



Technical report:
**Exchange Server 2007 with IBM
System Storage N series with
SnapManager for Exchange**

A performance-impact analysis

• • • • • • • • •

Document NS3598-0

June 22, 2008



Table of contents

Abstract	3
Introduction	3
Purpose and scope	3
Intended audience	3
Snapshot overview	4
Components of a VSS snapshot.....	4
VSS backup methods	5
Clone shadow copies	5
Snapshot shadow copies	5
Benefits of snapshot backups	5
Considerations for using copy-on-write snapshot backups.....	6
IBM Snapshot technology.....	6
How IBM Snapshot technology differs from other snapshot solutions.....	6
SnapManager for Exchange	7
Test overview	7
Exchange storage layout	8
SnapManager for Exchange configuration	8
LoadGen configuration	9
Results	9
Baseline performance characteristics.....	9
RPC-averaged latency	10
Database-read latency	11
CPU utilization.....	12
Performance impact of IBM Snapshot technology	13
RPC latency with snapshots	13
Database-read latency with snapshots.....	13
CPU utilization with snapshots	15
Performance impact of deleting snapshots	16
Averaged RPC latency with snapshot deletion.....	16
Database-read latencies with snapshot deletion	17
CPU utilization with snapshot deletion.....	19
Summary	20
Trademarks and special notices	21



Abstract

This paper documents performance characteristics for Microsoft Exchange Server 2007 while IBM System Storage N series with SnapManager for Exchange performs backups using highly efficient IBM System Storage N series with Snapshot technology.

Introduction

Microsoft® Exchange Server has become a vital source of communication as well as a repository for business-critical data. Because of the increasing importance of Microsoft Exchange Server, data protection, disaster recovery, and high availability play a critical role in Exchange Server deployments. As the importance of Exchange Server within the organization increases, companies expect quick recovery times with little or no data loss. With Exchange databases growing rapidly in size every day, it is increasingly difficult to complete time-consuming backup operations in a reasonable amount of time. When an outage occurs, it can take days to restore service from slower media such as tape, even assuming that all of the backup tapes are available and error free.

IBM® offers a comprehensive suite of data-management solutions that enable organizations to keep pace with the increasing data availability demands of an ever-expanding Exchange Server environment, as well as scale to accommodate future needs while reducing cost and complexity.

Purpose and scope

Since there are many different solutions that leverage Microsoft Volume Shadow Copy Services (VSS) for Exchange Server, it is important not only to understand the implementation but also to test performance, capacity and recovery implications of your solution to ensure that it meets your corporate messaging needs and meets service-level agreement (SLA) requirements.

The purpose of this paper is to present the performance implications of an IBM System Storage™ N series with SnapManager® for Exchange (SME) solution in an Exchange Server 2007 environment.

Intended audience

The audience of this paper should be familiar with the following products:

- Exchange Server backup and recovery solutions
- IBM N series storage solutions
- Microsoft VSS technology.



Snapshot overview

Exchange Server 2007 supports Microsoft VSS, which supports shadow copies that can be executed via a software or hardware provider. Generally, it takes much less time to restore a backup that was taken using a hardware-based snapshot. Therefore, depending on the hardware and software that was used in the solution, restoring Exchange data from these types of backups may make it easier to meet the requirements of an Exchange database SLA. Because larger databases can be restored more quickly, hardware-based snapshot restores can help support larger database sizes and still meet SLAs for restoring Exchange data.

VSS in Microsoft Windows Server 2003 provides an infrastructure that enables third-party storage management programs and hardware providers to create and manage shadow copies of Microsoft Exchange Server utilizing the VSS application plug in (API).

VSS coordinates communication between *requestors* (backup applications), *writers* (applications in Windows services like Exchange Server 2007), and *providers* (system, software or hardware components that create the shadow copies, such as IBM System Storage). To use the VSS feature to backup Exchange Server 2007, the backup program must include an Exchange Server 2007-aware VSS requestor. Because the backup program that is bundled with Windows Server has no such requestor, organizations must use third-party backup applications.

Components of a VSS snapshot

VSS ensures the consistency and integrity of a snapshot (which has been created by IBM System Storage N series with Snapshot™ technology) by orchestrating freeze/thaw, flushing outstanding input/output (I/O) to disk, and ensuring that all related logical unit numbers (LUNs) are brought together into a single, consistent snapshot.

There are four key components present in the snapshot process:

- VSS:** VSS is the Windows Server 2003 service that coordinates among other VSS components to provide snapshot-based backups and restores.
- VSS Requestor:** The VSS requestor is a backup application, such as IBM N series with SME. It initiates, or “requests” VSS backup and restore operations. The requestor also specifies snapshot attributes for the backups it initiates.
- VSS Writer:** The VSS writer owns and manages the data to be captured in the snapshot. In this case, the VSS writer is Exchange Server 2007.
- VSS Provider:** The VSS provider integrates storage array-specific snapshot and cloning functionality into the VSS framework. A provider can be either a hardware or a software provider. In this case, the IBM hardware provider is installed as part of the IBM System Storage N series with SnapDrive® software installation.



When the Exchange server receives a backup request, the Exchange writer prepares Exchange databases for backup. The writer does this by suspending all disk write I/O to the databases for up to 20 seconds. This is referred to as “freezing the databases.” The provider must be able to complete the shadow copy within this window or the backup will be aborted. After the backup is completed, the writer *thaws* the databases and resumes regular I/O operations. For more information about requestors, writers, and providers, see Microsoft technical support concerning basic VSS concepts; at the time of this report writing, information regarding shadow copies and shadow copy sets could be found at <http://msdn.microsoft.com/en-us/library/aa384612.aspx>. For additional information on VSS implementation in Exchange Server, see the Microsoft technical support site also; at the time of this report writing, information could be found at <http://www.microsoft.com/technet/prodtechnol/exchange/2003/vssbp.msp>.

Each VSS-compliant backup must have achieved the following:

- Exchange databases and transaction logs must be backed up through the Exchange VSS writer
- The backup application must validate the integrity of the backup set
- Restores must be done through the Exchange VSS writer.

VSS backup methods

The following is an overview of VSS backup methods and issues: clone shadow copies, snapshot shadow copies, benefits of snapshot backups, and considerations for using VSS backups.

Clone shadow copies

A “clone shadow copy” is a full copy of the volumes in a shadow copy set. A clone is a completely independent copy of the original data set.

Unlike a clone copy, a shadow copy is not an independent copy of the original data. A *shadow copy set* is a group of volume shadow copies that are synchronized at the same point in time.

Snapshot shadow copies

Generically, snapshots are created by defining a marker at a point in time and making sure that the data can be rolled back to that point in time. Multiple snapshots can be kept, and snapshots typically require much less additional disk space than clones.

Snapshots can be created in several different ways. The most common method is called *copy-on-write*.

The copy-on-write method defines a snapshot at a point in time, and then monitors the original dataset for changes. If a change is made to a block, the original contents of the block are copied to a separate location before the block is updated. Over time, the size of a snapshot can continue to grow, especially when a snapshot is made of a quickly-changing dataset.

Benefits of snapshot backups

VSS snapshots provide the ability to backup large amounts of data in a backup window with faster restore times.



Considerations for using copy-on-write snapshot backups

During the snapshot process, performance degradation can occur. As additional snapshots are taken and kept available, performance degradation can occur. And, deleting snapshot copies can negatively impact performance.

IBM Snapshot technology

IBM N series with Snapshot technology is unique in that it has all the benefits of snapshot technology in general but without additional performance overhead.

IBM with N series Snapshot technology is a feature of the IBM System Storage N series with WAFL® (Write Anywhere File Layout) storage virtualization technology, a component of IBM System Storage N series with Data ONTAP®. An IBM N series snapshot is a frozen, read-only view of a WAFL volume that provides easy access to old versions of files, directory hierarchies, and/or LUNs.

IBM Snapshot technology is very efficient as it only stores block-level changes between each incremental snapshot. There can be up to 255 snapshots per file system with little to no impact on performance.

How IBM Snapshot technology differs from other snapshot solutions

Most snapshot shadow copy implementations use a method called “copy-out snapshots.” The principle behind copy-out snapshots is to preserve snapshot data by relocating it from an active file system or LUN to a “special” pre-allocated copy-out area. After a snapshot is taken, when a data block in the active file system or LUN is about to be overwritten by new data, the following steps are taken:

1. The original data (that is about to be overwritten) is read from the underlying disks.
2. The original data is written to the copy-out area in order to preserve it.
3. The new data is written over the original data on the disks.

This process triples the I/O burden when data is being copied out (each new write operation becomes a read operation and two write operations) and typically results in a performance penalty of as much as 30-50%.

IBM N series SME utilizes IBM N series Snapshot technology as a basis for snapshots. IBM N series Snapshot technology differs from other vendors’ snapshot technology as it is not a copy-on-first-write implementation. This allows IBM N series Snapshot technology to store up to 255 snapshots per volume without incurring additional performance overhead.

When a snapshot is created on an IBM storage system it does not copy changed blocks. It instead identifies all blocks as belonging to both the active file system and the snapshot. As data changes, data is written to new disk blocks. This creates two copies of the changed data: the newly written copy in the active file system and the old data in the snapshot copy.

This enables the IBM N series storage system to create near-instantaneous snapshot copies without incurring additional performance overhead.



SnapManager for Exchange

IBM N series SME 4.0 software is available for Microsoft Exchange Server 2003 and 2007. SME 4.0 has achieved a “[Certified for Windows](#)” logo for Windows® Server 2003 Standard and Enterprise editions. SME is also a “Microsoft [SimpleSAN](#)” designated and a Windows Server 2003 certified backup and recovery solution for Exchange Server. SME is tightly integrated with Microsoft Exchange, which allows for consistent online backups of Exchange environments while leveraging IBM N series with Snapshot copy technology. SME 4.0 supports Microsoft VSS, Virtual Disk Services (VDS), and other Microsoft storage technologies. (For details on VSS, see [Microsoft knowledgebase, including article 822896, which, at the time of this report writing, was located at http://support.microsoft.com/kb/822896](#).)

SME is a VSS (snapshot copy) requestor, which means that it uses the Microsoft-supported VSS subsystem to initiate backups. SME provides a complementary feature set for the new Microsoft Exchange 2007 data replication features. SME works with Local Continuous Replication (LCR) and Cluster Continuous Replication (CCR) replica databases and provides a rich feature set to leverage these new technologies.

SME benefits at a glance:

- Quickly backup and restore entire Exchange Storage Groups using IBM highly efficient N series with Snapshot technology.
- Exchange data stored in snapshots can be mirrored to one or more locations for archival or disaster recovery purposes.

Test overview

To validate the use of IBM N series Snapshots technology in an Exchange Server 2007 environment, a test scenario was built using Microsoft Exchange Load Generator (LoadGen) to simulate a 3,000-user environment. LoadGen is a simulation tool that tests Exchange servers to see how they respond to email loads.

Multiple runs of LoadGen testing were completed during the testing periods. The first day was run without any scheduled snapshots for baseline performance. Over subsequent days, SME snapshots were scheduled every two hours with 24 snapshots retained. Snapshots were verified nightly using a separate verification server.

During the simulation, data was collected using Windows Performance Monitor (perfmon). Upon completion of the test, the data was analyzed to determine affects of snapshots on server performance and to ensure that server response times remained within acceptable limits.

The following sections detail server and storage configurations used in the test environment.

Storage:

(1) IBM System Storage N series N5500 storage system

Composed of 2 N5500 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 144 GB 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 8 GB RAM
- Windows Server 2003 Enterprise Edition (64-bit) Service Pack 1
- (2)1 GbE iSCSI connections per server
- Intel Pro/1000 MT driver 8.1.8.0

Software:

- Microsoft iSCSI Software Initiator v2.02
- Microsoft Exchange Server 2007
- IBM System Storage N series with SnapDrive for Windows 4.1 with multipathing I/O (MPIO) solution
- IBM System Storage N series with SME 4.0

Network:

- (1) NetGear GSM7224 GbE Ethernet switch (jumbo frames enabled)

Exchange storage layout

Storage layout was sized according to IBM best practices. The N5500 storage controller was configured with separate aggregates for database and transaction log volumes. Each database aggregate on the storage controller contained two IBM System Storage N series with FlexVol™ volumes (one volume for each storage group). Both volumes were configured as RAID-DP (double parity) volumes. Each volume contained one LUN on which mailbox databases were located. The transaction log aggregate on each storage controller contained one FlexVol volume with three 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	3600 GB	2880 GB
LOGAGGR	9	LOGVOL	340 GB	272 GB

Table 1) Storage layout

SnapManager for Exchange configuration

SME backups were scheduled every two hours during the test, providing a recovery point of two hours. The backup job was configured to complete a full backup of all storage groups and mailbox databases. A retention period of seven days was maintained. A separate verification server was set up in order to isolate database checksum I/O from the active server.

Each SME backup was configured to do a full backup of all three storage groups, with up to the minute recovery capability. When up to the minute recovery is used, transaction logs corresponding to the current backup set are moved to the SnapInfo location. In addition to corresponding transaction log files the SnapInfo directory contains information regarding each backup taken. The SnapInfo directory is contained on the same volume as the transaction logs.



When the transaction log directory resides on the same NT file system (NTFS) volume as the SnapInfo directory, SME will utilize NTFS hard links to avoid the file copy operation and maximize storage utilization. During the backup process, SME will edit the file pointers for a given transaction log. SME inserts a pointer to the SnapInfo directory and removes the pointer to the Exchange transaction log directory. As a result, it will appear as if a transaction log was moved from one directory to the other. However, a physical move operation was not performed, and no I/O overhead was incurred on the Exchange server. Only the pointers for a transaction log file were updated, thus saving the additional space required to copy the transaction log files into the SnapInfo directory.

LoadGen configuration

The following configuration was used during LoadGen testing:

Exchange Server 2007:

- 3,000 mailboxes
- 3 Storage Groups
- 1 Mailbox database per storage group

Mailbox characteristics:

- 1,000 Users/Mailbox database
- Average mailbox size of 225MB
- Heavy mailbox usage profile

For more detailed information mailbox profiles, please refer to the Microsoft technical site; at the time of this report writing, information is located at:

<http://www.microsoft.com/technet/prodtechnol/exchange/2007/plan/hardware.mspx>.

Results

Analysis of the data collected shows the following performance characteristics.

Baseline performance characteristics

In order to establish baseline performance characteristics, Exchange server performance while under LoadGen testing was observed without performing any snapshot backups. LoadGen tests were run for 10 hours. In order to gather consistent data points, the following charts display the last four hours of the test period.



RPC-averaged latency

Average remote procedure call (RPC) latency is a measure of the average amount of time it takes for a server to respond to a client request. Under normal circumstances, this number should be below 50 milliseconds (ms). RPC latencies above 50 ms can cause Outlook clients to temporarily lose connectivity to the Exchange Server. The RPC-averaged latency during the baseline testing period was nine ms.

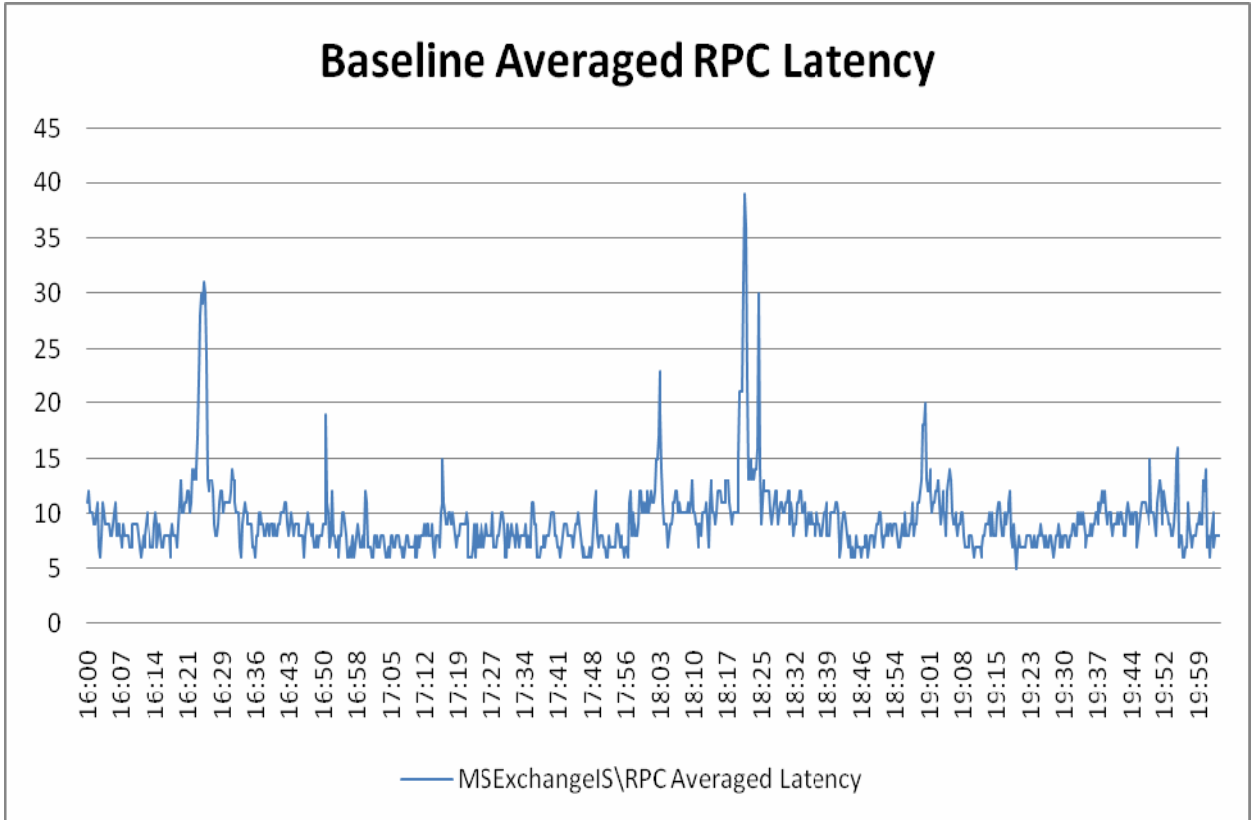


Figure 1: Average RPC latency for MS Exchange Server.



Database-read latency

In an Exchange Server environment, database read latency equates to degraded system performance. Microsoft recommends that average read latency for Exchange 2007 servers remain under 20 ms, and that latency spikes over 50 ms should not last more than a few seconds.

The chart below shows the average database read latency of the database LUNs during the 10-hour test period. Average database read latency on the database LUNs was 8 ms on the G drive and 9 ms on the R drive. Graph of average read latency to ensure performance is within acceptable limits during the entire test.

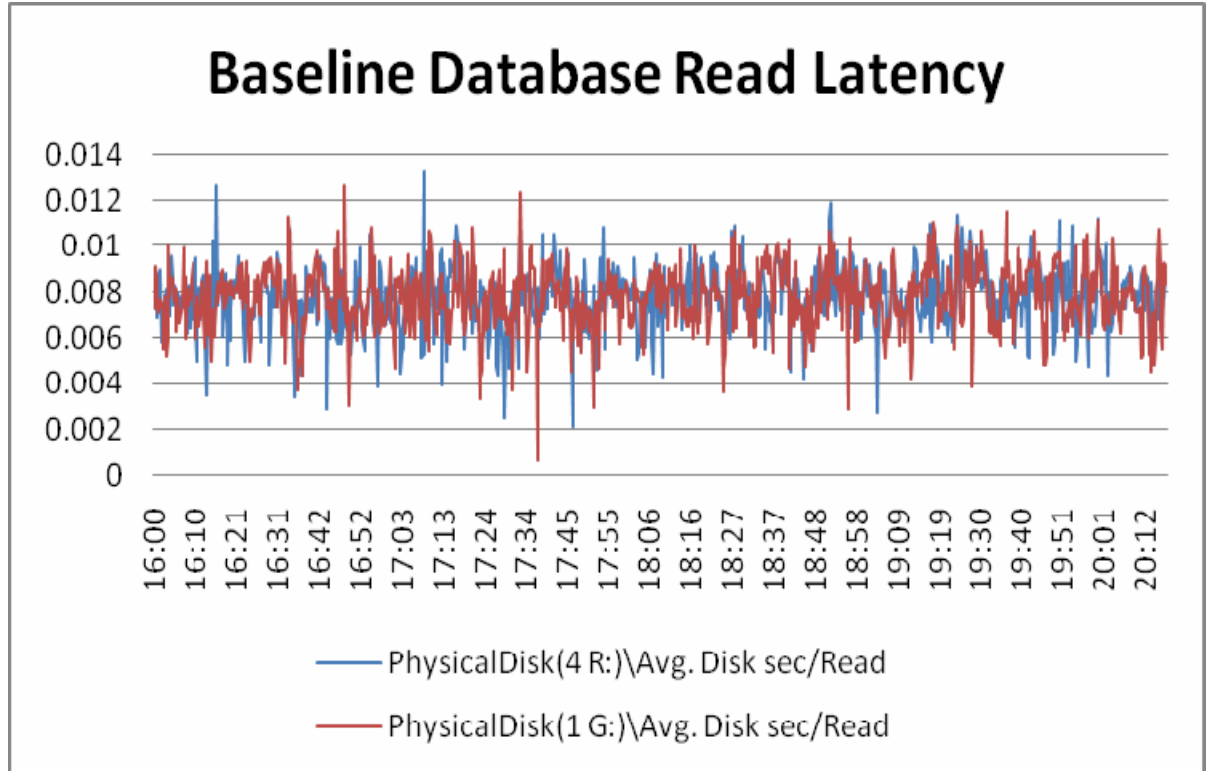


Figure 2: Baseline Microsoft Exchange database read latency.



CPU utilization

For an optimal Exchange Server environment, central processing unit (CPU) utilization should average below 70%. The following chart shows CPU utilization during a baseline run of 10 hours. During the baseline run, with no scheduled snapshots, the average CPU utilization was 19%.

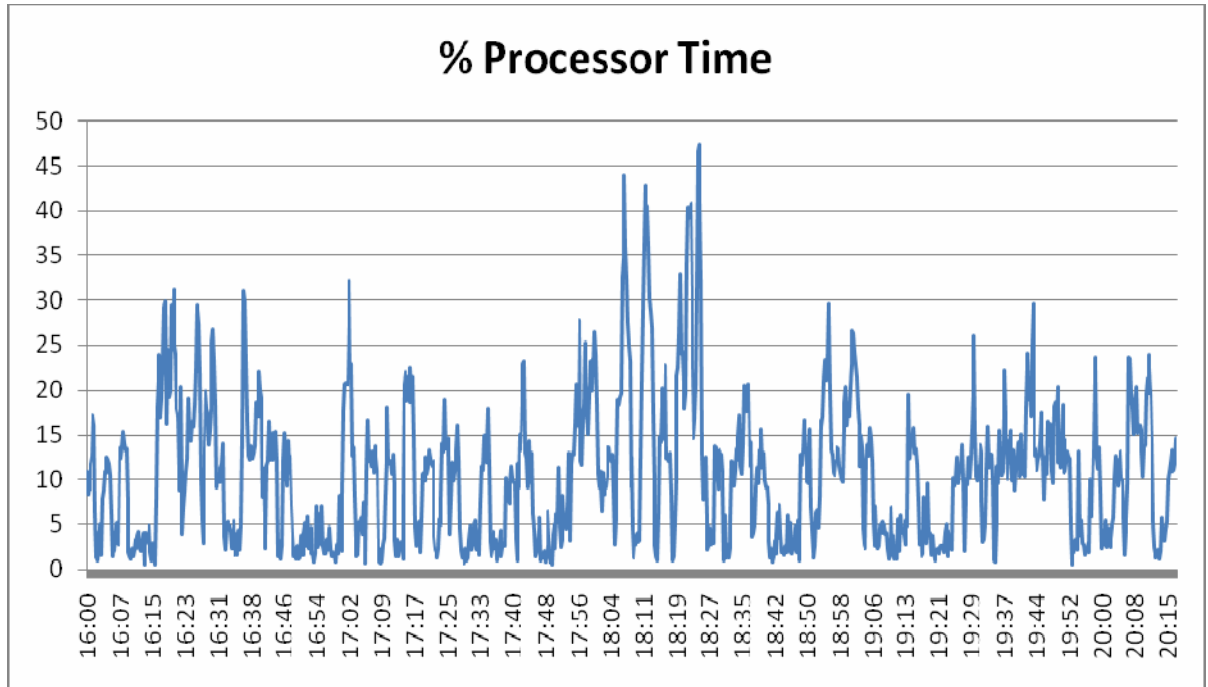


Figure 3: CPU utilization.



Performance impact of IBM Snapshot technology

The following charts show the impact IBM N series Snapshot technology has on Exchange Server performance. During the 10-hour test period, SME was scheduled to take a snapshot backup every two hours with 24 snapshots kept online. The shaded areas in the charts below indicate periods where snapshots took place.

NOTE: Verification times and performance were not a part of the testing performed.

RPC latency with snapshots

The following chart shows RPC average latency with snapshots occurring every two hours. No significant changes were apparent in averaged RPC latency numbers. With snapshot backups occurring every two hours, averaged RPC latency is still well within acceptable ranges, with no sustained spikes occurring during the snapshot backup process.

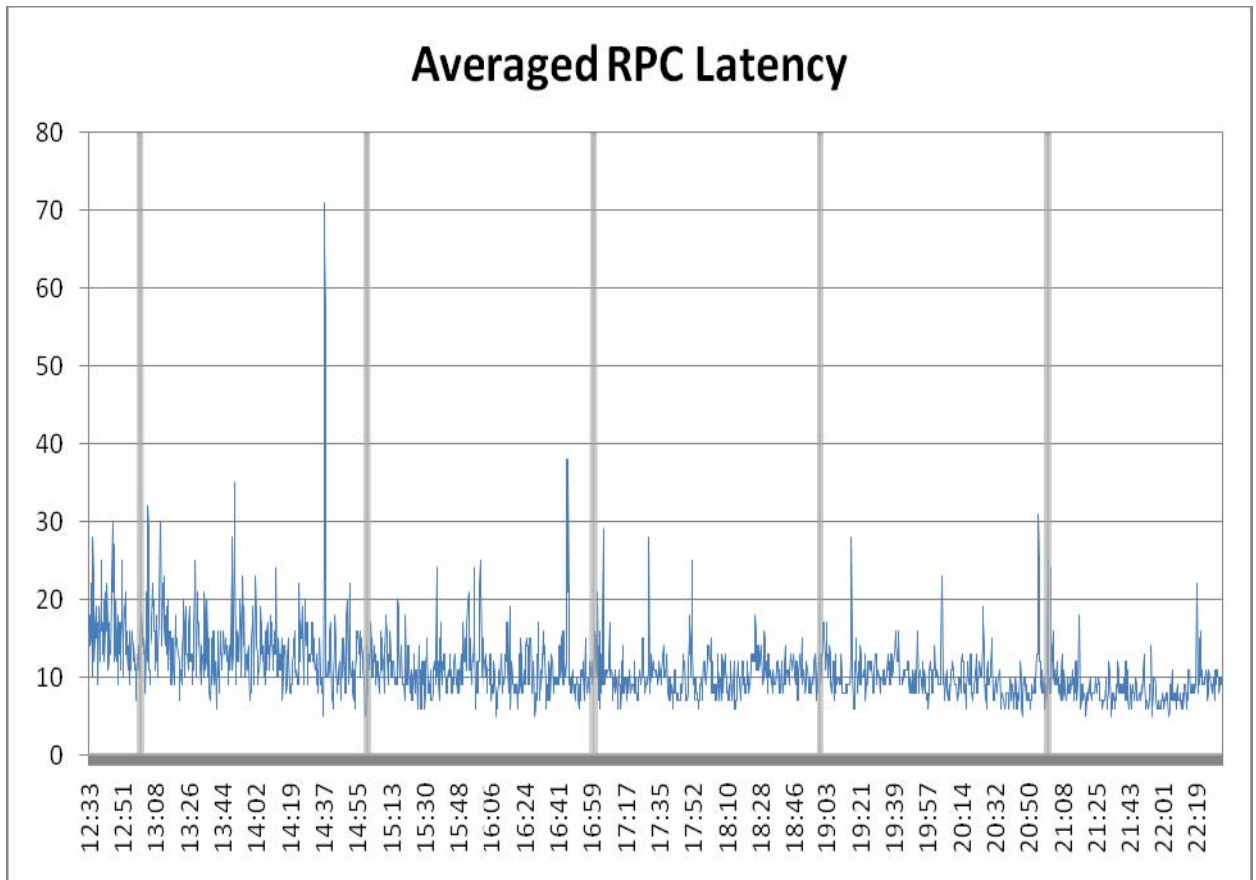


Figure 4: Averaged RPC latency with snapshots.

Database-read latency with snapshots

The following chart shows the database read latency with snapshots occurring every two hours. During the specified run period, database read latencies were eight milliseconds on average with a maximum



spike of up to 22 milliseconds. Note that during the duration of the snapshot backup operations, no sustained performance spikes occurred in database read latency.

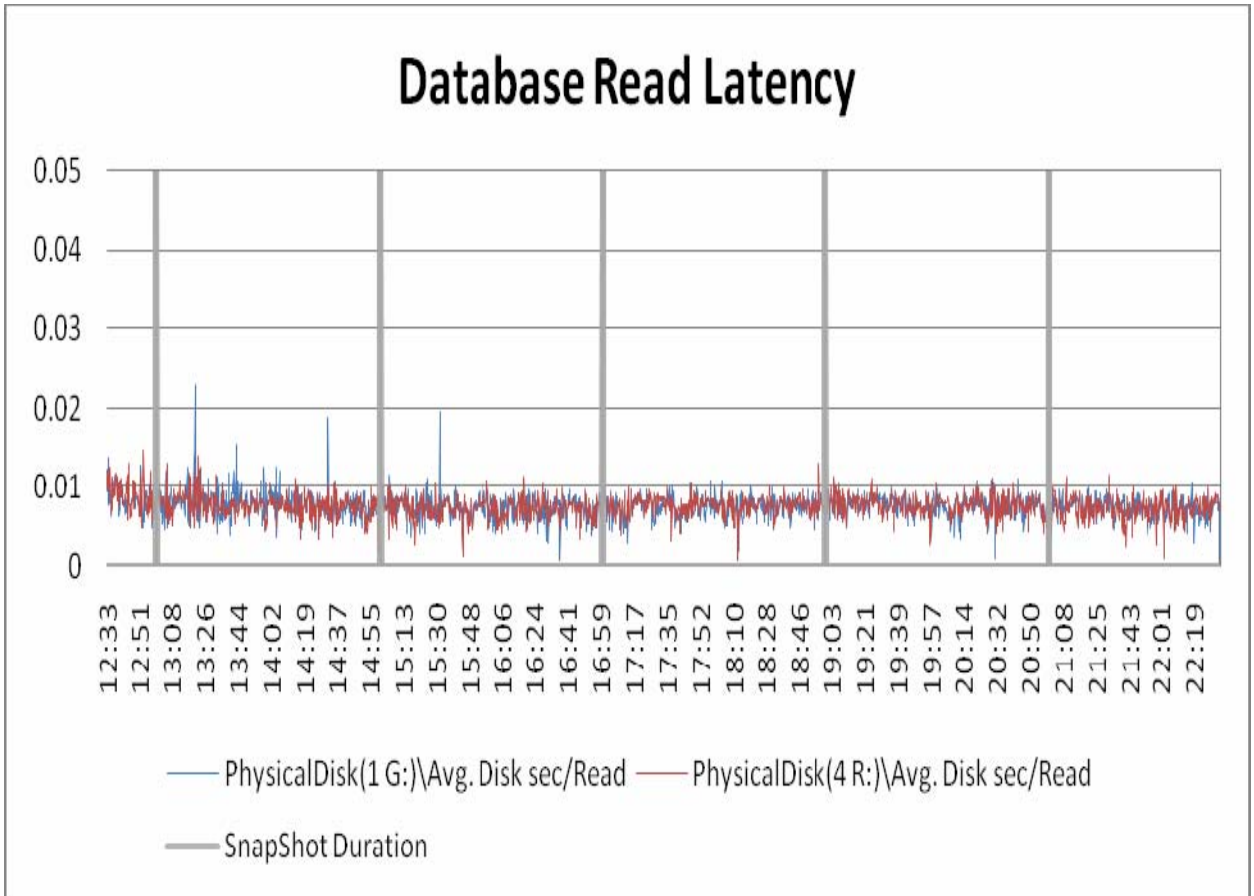


Figure 5: Database read latency with snapshots.



CPU utilization with snapshots

The following chart shows the CPU utilization during the testing period. Average CPU utilization during this testing period was 11% with a maximum spike of 51%.

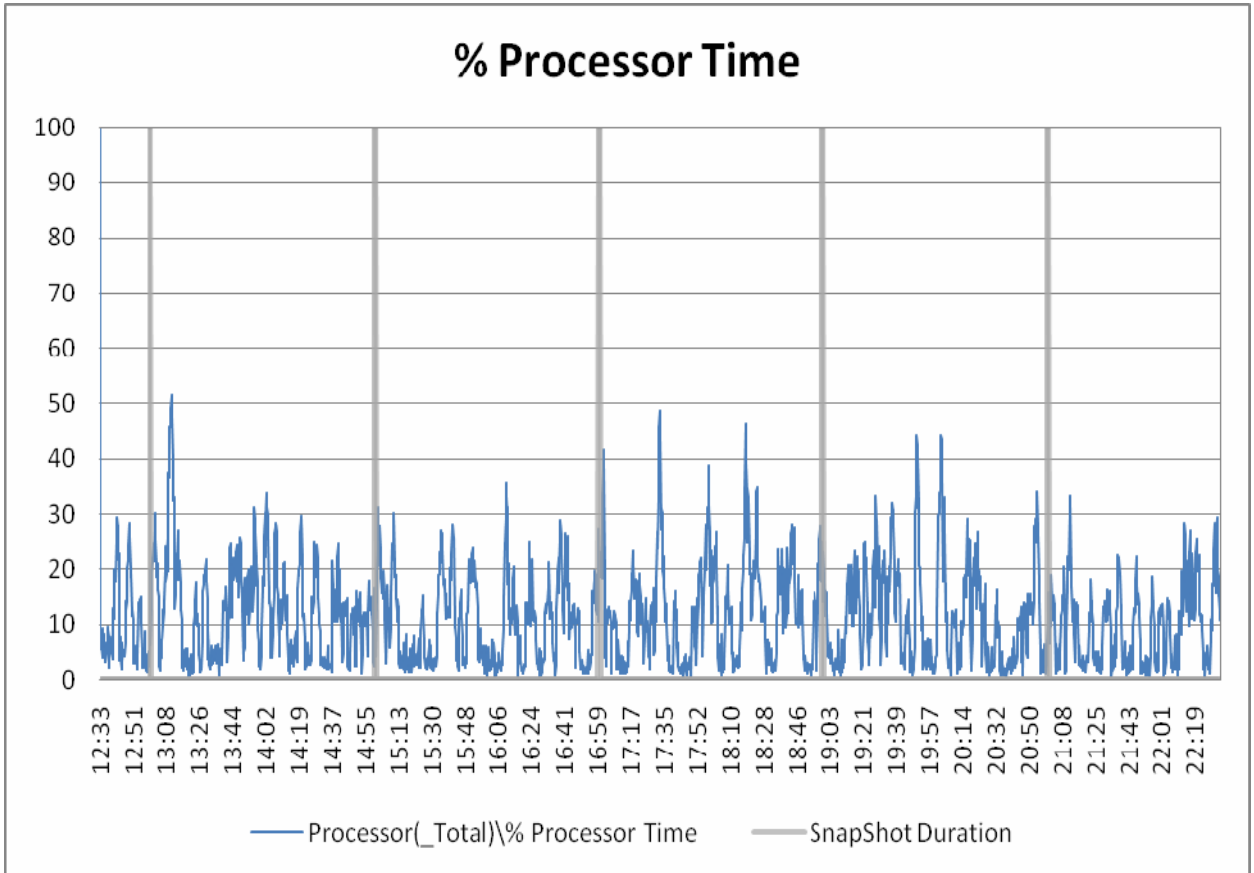


Figure 6: CPU utilization with snapshots.



Performance impact of deleting snapshots

Tests were configured to keep a total of 24 snapshots online. Once the number of online snapshots threshold is met, the oldest snapshot is deleted after successful completion of the next snapshot backup operation. This section examines Exchange Server performance characteristics when deleting a VSS snapshot.

A base of 24 snapshots were kept online to provide a base for this testing period. SME was configured to keep a total of 24 snapshot backups online.

As part of a scheduled database backup job, backup deletion criteria (based on snapshot copy age or quantity) can be specified. SnapManager automatically deletes the oldest backups for that storage group and backup management group after the backup is finished.

In order to determine any impact deleting snapshots may have on performance, including the impact that IBM N series with Snapshot technology has on performance, performance counters were examined for a period of 4 minutes before and 20 minutes after each snapshot operation. The following charts show the performance characteristics of keeping 24 snapshot backups online. All counters examined showed no dramatic increase while deleting the snapshot backups.

Averaged RPC latency with snapshot deletion

The chart below shows a representative sample of averaged RPC latency surrounding a snapshot backup. The graph below shows temporary increases in RPC latency after the backup was performed and deleting the snapshot after it was completed. It is important to note that, while a temporary increase is shown, the averaged RPC latency did not go beyond the 50-ms threshold.

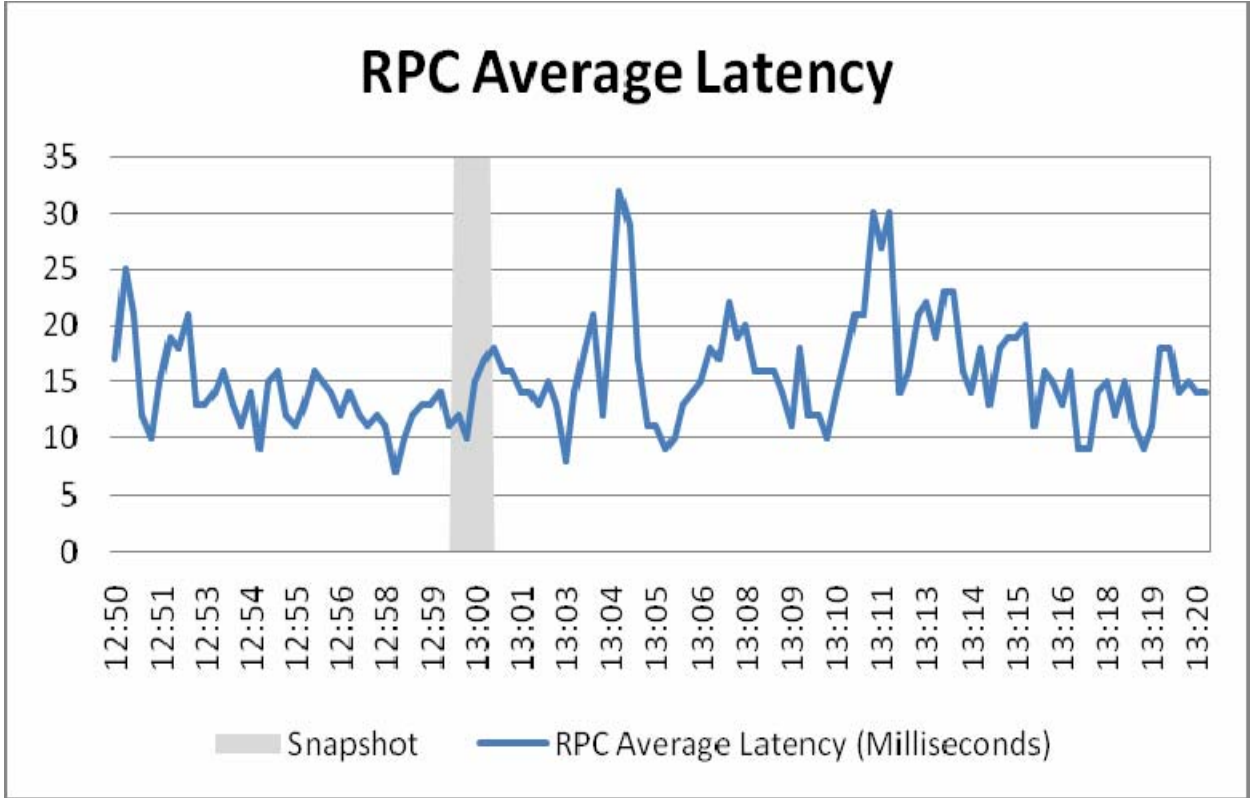


Figure 7: RPC latency with snapshot deletion.

Database-read latencies with snapshot deletion

The chart below shows a representative sampling of read latencies of the database LUNs. No sustained increase in database read latency is experienced as part of deleting the snapshot backup.

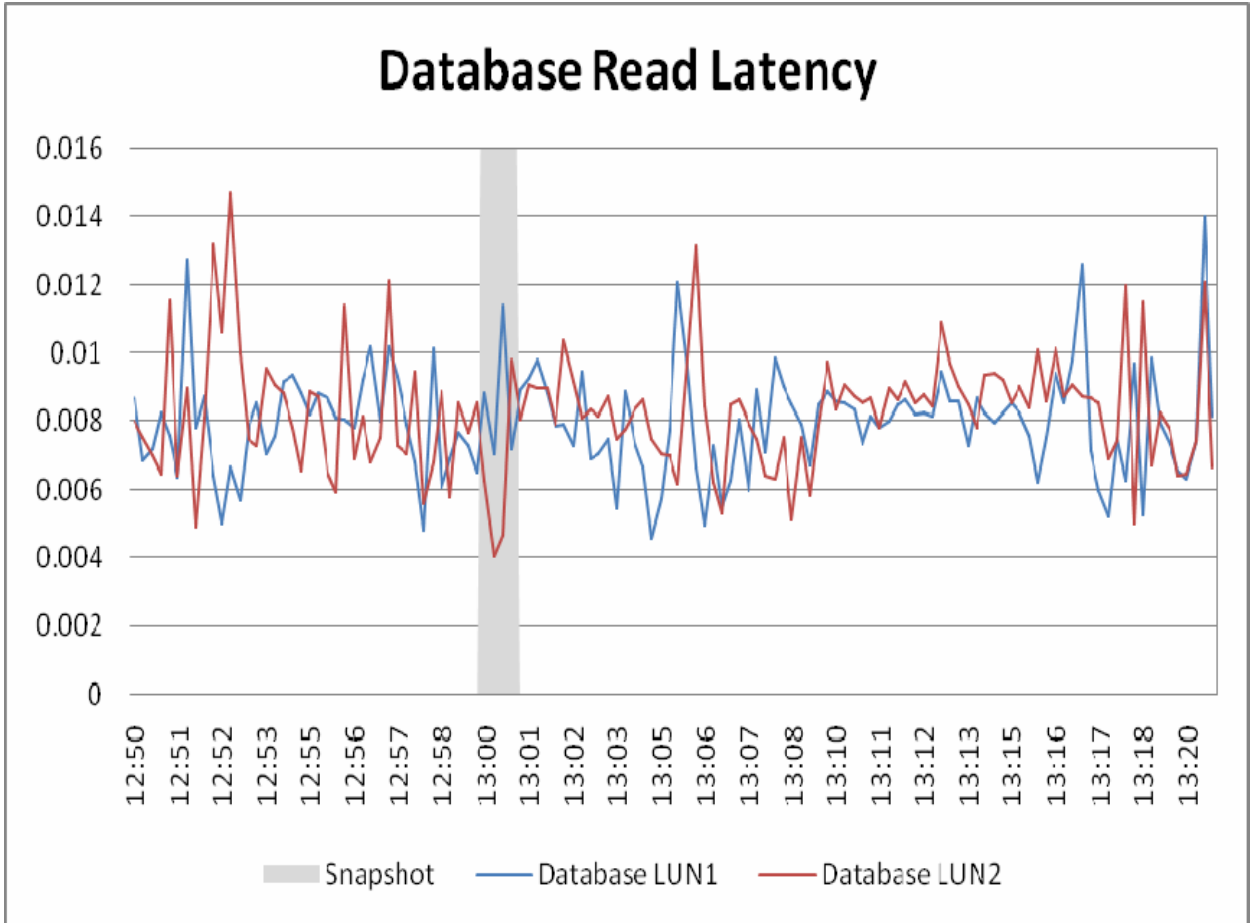


Figure 8: Database read latency with snapshot deletion.

CPU utilization with snapshot deletion

The chart below shows a temporary increase in CPU utilization in the minutes after the snapshot process. CPU utilization increased to a maximum of 50% then returned to the normal range.

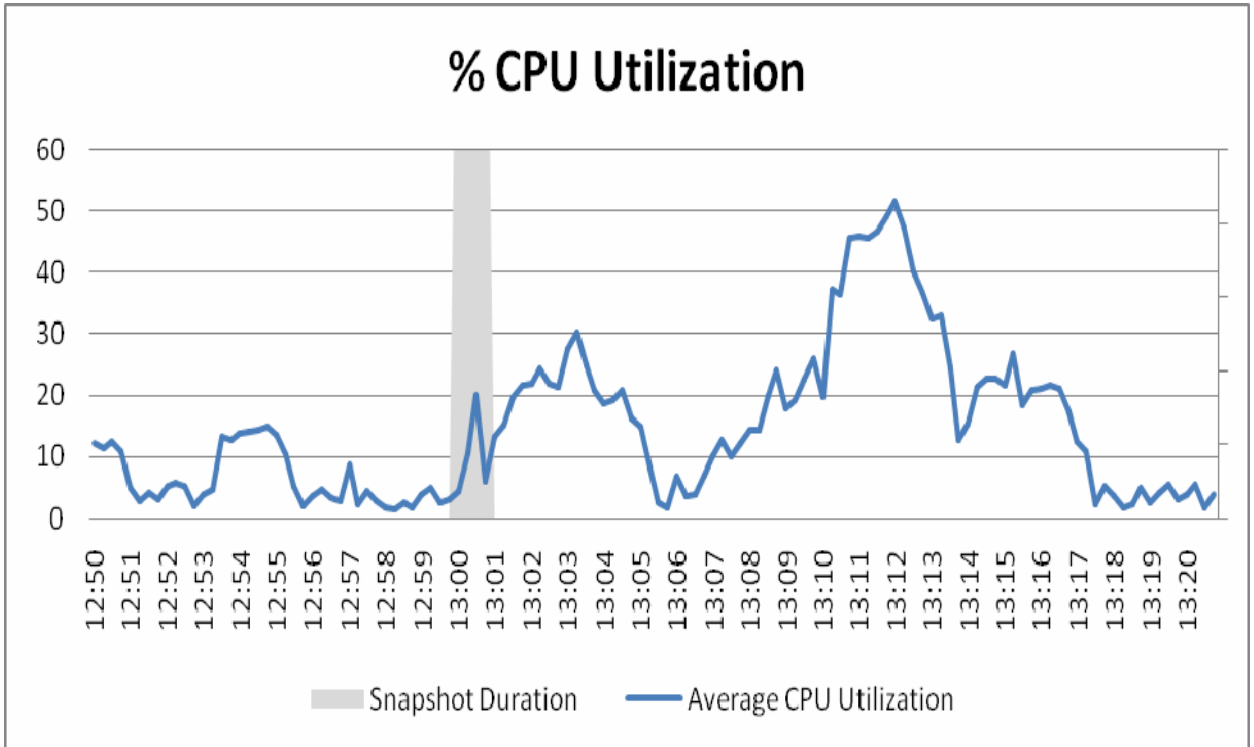


Figure 9: CPU utilization with snapshot deletion.



Summary

After examining Exchange Server 2007 performance characteristics during periods using IBM N series with Snapshot backups, it is clear that there is no performance impact on an Exchange Server. All performance counters, on average, remained virtually unchanged. It is also important to note that while running the test scenario with 24 snapshot backups kept online, performance not only remained consistent but also remained well within accepted performance thresholds and Microsoft best practices.

The combined features of Exchange Server 2007 and IBM N series with SME provide unparalleled ease of use and performance results. The solution also provides IT departments with a method of meeting the backup demands of ever-increasing Exchange mailbox databases.



Trademarks and special notices

© International Business Machines 1994-2008. IBM, the IBM logo, Lotus, Notes, System Storage, and other referenced IBM products and services are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. All rights reserved.

FlexVol, Data ONTAP, Network Appliance, the Network Appliance logo, SnapDrive, SnapManager, Snapshot and WAFL are trademarks or registered trademarks of Network Appliance, Inc., in the U.S. and other countries.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Other company, product, or service names may be trademarks or service marks of others.

References in this document to IBM products or services do not imply that IBM intends to make them available in every country.

Information is provided "AS IS" without warranty of any kind.

Information concerning non-IBM products was obtained from a supplier of these products, published announcement material, or other publicly available sources and does not constitute an endorsement of such products by IBM. Sources for non-IBM list prices and performance numbers are taken from publicly available information, including vendor announcements and vendor worldwide homepages. IBM has not tested these products and cannot confirm the accuracy of performance, capability, or any other claims related to non-IBM products. Questions on the capability of non-IBM products should be addressed to the supplier of those products.