1 for Exchange Server 2003 testing
- Windows Server 2003 Enterprise Edition (64-bit) service pack 1 for Exchange Server 2007 testing
- NetApp SnapDrive® for Windows 4.1 with MPIO (NetApp multipathing I/O solution)
- 1 GbE connection per server
- Microsoft iSCSI Software Initiator v2.02

Network

- 1 NetGear GSM7224 GbE Ethernet switch (Jumbo Frames enabled)

4.2 Exchange Storage Group and Storage Layout

Each user count configuration for JetStress 2003 was sized according to NetApp best practices. Each FAS3050 storage controller was configured with separate aggregates for database and transaction log volumes. Each database aggregate on each storage controller contained two NetApp FlexVol volumes (one volume for each JetStress Storage Group). Each volume contained one LUN, on which JetStress databases were located. The transaction log aggregate on each storage controller contained one FlexVol volume with four LUNs (one LUN per storage group).

The following table shows the storage configuration used in each test:

Aggregate	# of Drives in Aggregate	Volume	Total Space Available	Used Space
DBAGGR	39	DB	3600GB	2880GB
LOGAGGR	9	LOGVOL	340GB	272GB

Table 1) Storage Layout

5. Test Results: Exchange 2003 and Exchange 2007 Performance Comparison

For each user count and configuration in the preceding sections, the following counters were compared for both Exchange 2003 and Exchange 2007 workloads:

- Memory utilization
- Database cache
- Host-side read:write ratio
- Storage system read:write ratio

5.1 Memory Utilization

Figure 1 compares available system memory between Exchange 2003 and Exchange 2007. The 4GB limitation imposed by 32-bit platforms limits Exchange 2003 to only 2GB for application memory or 3GB when utilizing a 3GB switch. The x64 platform gives Exchange 2007 access to memory greater than the previous 4GB limit.

The data shows that Exchange 2003 utilized only 1.5GB of memory, leaving 6.5GB of memory available. The 64-bit platform running the Exchange 2007 workload was able to utilize additional memory for database cache, leaving only 2.4GB of available system memory.




Figure 1) Available system memory.

NOTE: As Exchange workloads increase, additional memory can be utilized, increasing the total amount of database cache available.

5.2 Database Cache

As a result of larger amounts of memory available in the x64 platform, the Exchange Server 2007 database cache can be larger. This can reduce the number of reads and writes to disk. In Exchange Server 2003, the maximum was 900MB of database cache. With additional memory available in Exchange Server 2007, this amount can increase from 900MB to multiple GBs, depending on the amount of server system memory installed. As additional memory is added to the host server, the Exchange Server database cache can increase, further decreasing the number of disk reads.

Figure 2 compares database cache sizes between Exchange 2003 and Exchange 2007. Due to additional memory in Exchange 2007, the database cache size increased from 900MB to 2438MB in Exchange 2007.

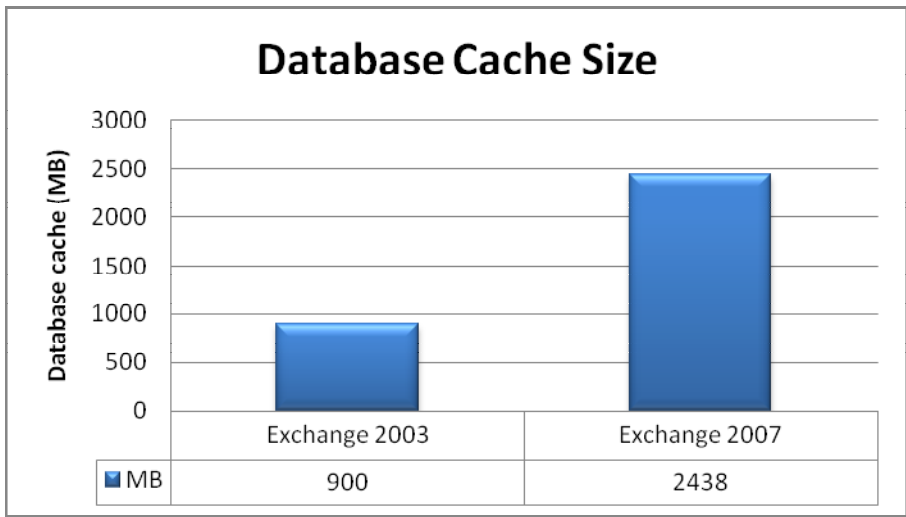


Figure 2) Database cache size.

5.3 Read/Write Comparison

With Exchange Server 2003, a read:write ratio of 2:1 was typical. With Exchange Server 2007, additional database cache, coupled with an increase in database page size from 4KB to 8KB, reduces the number of database reads to the storage subsystem. Following Microsoft-recommended guidelines for Exchange Server 2007, the read:write ratio can be reduced to approximately 1:1.

Figure 3 compares the read:write ratios for Exchange 2003 and Exchange 2007 workloads as generated by JetStress testing. The number of reads decreased from 2,897 average reads/sec in Exchange 2003 to approximately 641 average reads/sec in Exchange 2007. The reduction of reads in Exchange 2007 shows the decrease in the database read:write ratio from 2.25:1 to approximately 1.24:1.

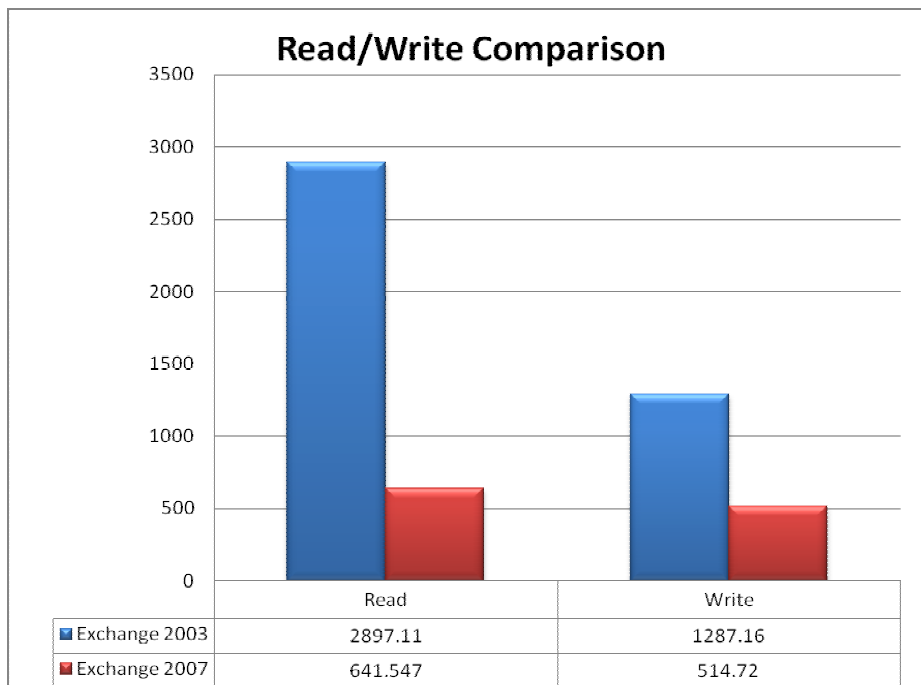


Figure 3) Read/write comparison.

5.4 Read:Write Ratio per User

The read/write comparison can be further examined as the database read:write ratio per user. Figure 4 displays the previous data as the average number of database reads and writes per user. In the Exchange Server 2007 environment, the number of database reads decreased from 0.72 to 0.16 reads per user.




Figure 4) Read:write ratio per user.

5.4 Effects of I/O Reduction on Total iSCSI IOPS

Figures 5 and 6 show that the number of reads over the iSCSI interconnect decreased from approximately 12,000 kb/sec from Exchange 2003 to only 6,000 kb/sec with the Exchange 2007 workload. This is a 50% reduction in read traffic to the storage subsystem. This reduction in IOPS enables the storage system to accommodate the larger mailbox sizes and additional storage groups available in Exchange Server 2007.

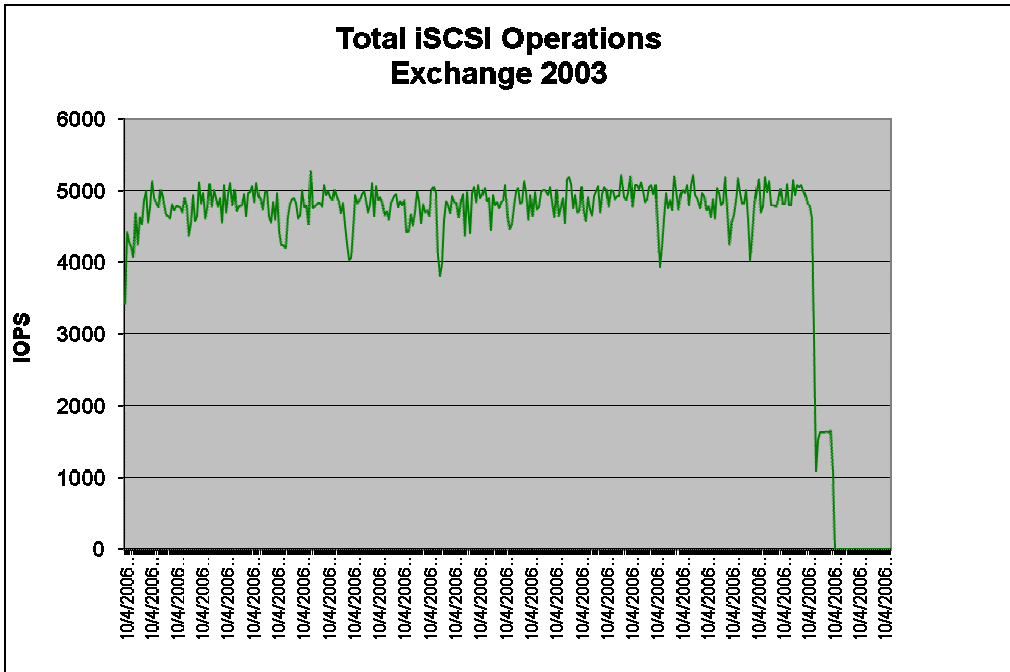


Figure 5) Total iSCSI operations, Exchange 2003.




Figure 6) Total iSCSI operations, Exchange 2007.

6. Online Versus Cached Mode Performance Characteristics

Factors aside from those discussed can cause increased I/O in an Exchange Server implementation. One such factor is the mode in which Outlook clients are configured.

Outlook clients can be connected to the Exchange Server in one of two ways: online mode or cached mode. Online mode requires a constant connection with the server. All transactions take place on the server, which increases the workload on the Exchange Server to service user transactions.

With cached mode operation, messages are copied from the server to the Outlook client. When these messages have been cached locally, typical user actions such as reopening the message do not require interaction with the server. Only small amounts of data are synchronized with the server in the background to keep the client information up to date. This reduces the load placed on the server.

To further determine how online versus cached mode operation affects I/O as well as other workload characteristics of Exchange 2007, LoadGen was used to simulate user traffic over a period of time.

6.1 LoadGen Tests

LoadGen is a load simulation tool from Microsoft that simulates the impact of MAPI-based clients on an Exchange Server. For additional information on LoadGen, visit <http://www.microsoft.com/downloads/details.aspx?familyid=DDEC1642-F6E3-4D66-A82F-8D3062C6FA98&displaylang=en>.

The following tests were performed to get an accurate representation of Exchange Server 2007 performance characteristics under different circumstances:

- 4000 Users 8GB system memory – 10-hour Outlook 2007 Online mode test
- 4000 Users 8GB system memory – 10-hour Outlook 2007 Cached mode test

- 4000 Users 16GB system memory – 10-hour Outlook 2007 Online mode test
- 4000 Users 16GB system memory – 10-hour Outlook 2007 Cached mode test

6.2 LoadGen Configuration

The following configuration was used during LoadGen testing:

Exchange Server 2007

- 4000 mailboxes
- 4 storage groups
- 1 mailbox database per storage group

Mailbox Characteristics

- 1000 users per mailbox database
- Average mailbox size of 225MB

6.3 LoadGen Results

Read:write ratio and user I/O summary

Results collected using Windows Performance Monitor (PERFMON) from online mode and cached mode tests were compared to determine differences in I/O between online and cached modes. Table 2 compares read:write ratio and user IOPS for online and cached mode workloads. When performing the test in online mode, the read:write activity increased from a ratio of 1.63 to 3.42. The total user IOPS per user also increased from 0.07 to 0.15 when running in online mode.

	Online mode	Cached mode
Read/sec	477.74	190.547
Write/sec	144.682	117.083
IOPS	622.422	307.63
R:W	3.424545	1.632587
IOPS/user	0.155606	0.076908

Table 2 (Online vs. Cached mode comparison)

Figure 7 shows the overall reduction in IOPS from 622.422 to 307.63 when using online mode. The reduction in overall IOPS in this comparison is due to the 42% reduction in reads. The reduction is seen in both IOPS per user and read:write ratio when using cached mode.

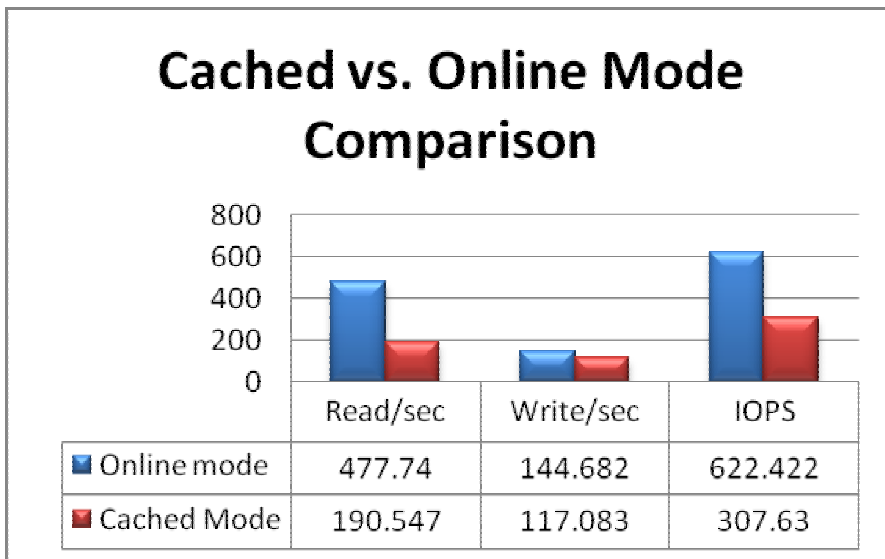


Figure 7) Cached versus online mode comparison.

Database memory cache

Increasing the amount of database cache also affects IOPS. When testing 4000 users with 8GB of server memory, database cache is limited to 1.48MB/user $[(8GB - 2048) / 4000]$. Increasing the amount of server memory to 16GB increases the amount of database cache to 3.48MB/user $[(16GB - 2048) / 4000]$.

Figure 8 shows the amount of reads and writes to the storage subsystem with different amounts of host-side memory. When host-side memory was increased from 8GB to 16GB, the amount of database cache was increased from 1.48 MB/user to 3.48 MB/user. With additional database cache available, reads/sec decreased from 200.414 to 190.923. Average writes/sec decreased from 166.024 to 138.592. Increasing database cache also decreased total IOPS from 366.434 to 329.515 in the online mode comparison.




Figure 8) Memory comparison in online mode.

7.0 Summary

Organizations deploying Exchange Server 2007 in their IT infrastructure can use the data in this document for insight into how Exchange Server 2007 performs when deployed in NetApp iSCSI environments.

The key point that emerges from the data presented is that the iSCSI storage networking protocol is a viable and high-performing storage interconnect for Exchange Server 2007 environments. One of the most notable data points from the test data shows the decrease in total iSCSI IOPS. In the Exchange 2007 workload, total iSCSI IOPS were 50% less than in the Exchange Server 2003 environment. In all tests performed over iSCSI, the storage interconnect provided sufficient bandwidth to serve I/O requirements of the Exchange mailboxes while taking advantage of the performance benefits in the Exchange Server 2007 environment.

By utilizing NetApp storage solutions, organizations deploying Exchange Server 2007 are positioned to provide unparalleled flexibility and scalability to enterprise-class Exchange Server deployments. NetApp IP SAN solutions enable customers to decrease TCO without sacrificing the performance, flexibility, and management requirements of enterprise-level Exchange Server deployments.

© 2007 Network Appliance, Inc. All rights reserved. Specifications subject to change without notice. NetApp, the Network Appliance logo, Data ONTAP, and SnapDrive are registered trademarks and Network Appliance, FlexVol, and Snapshot are trademarks of Network Appliance, Inc. in the U.S. and other countries. Microsoft is a registered trademark and Windows is a trademark of Microsoft Corporation. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such.