

Microsoft® Exchange Server 2003, Microsoft Windows Storage Server 2003, and HP ProLiant Storage Server — white paper



About this document.....	2
Intended audience.....	2
Prerequisites	2
Introduction.....	2
Overview	3
Related documentation	4
Recommended design principles.....	5
Network design	5
Exchange storage design	6
Clustering.....	7
User profile definition and sizing	7
Tested configurations	8
Recommended low-capacity configuration.....	8
Recommended medium-capacity configuration	10
Recommended high-capacity configuration	13
Test results summary.....	17
Installing the Feature Pack.....	18
Creating shares	20
Relocating the Exchange files	21
Managing Exchange 2003 on NAS	22
Monitoring	23
Managing Exchange files and services	26
Antivirus	27
Full-text indexing	27
Disaster recovery	27
Backup and restore testing	27
Recovery storage groups	28
Windows Storage Server future	28
For more information.....	29
HP storage website	29
HP authorized reseller	29
HP technical support.....	29

About this document

This document provides validated configurations for Microsoft® Exchange Server 2003 on an HP ProLiant Storage Server running Microsoft Windows® Storage Server 2003 through an addition from Microsoft called the Feature Pack. The Feature Pack enables placement of the databases and logs for one or two Exchange servers (either Standard or Enterprise edition) and up to 1,500 mailboxes on a single Windows Storage Server device.

This paper provides solution details, performance testing results, and best-practices. Demonstration (for a limited time) and consultation for this solution are available by asking your HP sales channel to contact microsoft_storage_solutions@hp.com.

Intended audience

The primary audience for this document is:

- System architects, consultants, and network engineers who plan the components and the deployment
- Storage administrators who perform the design, installation, and management of network attached storage (NAS) using Windows Storage Server
- Exchange administrators who perform the design, installation, and management of Exchange servers and backup and restore of Exchange database files

Prerequisites

This document assumes that the reader:

- Is familiar with Exchange server administration concepts
- Has experience with Windows Storage Server 2003

Introduction

With the new Feature Pack, Windows Storage Server 2003 enables you to use NAS for applications, file serving, and now Exchange server. Microsoft announced support for Exchange Server 2003 databases and log files on Windows Storage Server 2003 Standard Edition in April, 2004. The Feature Pack for Windows Storage Server 2003 enables Exchange Server 2003 databases and transaction logs to be stored on an HP ProLiant Storage Server running Windows Storage Server 2003.

The Feature Pack is installed on both the Windows Storage Server computer (NAS device) and the Exchange Server 2003 computer to add new components and functionality to each. These components provide tools and services that enable Exchange databases and transaction logs to be moved to a Windows Storage Server computer and to configure Exchange Server 2003 access to the remotely stored files.

There is no additional charge for the Feature Pack. A CD ships with the HP ProLiant DL100 Storage Server or HP ProLiant DL380 G4 Storage Server, which includes Microsoft hotfixes, security updates, and the Feature Pack.

The *HP StorageWorks Service Release 1 Feature Pack Deployment Guide* (available from ftp://ftp.hp.com/pub/information_storage/software/nsas/nas/366581-001.pdf) defines the following scenarios:

- Low-capacity scenario (for up to 250 mailboxes)
- Medium-capacity scenario (for up to 750 mailboxes)

- High-capacity scenario (for up to 1,500 mailboxes)

The results of testing these scenarios are presented here, including definition of the bill of materials for each configuration. The actual configurations recommended in this paper may differ from the deployment guide (see Table 1 for a list of documentation and resources). The configurations in this paper are based on stress-tested configurations, as opposed to bare-minimum configurations, that should be fully deployable in the majority of production environments. This document provides best practices, for example, to size systems for peak loads and to accommodate future growth, thus avoiding the penalty of downtime when a system is undersized and must be upgraded in the immediate future.

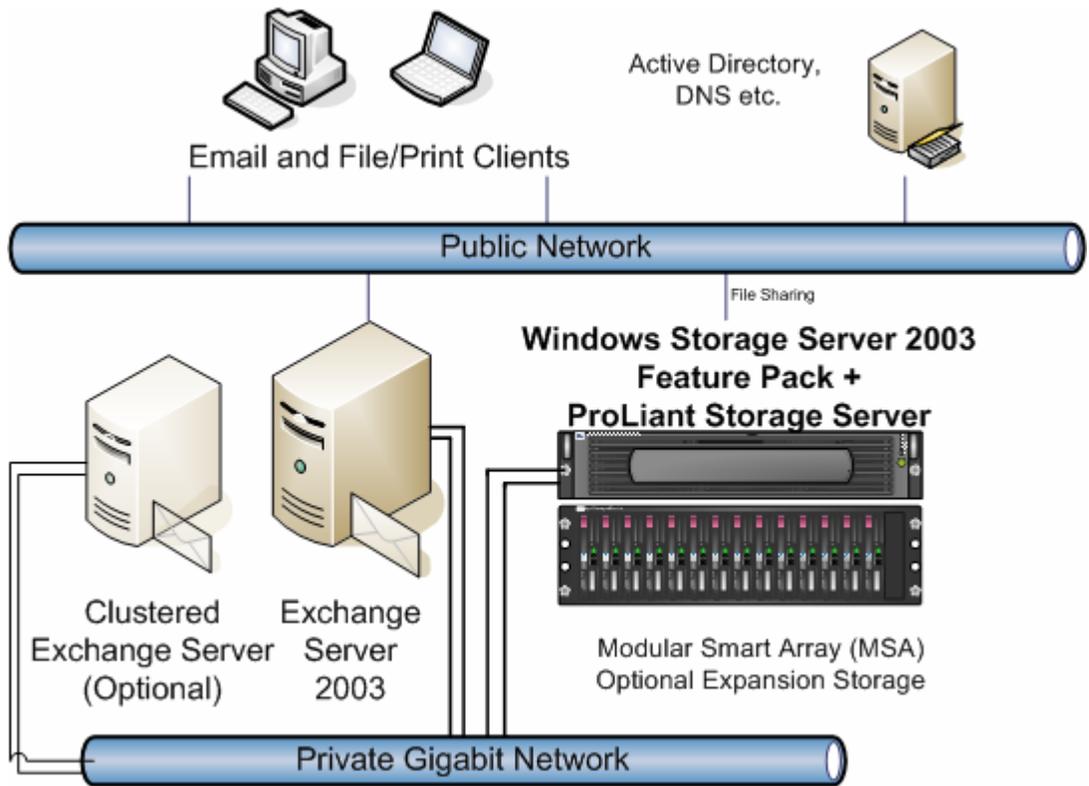
Overview

A single Windows Storage Server computer running the Feature Pack can host the databases and transaction logs of up to two Exchange servers and up to 1,500 Exchange mailboxes, according to the Microsoft design specifications. This paper details the testing and validation of several configurations of server and NAS hardware.

Exchange database tools and applications that are currently being used by an organization, for example, backup and anti-virus software, can still be used in conjunction with the Windows Storage Server 2003 NAS and the Feature Pack.

Figure 1 helps to illustrate the basic configuration. Each element is discussed in detail later.

Figure 1. Basic network configuration of Exchange Server 2003 and Windows Storage Server 2003



Related documentation

Many of the resources and documents available to assist you in your deployment are compiled in Table 1.

Table 1. Additional documentation and resources

Resource	Documentation	Link
Microsoft knowledge base article	#839687, "Microsoft support policy on the use of network-attached storage devices with Exchange Server 2003"	http://support.microsoft.com/default.aspx?kbid=839687
Microsoft knowledge base article	#328879, "Using Exchange Server with Storage Attached Network and network-attached storage devices"	http://support.microsoft.com/default.aspx?scid=kb;en-us;328879
Microsoft white paper	Managing Exchange Storage with the Windows Storage Server 2003 Feature Pack	http://www.microsoft.com/windowserversystem/wss/2003/techinfo/plandeploy/exchange_fp.mspx
HP ProLiant Storage Server	Information on ordering HP ProLiant Storage Server configurations	http://h18006.www1.hp.com/storage/networkattached.html
Help systems	Context-sensitive help for the Remote Storage wizard, remote storage tools for Exchange, and the Windows Storage Server WebUI	Included in default installation
Installation guide	Procedures for use during the installation of the software and hardware	Included on the CD that contains the Feature Pack. Installed on both the Windows Storage Server computer and the Exchange server during the installation process
Release notes	Information about late-breaking bugs or workarounds for known issues	Included on the CD that contains the installation files for the Feature Pack
Feature Pack deployment guide	Information to help you install the Windows Storage Server 2003 Feature Pack on HP ProLiant DL100 Storage Server or HP ProLiant DL380 G4 Storage Server	ftp://ftp.hp.com/pub/information_storage/software/nsas/nas/366581-001.pdf

Recommended design principles

This section outlines the HP recommended design principles for any size deployment. The next sections detail the sized and tested configurations.

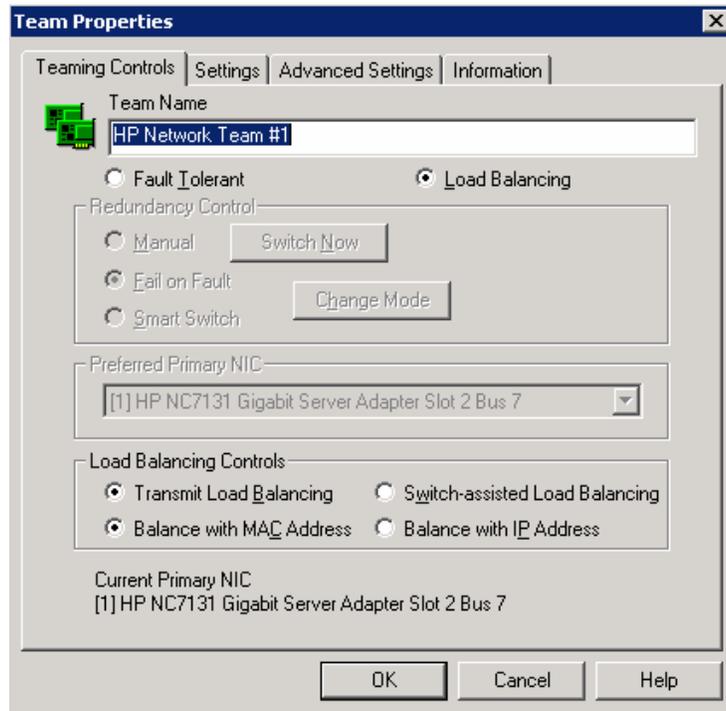
Network design

In any deployment, it is recommended that a dedicated Gigabit Ethernet network be used between the Exchange server and the Windows Storage Server NAS to ensure adequate performance as well as provide data security against network sniffing of Exchange data. An alternative would be to use IPSec to secure the connection if it is not possible to use a private, secured network—but there will be a performance overhead for using the IPSec.

In some scenarios, the Microsoft NAS Feature Pack deployment guide recommends a crossover network cable between the NAS and the Exchange server. Adding an additional Gigabit network card and multiple Gigabit switches provides fault tolerance. In the test environment, the HP network teaming feature (Figure 2) was enabled, providing a 2-Gbps link in addition to the benefit of fault tolerance in case of network interface, link, or switch failure. The Gigabit Ethernet cables were connected to a private HP ProCurve switch.

If you plan to locate your Exchange server any distance away from the Windows Storage Server NAS, check with your networking hardware vendors on the specifics regarding the maximum supported distance. The maximum distance will vary according to the cable type and specifications.

Figure 2. HP network teaming combines Gigabit adapters for fault tolerance and throughput



Exchange storage design

This section discusses the most important criteria for design and selection of Exchange storage:

- Isolation of Exchange transaction logs from databases
- Selection of best RAID protection for performance and fault tolerance
- Hardware RAID controller with sufficient write-back caching for performance

Exchange transaction logs and databases must be stored on separate disk volumes to provide both data protection and efficiency (separation of sequential writes and random read/write access, respectively). The Feature Pack and NAS can be configured to place the Exchange transaction logs on the Exchange server and the databases on the Windows Storage Server NAS. However, you must consider how to recover the transaction logs if the Exchange server fails and is replaced with new hardware; in that case, the log drives would need to be moved to the new server.

Transaction logs should be placed on a RAID 1 (mirror pair) array (volume), or for additional disk space, four or more disks in a RAID 1+0 (striped mirror) array. The number of spindles in the array determines storage space, but not performance. Performance on the transaction log volume is enhanced by decreasing the response time, which is accomplished by write-back caching. On HP Smart Array controllers with battery-backed write cache, such as the Smart Array 5i Plus and later, the write cache percentage should be set at 100%. (Dedicated read cache memory is built into the controller.) This setting will also benefit performance on the database arrays (volumes).

For database arrays (volumes), the choice of RAID protection on the disk arrays is often a trade-off between maximum storage and performance. While RAID 5 can provide data protection, it does so at the cost of performance. RAID 1+0 was shown in the testing to provide the absolute best performance given the same number of disk spindles. For example, the descriptions of the tests explain how six spindles in a RAID 5 array cannot provide sufficient I/O rate to support the same

number of Exchange mailboxes as the same number of disks in a RAID 1+0 array. RAID 1+0 is preferable for the database volume—the tests used 36-GB drives in an array, and the number of disks required to support the I/O also provided ample storage for 100-MB mailboxes for the users. Exercise caution when sizing if you are using the newest disk drives (for example, 146 GB or larger), as a few spindles can support the required database storage capacity but may not be able to support the required I/O performance.

As the testing results described later in this document indicate, it is most important to place the Exchange logs and database files on a RAID 1+0 array on a hardware RAID controller with sufficient write-back caching. Even if the logs need to be accessed over the gigabit network (on the NAS), there should be adequate network performance, and the RAID controller performance is paramount.

Clustering

Note

All references to clustering and the Feature Pack in this document refer to clustering of the Exchange servers, as opposed to clustering of the NAS servers (often referred to as NAS heads)—the Feature Pack does not support clustering the Windows Storage Server devices. The Feature Pack enables two-node Exchange server clusters with a Windows Storage Server 2003 device to provide shared storage for the cluster nodes.

For more information on clustering Exchange servers, see the HP white paper “HP Best Practices for Microsoft Exchange Server 2000 and 2003 Cluster Deployments” at <http://h71019.www7.hp.com/ActiveAnswers/Render/1,1027,6638-6-100-225-1,00.htm>.

In summary, the Feature Pack must be installed on each Exchange server in the cluster. When the Feature Pack is installed on an Exchange server that is clustered, the Distributed File System (DFS) root and mapped drive are created and the DFS information is stored in the cluster key of the registry. The DFS root and registry information is then replicated between cluster nodes so that both nodes use the same configuration. When the Feature Pack is installed on the second node of the Exchange cluster, the DFS root and registry information already exists and does not need to be recreated. When performing a full move of Exchange databases and/or logs, or just a configuration update for an Exchange cluster, all nodes in the cluster must be online, otherwise the move will not proceed.

The following installation dependencies should be noted: The cluster must be configured before Exchange is installed, including the Distributed Transaction Coordinator (DTC), which requires a physical disk resource. (For more information on placement of the DTC, see the HP white paper previously mentioned.) After Exchange is installed, the Feature Pack can be installed. The Exchange databases can be located on the Windows Storage Server NAS device. However, the physical disk resource required for the quorum and the DTC will need to be either a shared SCSI or Fibre Channel array such as the HP StorageWorks Modular Smart Array 500 (MSA500) or HP StorageWorks Modular Smart Array 1000 (MSA1000), respectively.

User profile definition and sizing

In any test environment measuring user load, the evaluation should closely match what will be deployed in production. For the tests in this paper, LoadSim 2003 was used to simulate the user load against the Exchange server. To most accurately represent an organization that relies on e-mail, the user profile selected within LoadSim was a mix of medium and heavy users, at a ratio of 2:1, respectively. These profiles represent quite active e-mail users and are used to simulate Enterprise environments. For the smaller organizations, where e-mail may not be so heavily used, these profiles may overstate the server impact, but they do provide assurance that the following configurations are properly sized to meet the most stressful peak-period demands.

In addition to the level of activity by the e-mail users, another sizing consideration is the size of the mailboxes. The medium profile creates on average a 60-MB mailbox, and the heavy profile creates on average a 100-MB mailbox. Larger mailboxes do affect not only your storage sizing but also performance criteria, as Exchange performance is affected by managing the larger mailboxes.

The final measure of LoadSim results is the 95th-percentile response time (a measure of client-side latency), which should be well below the desired goal of 1 second (1,000 milliseconds).

Tested configurations

The following sections define the hardware recommended for each scenario: low-capacity, medium-capacity, and high-capacity.

Recommended low-capacity configuration

The low-capacity scenario supports up to 250 mailboxes on one Exchange Server 2003 server. The deployment guide specifies the HP ProLiant DL100 Storage Server for the low-capacity configuration. Performance testing using LoadSim was completed with the HP ProLiant DL100 Storage Server, and it was found to support the low-capacity scenario supports up to 250 mailboxes on one Exchange Server 2003 server. However, the following discussion will focus on the HP ProLiant DL380 G4 Storage Server because of its availability, performance, and expansion capabilities.

The test results listed as follows show the HP StorageWorks NAS 2000S (now replaced by the HP ProLiant DL380 G4 Storage Server). The primary storage difference between the HP ProLiant Storage Server models is that the DL380 G4 Storage Server offers a Smart Array RAID controller with hot-swappable SCSI drives instead of ATA drives. For more information, see <http://h18006.www1.hp.com/storage/networkattached.html> or go to <http://www.hp.com/storage> and click network attached storage.

Figure 3 illustrates the sample configuration for a low-capacity scenario (up to 250 mailboxes). Note that for the Exchange server the HP ProLiant DL380 server is chosen because it can be configured with fault-tolerant options for redundant fans and power supplies. However, another HP ProLiant server could easily support the 250 Exchange users. The main sizing concerns are sufficient RAM (1 GB) and enough disks for the Exchange database volume—especially if using RAID 5 (or Advanced Data Guarding [ADG]). As discussed earlier, RAID 1+0 is preferable for the database volume, especially if you want to support up to the full 250 mailboxes. To make that point very clear: RAID 1+0 is critical for best performance of Exchange database volumes. Note also that the transaction logs are placed in the Exchange server, which takes advantage of the additional storage capacity of the ProLiant DL380 server.

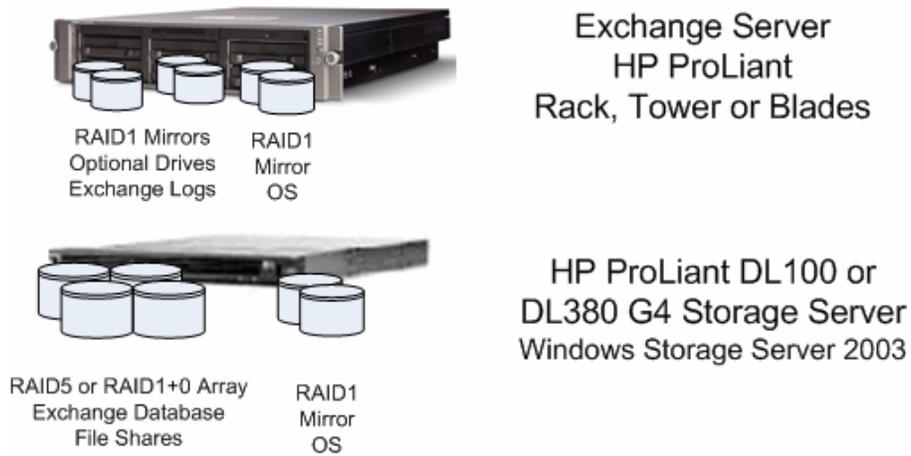
The following design was used in the validation testing: An HP ProLiant DL380 G2 server with dual Intel® Pentium® III 1.266-GHz processors was used for the testing—with the deliberate intention of not matching current generations of processors running around or above 3 GHz, as some customers may deploy existing HP ProLiant servers with the new DL380 G4 Storage Server. For example, the DL380 G4 Storage Server can be used to expand the current six-drive capacity of the HP ProLiant DL380 server. In addition, the HP ProLiant DL380 G2 server was tested without the option kit for the Smart Array 5i RAID controller—this allows a comparison with the Smart Array 5i Plus controller in the DL380 G4 Storage Server. The base Smart Array 5i RAID controller without the option kit does not allow the cache to be set at 100% write-back, which is important for performance, as discussed earlier in this document.

The HP ProLiant DL380 G2 server in combination with the NAS 2000s was found to provide a baseline 95th-percentile LoadSim response time of 160 milliseconds (well below the desired 1 second response rate) even with the addition of anti-virus scanning (Store VSAPI-based). The key to this response time was found to be using RAID 1+0 instead of RAID 5 for the database array on the Smart Array 5i Plus RAID controller. Only by adding a Smart Array 6402 controller (to the NAS 2000s)

was a four-disk RAID 5 array sufficient to deliver a 95th-percentile LoadSim response time of 280 milliseconds.

Results are summarized later in this document in Table 5.

Figure 3. Sample configuration for low-capacity scenario (up to 250 mailboxes)



A sample bill of materials for the low-capacity configuration shown in Figure 3 appears in Table 2. You can use it as a starting point for defining your own deployment.

Table 2. Configuration (bill of materials) for low-capacity scenario

Configuration component	Description	HP part number (NA) = North America
Exchange server	HP ProLiant DL380 G4 server	
1. RAID controller	1. Ultra3-based Smart Array 6i with optional 128-MB Battery Backed Write Cache (BBWC)	
2. Network adapter	2. Two-port NC7782 10/100/1000 embedded NIC	
3. Memory	3. 1-GB base memory (2 x 512 MB)	311144-xx1
4. Processor options	4. Intel® Xeon™ 3.6-1MB cache /800-MHz FSB Or other model such as Xeon 3.2 GHz 1-MB cache/533-MHz FSB, 1-GB RAM	370596-xx1
• Disks—See configuration for exact number and placement	Up to 14 Ultra320 hot-pluggable SCSI drives <ul style="list-style-type: none"> • 300 GB 10,000 rpm U320 Universal Hard Drive, 1 in. • 146.8 GB 10,000 rpm, U320 Universal Hard Drive, 1 in. • 72.8 GB 10,000 rpm, U320 Universal Hard Drive, 1 in. • 72.8 GB 15,000 rpm, U320 Universal Hard Drive, 1 in. • 36.4 GB 15,000 rpm, U320 Universal Hard Drive, 1 in. 	<ul style="list-style-type: none"> • 350964-B22 • 286716-B22 • 286714-B22 • 286778-B22 • 286776-B22
• Optional redundant power supply	Hot Plug AC Redundant Power Supply Module (NEMA and IEC cords) (NA and Japan) Hot Plug AC Redundant Power Supply Module (IEC cord) (WW)	355892-001 355892-B21
• Uninterruptible power systems	See http://h18002.www1.hp.com/products/quickspecs/11473_div/11473_div.HTML	

Configuration component	Description	HP part number (NA) = North America
Storage server	<p>HP ProLiant DL380 G4 Storage Server</p> <ul style="list-style-type: none"> 3.2GHz Intel® Xeon™ processor with EM64T, 800 MHz FSB Dual Port NC7782 PCI-X Gigabit NICs (embedded) 10/100/1000 Wake on LAN (WOL) Hot Plug Fans with optional full redundancy (Included in High performance models) Hot Plug Power Supply with optional redundancy (Included in High performance models) 	370596-xx1
Network switch	HP Gigabit Switch: ProCurve 2626	MFG #J4900A#ABA
Backup device	Options include <ul style="list-style-type: none"> Tape drives (S/DLT, LTO, and so on) Disk array (for example, on MSA) 	
Tested software:		
1. Operating system	1. Windows Server 2003 Standard Edition	See third-party vendor for ordering information
2. Messaging application	2. Exchange Server 2003 Standard Edition	
3. Management and monitoring	3. HP Systems Insight Manager	
	4. HP OpenView Storage Data Protector 5.1	
4. Backup	5. Sybari Antigen 7.5	
5. Antivirus	6. LoadSim 2003 (free download)	
6. Load testing		

Notes:

Part numbers were obtained from and can be updated at <http://h18002.www1.hp.com/products/quickspecs/productbulletin.html>.

Recommended medium-capacity configuration

The medium-capacity scenario supports up to 750 mailboxes on one Exchange Server 2003 server. Figure 4 illustrates the sample configuration. The HP ProLiant DL380 server is chosen because it can be configured with fault-tolerant options for redundant fans and power supplies. In addition, a second processor can be added. The main sizing concerns are sufficient RAM (1 GB) and enough disks for the Exchange database volume—especially if using RAID 5 (or ADG). As discussed earlier, RAID 1+0 is preferable for the database volume if you are using large disk drives and want to support up to the full 750 mailboxes. To make that point very clear: RAID 1+0 is critical for best performance of Exchange database volumes. To provide the additional disks, a Smart Array 6402 controller (in the DL380 G4 Storage Server¹) and a rack-mountable HP StorageWorks Modular Smart Array 30 (MSA30) disk enclosure (up to 14 drives with redundant power supplies) were added. Note also that the transaction logs are placed in the Exchange server as a RAID 1+0 array, which takes advantage of the additional storage capacity of the ProLiant DL380 server.

The design shown in the figure was used in validation testing with a single-processor ProLiant DL380 G2 server with 1 GB of RAM. Without the addition of anti-virus scanning (Store VSAPI-based), the tests delivered a 95th-percentile LoadSim response time of 438 milliseconds, with CPU utilization averaging 45%. The addition of anti-virus scanning required adding RAM to 2 GB, otherwise the tests failed, with scores exceeding 1 second and the Stores going offline. For more information, see the later sections on management and troubleshooting. With 2 GB of RAM and the addition of anti-virus

¹ In actual testing, the HP ProLiant DL380 G3 Storage Server was used, but test results with the new generation of Storage Server are assumed to offer improved performance.

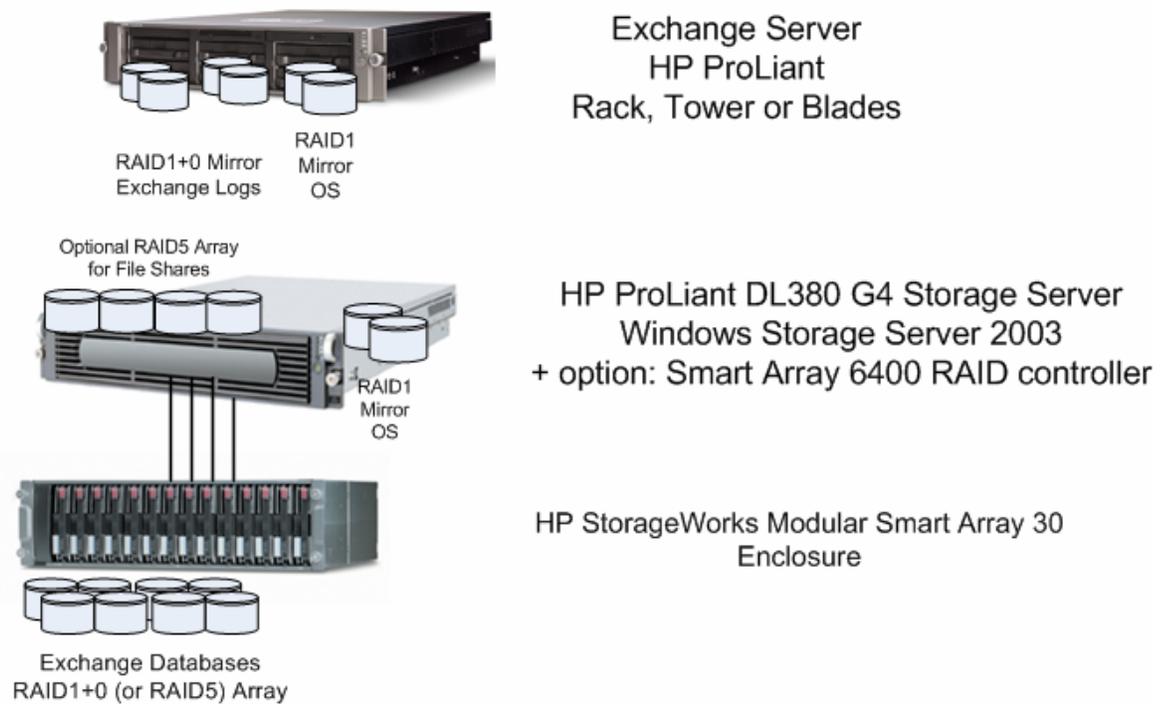
scanning, the system delivered a 95th-percentile LoadSim response time of 495 milliseconds, with CPU utilization averaging 50%.

When RAM was added to reach 3 GB, the gains were not that significant: With the addition of the anti-virus scanning, the system scored a 95th-percentile LoadSim response time of 454 milliseconds, and the CPU utilization was at 60%. Without the anti-virus scanning, the CPU utilization averaged 49%.

The key to these response times was found to be the use of RAID 1+0 instead of RAID 5 for the database array. The tests using RAID 5 with the same number of spindles (six) in the 2-GB configuration delivered a 95th-percentile LoadSim response time of 738 milliseconds, with CPU utilization averaging 54%. Adding anti-virus scanning to this configuration would require more memory and additional disks to maintain the same level of responsiveness.

Results are summarized later in this document in Table 5.

Figure 4. Sample configuration for medium-capacity scenario (up to 750 mailboxes)



A sample bill of materials for the medium-capacity configuration shown in Figure 4 appears in Table 3. You can use it as a starting point for defining your own deployment.

Table 3. Configuration (bill of materials) for medium-capacity scenario

Configuration component	Description	HP part number (NA) = North America
Exchange server	HP ProLiant DL380 G4 server	
1. RAID controller	1. Ultra3-based Smart Array 6i with optional 128-MB Battery Backed Write Cache (BBWC)	
2. Network adapter	2. Two-port NC7781 10/100/1000 embedded NIC	
3. Memory	3. 1-GB base memory (2 x 512 MB)	311144-xx1
4. Processor options	4. Intel® Xeon™ 3.6-1MB cache /800-MHz FSB Or other model such as Xeon 3.2 GHz 1-MB cache/533-MHz FSB, 1-GB RAM	370596-xx1
• Added RAM	HP ProLiant 2GB PC2-3200 2x1GB FIO BASE	343056-L21
• Disks—See configuration for exact number and placement	Ultra320 hot-pluggable SCSI drives <ul style="list-style-type: none"> • 300 GB 10,000 rpm U320 Universal Hard Drive, 1 in. • 146.8 GB 10,000 rpm, U320 Universal Hard Drive, 1 in. • 72.8 GB 10,000 rpm, U320 Universal Hard Drive, 1 in. • 72.8 GB 15,000 rpm, U320 Universal Hard Drive, 1 in. • 36.4 GB 15,000 rpm, U320 Universal Hard Drive, 1 in. 	<ul style="list-style-type: none"> • 350964-B22 • 286716-B22 • 286714-B22 • 286778-B22 • 286776-B22
• Second processor option kit	Intel X3.2GHz/800-1MB 380G4 FIO BASE Intel X3.4GHz/800-1MB 380G4 FIO BASE Intel X3.6GHz/800-1MB 380G4 FIO BASE	374492-L21 311583-L22 311584-L22
• Battery-backed write cache for RAID controller	Enabler option kit	346914-L21
• Optional redundant power supply	Hot Plug AC Redundant Power Supply Module (NEMA and IEC cords) (NA and Japan) Hot Plug AC Redundant Power Supply Module (IEC cord) (WW)	355892-001 355892-B21
• Uninterruptible power systems	See http://h18002.www1.hp.com/products/quickspecs/11473_div/11473_div.HTML	
Storage server	HP ProLiant DL380 G4 Storage Server	311144-xx1
	<ul style="list-style-type: none"> • 3.6-GHz Intel® Xeon™ processor with EM64T, 800 MHz FSB • Dual Port NC7782 PCI-X Gigabit NICs (embedded) 10/100/1000 Wake on LAN (WOL) • Hot Plug Fans with optional full redundancy (Included in High performance models) • Hot Plug Power Supply with optional redundancy (Included in High performance models) 	
• Smart Array RAID controller	<ul style="list-style-type: none"> • Smart Array 6402/128 controller • Or Smart Array 6404/256 controller 	273915-B21 (128 MB) 273914-B21 (256 MB)
• Drives for optional RAID 5 array	• Add Ultra320 15K SCSI drives	(See below)

Configuration component	Description	HP part number (NA) = North America
External storage	MSA30—Rack-mountable 14-drive enclosure with redundant power supplies Attached to the Smart Array 6402 controller <ul style="list-style-type: none"> Single bus Dual bus 	<ul style="list-style-type: none"> 302969-B21 302970-B21
<ul style="list-style-type: none"> Added disks 	Up to 14 Ultra320 hot-pluggable SCSI drives <ul style="list-style-type: none"> 300 GB 10,000 rpm U320 Universal Hard Drive, 1 in. 146.8 GB 10,000 rpm, U320 Universal Hard Drive, 1 in. 72.8 GB 10,000 rpm, U320 Universal Hard Drive, 1 in. 72.8 GB 15,000 rpm, U320 Universal Hard Drive, 1 in. 36.4 GB 15,000 rpm, U320 Universal Hard Drive, 1 in. 	<ul style="list-style-type: none"> 350964-B22 286716-B22 286714-B22 286778-B22 286776-B22
Network switch	HP Gigabit Switch: ProCurve 2626	MFG #J4900A#ABA
Backup device	Options include <ul style="list-style-type: none"> Tape drives (S/DLT, LTO, and so on) Disk array (for example, on MSA) 	
Tested software		
1. Operating system	1. Windows Server 2003 Enterprise Edition (to support clustering)	See third-party vendor for ordering information
2. Messaging application	2. Exchange Server 2003 Enterprise Edition (to support clustering and multiple storage groups)	
3. Management and monitoring	3. HP Systems Insight Manager	
	4. HP OpenView Storage Data Protector 5.1	
4. Backup	5. Sybari Antigen 7.5	
5. Antivirus	6. LoadSim 2003 (free download)	
6. Load testing		

Notes:

Part numbers were obtained from and can be updated at <http://h18002.www1.hp.com/products/quickspecs/productbulletin.html>.

Recommended high-capacity configuration

The high-capacity scenario supports up to 1,500 mailboxes on one or two Exchange Server 2003 servers. Figure 5 illustrates the sample configuration. The ProLiant DL380 server is chosen because it can be configured with fault-tolerant options for redundant fans and power supplies. In addition, a second processor was added. The main sizing concerns are sufficient RAM (2 GB or more) and enough disks for the Exchange database volume—especially if using RAID 5 (or ADG). As discussed earlier, RAID 1+0 is preferable for the database volume if you are using large disk drives and want to support up to the full 1,500 mailboxes. To make that point very clear: RAID 1+0 is critical for best performance of Exchange database volumes. A rack-mountable MSA30 disk enclosure (14 drives with redundant power supplies) was configured to provide the additional disks.

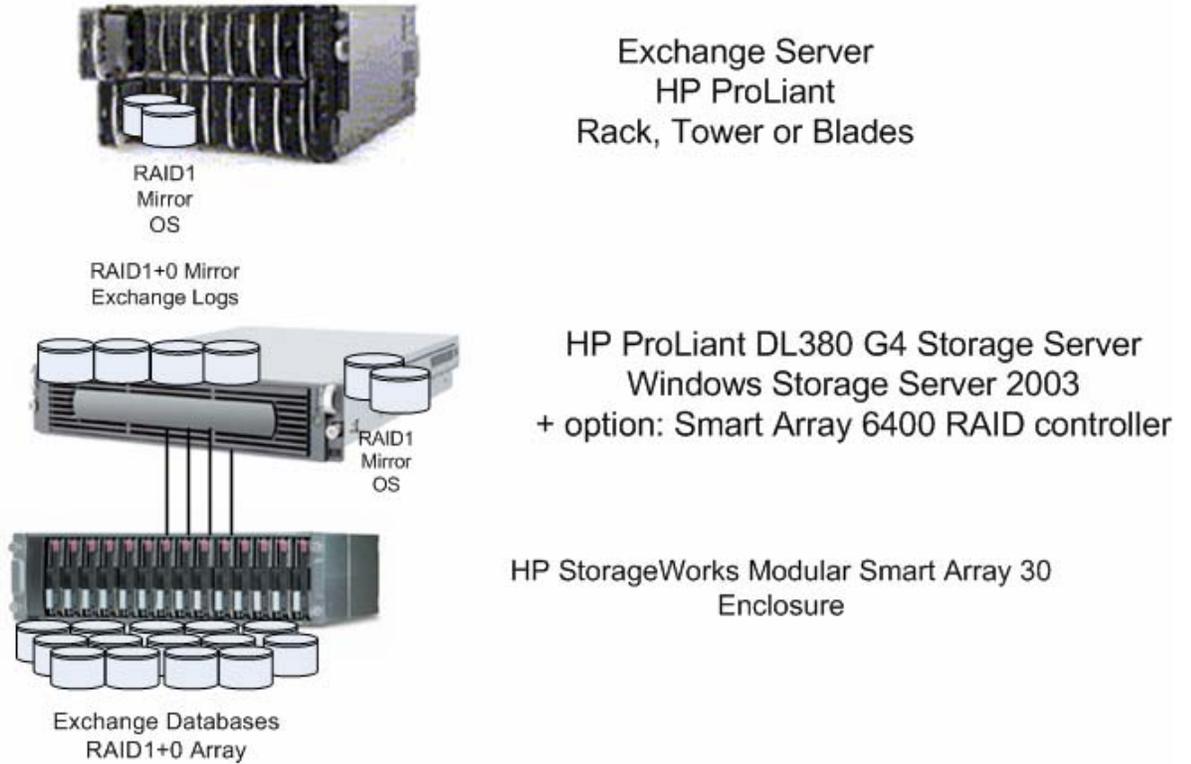
Note also that the transaction logs are placed in the Exchange server as a RAID 1+0 array, which takes advantage of the additional storage capacity of the HP ProLiant DL380 server.

The following design was used in validation testing: A dual processor ProLiant DL380 G2 server with 3 GB of RAM was found to provide a baseline 95th-percentile LoadSim response time of 485 milliseconds (well below the desired 1-second response rate) at 59% CPU utilization. The addition of anti-virus scanning (Store VSAPI-based) increased that to 536 milliseconds at 62% CPU utilization.

The key to this response time was found to be using RAID 1+0 instead of RAID 5 for the database array.

Results are summarized later in this document in Table 5.

Figure 5. Sample configuration for high-capacity scenario (up to 1,500 mailboxes)



A sample bill of materials for the high-capacity configuration shown in Figure 5 appears in Table 4. You can use it as a starting point for defining your own deployment.

Table 4. Configuration (bill of materials) for high-capacity scenario

Configuration component	Description	HP part number (NA) = North America
Exchange servers: Blade Systems	HP ProLiant BL40p P3.0-4MB, 1 GB (2P)	
1. Processor		344286-B21
2. Memory	<ul style="list-style-type: none"> (2) Xeon MP 3.0 GHz standard (up to 4 supported) 	
3. Network adapter	<ul style="list-style-type: none"> Cache Memory 4-MB Level 3 Cache (per processor) 	
4. RAID controller		
5. Options	<ul style="list-style-type: none"> Memory 1024 MB (Standard) of 2-way interleaved PC2100 ECC DDR (2 X 512 MB) Network Controller Five NC7781 PCI-X Gigabit NICs (embedded) with WOL and PXE One additional 10/100T NIC dedicated to iLO Management Storage Controller Smart Array 5i Plus Controller (integrated on system board) with BBWC 	
• Memory (RAM)	4096MB of PC2100 ECC DDR DIMM Memory Kit (2x2048MB)	300682-B21
• Disks—See configuration for exact number and placement	Up to 14 Ultra320 hot-pluggable SCSI drives	<ul style="list-style-type: none"> 350964-B22 286716-B22 286714-B22 286778-B22 286776-B22
	<ul style="list-style-type: none"> 300 GB 10,000 rpm U320 Universal Hard Drive, 1 in. 146.8GB 10,000 rpm, U320 Universal Hard Drive, 1 in. 72.8GB 10,000 rpm, U320 Universal Hard Drive, 1 in. 72.8GB 15,000 rpm, U320 Universal Hard Drive, 1 in. 36.4GB 15,000 rpm, U320 Universal Hard Drive, 1 in. 	
• Processor option kit	Xeon MP 3.0GHz 4MB Processor Option Kit	344287-B21
• Battery-backed write cache for RAID controller	Enabler option kit	255514-B21
• Optional redundant power supply		
• Uninterruptible power systems	See http://h18002.www1.hp.com/products/quickspecs/11473_div/11473_div.HTML	
Storage server	HP ProLiant DL380 G4 Storage Server	311144-xx1
	<ul style="list-style-type: none"> 3.6GHz Intel® Xeon™ processor with EM64T, 800 MHz FSB Dual Port NC7782 PCI-X Gigabit NICs (embedded) 10/100/1000 Wake on LAN (WOL) Hot Plug Fans with optional full redundancy (Included in High performance models) Hot Plug Power Supply with optional redundancy (Included in High performance models) 	

Configuration component	Description	HP part number (NA) = North America
<ul style="list-style-type: none"> RAID controller 	<ul style="list-style-type: none"> Smart Array 6402/128 controller Or Smart Array 6404/256 controller 	273915-B21 (128 MB) 273914-B21 (256 MB)
<ul style="list-style-type: none"> Drives for optional RAID 5 array 	<ul style="list-style-type: none"> Add Ultra320 15K SCSI drives 	(See below)
External storage	MSA30—Rack-mountable 14-drive enclosure with redundant power supplies. Attached to the Smart Array 6402 controller	
	<ul style="list-style-type: none"> Single bus Dual bus 	302969-B21 302970-B21
<ul style="list-style-type: none"> Added disks 	Up to 14 Ultra320 hot-pluggable SCSI drives <ul style="list-style-type: none"> 300 GB 10,000 rpm U320 Universal Hard Drive, 1 in. 146.8GB 10,000 rpm, U320 Universal Hard Drive, 1 in. 72.8GB 10,000 rpm, U320 Universal Hard Drive, 1 in. 72.8GB 15,000 rpm, U320 Universal Hard Drive, 1 in. 36.4GB 15,000 rpm, U320 Universal Hard Drive, 1 in. 	<ul style="list-style-type: none"> 350964-B22 286716-B22 286714-B22 286778-B22 286776-B22
Network switch	HP Gigabit Switch: ProCurve 2626	MFG #J4900A#ABA
Backup device	Options include <ul style="list-style-type: none"> Tape drives (S/DLT, LTO, and so on) Disk array (for example, on MSA) 	
Tested software: <ol style="list-style-type: none"> Operating system Messaging application Management and monitoring Backup Antivirus Load testing 	<ol style="list-style-type: none"> Windows Server 2003 Enterprise Edition (to support clustering) Exchange Server 2003 Enterprise Edition (to support clustering and multiple storage groups) HP Systems Insight Manager HP OpenView Storage Data Protector 5.1 Sybari Antigen 7.5 LoadSim 2003 (free download) 	See third-party vendor for ordering information

Notes:

Part numbers were obtained from and can be updated at <http://h18002.www1.hp.com/products/quickspecs/productbulletin.html>.

Test results summary

The results listed in the previous sections are summarized in Table 5. Note that the tests were performed using an HP StorageWorks NAS 2000s, which has now been replaced by a newer model, the HP ProLiant DL380 G4 Storage Server.

Table 5. Summary of tested configurations and progression of changes

Tested number of users ¹	Configuration change	Key performance results	
	Exchange server: HP ProLiant DL380 G2 server 1.266 GHz ² Baseline = 1 CPU 1-GB RAM unless noted Smart Array 5i RAID controller (tested without option kit) NAS 2000S with Smart Array 5i+ RAID controller	95th percentile LoadSim latency (milliseconds)	Average Exchange server CPU ²
250	Baseline: 4 disks RAID 1+0 on Smart Array 5i+	160	17%
250	4 disk RAID 5 array on Smart Array 6402 controller (in the NAS 2000S)	280	11%
750	MSA30 Storage Expansion attached to Smart Array 6402 controller in NAS 2000S with 6 disks RAID 5 array	738	54%
750	6 disks RAID 1+0 on Smart Array 6402 controller (in the NAS 2000S)	438	45%
750	Addition of antivirus scanning (Store VSAPI-based), requiring adding RAM to 2 GB	495	50%
750	Added RAM to 3 GB	454	60%
1,500	14 disks RAID 1+0 on Smart Array 6402 controller (in the NAS 2000S). 2 CPU with 3-GB RAM	485	59%
1,500	Addition of anti-virus scanning (Store VSAPI-based)	536	62%

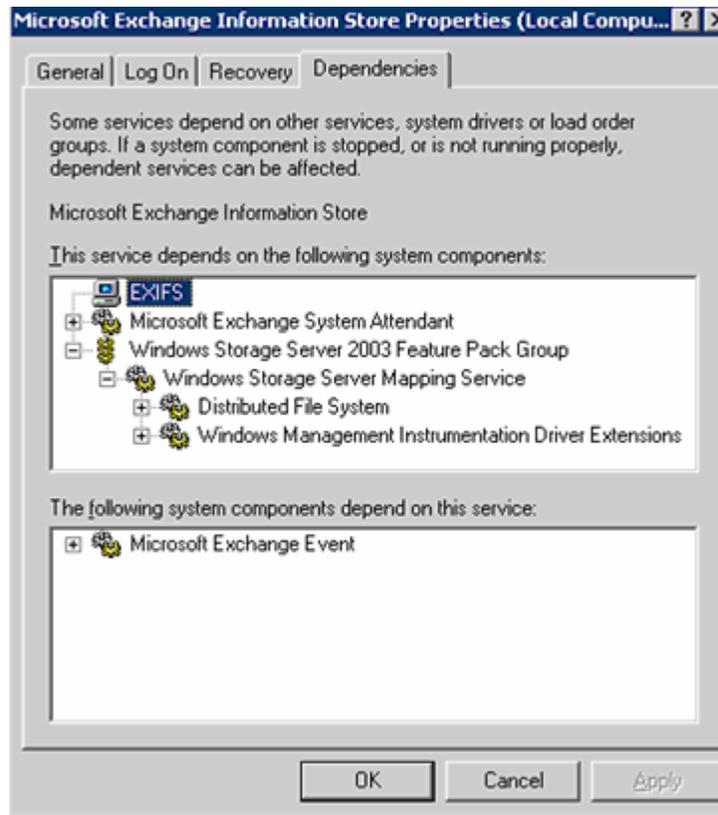
Notes:

1. See the section on User profile definition and sizing. Testing used a LoadSim 2003 user profile mix of medium and heavy users, at a ratio of 2:1, respectively.
2. An HP ProLiant DL380 G2 server with dual Intel Pentium III 1.266-GHz processors was used for the testing—with the deliberate intention of not matching the latest generations of processors running around or above 3 GHz, as some customers might deploy existing HP ProLiant servers with the new HP ProLiant DL380 G4 Storage Server.

Installing the Feature Pack

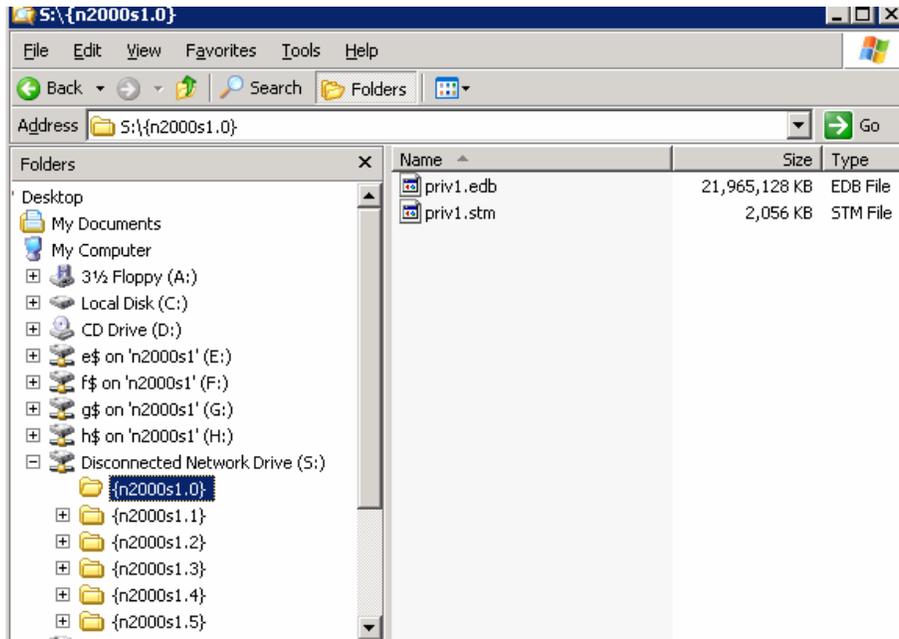
Installing the Feature Pack on the Exchange server is critical to enable Exchange Server 2003 to access the database and log files on the NAS. Note that at this point, other Exchange files, including the SMTP mailroot files and the transaction log checkpoint file, cannot be relocated to the NAS. As part of the Exchange server configuration, the Microsoft Exchange Information Store service dependencies are modified so that the configuration depends on the Windows Storage Server 2003 Feature Pack Group, as shown in Figure 6. This ensures that the NAS storage is online and accessed by way of the DFS, as shown in the later figures.

Figure 6. Information Store dependencies changed by Feature Pack installation



When the Exchange files are relocated to the NAS, DFS pointers are created, as illustrated in the view of Windows Explorer shown in Figure 7. Note that the S drive is used by default, and if it is not available on that system, the next available drive letter is attempted. If Exchange utilities such as ESEUTIL need to be run, it is more efficient to run them directly on the Windows Storage Server, rather than accessing them over the network drives.

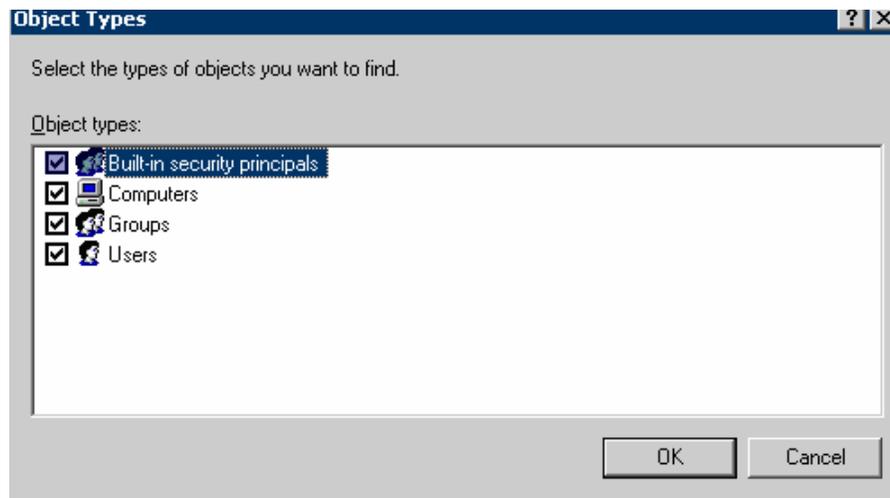
Figure 7. Sample DFS locations of Exchange database files



Creating shares

See the `Installation_Guide.doc` in the Feature Pack Exchange share (located in, for example, `C:\Program Files\Windows Storage Server\Exchange` or `\\nas\exchange`) for configuration of the shares on the Windows Storage Server. In summary, create the folder and share it, giving the Exchange administrator account, the backup operators, and the Exchange server computer account (for example, `servername$`) full control for the share and the underlying folder. Note that if you browse or even type in the `server$` account, you must select the computers Object type, as shown in Figure 8, or you will not be able to add computer accounts to the permissions. By default, the Everyone group will have Read permission, which can be removed.

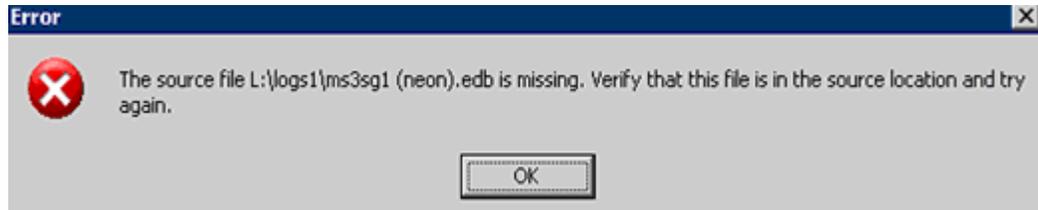
Figure 8. Select the computers Object type to add computer accounts to the permissions



Relocating the Exchange files

When the Feature Pack is installed on both the Exchange server and NAS head and the shares created, existing Exchange files can be relocated using the Remote Storage wizard. To access the wizard, right-click a storage group, and select All Tasks and Remote Storage wizard. Note that if you create a new Store, the Store must first be mounted (to create the database files) before Remote Storage wizard can relocate it. If not, you will see an error message similar to the one shown in Figure 9.

Figure 9. Error message when attempting to relocate Exchange files for new, unmounted Store

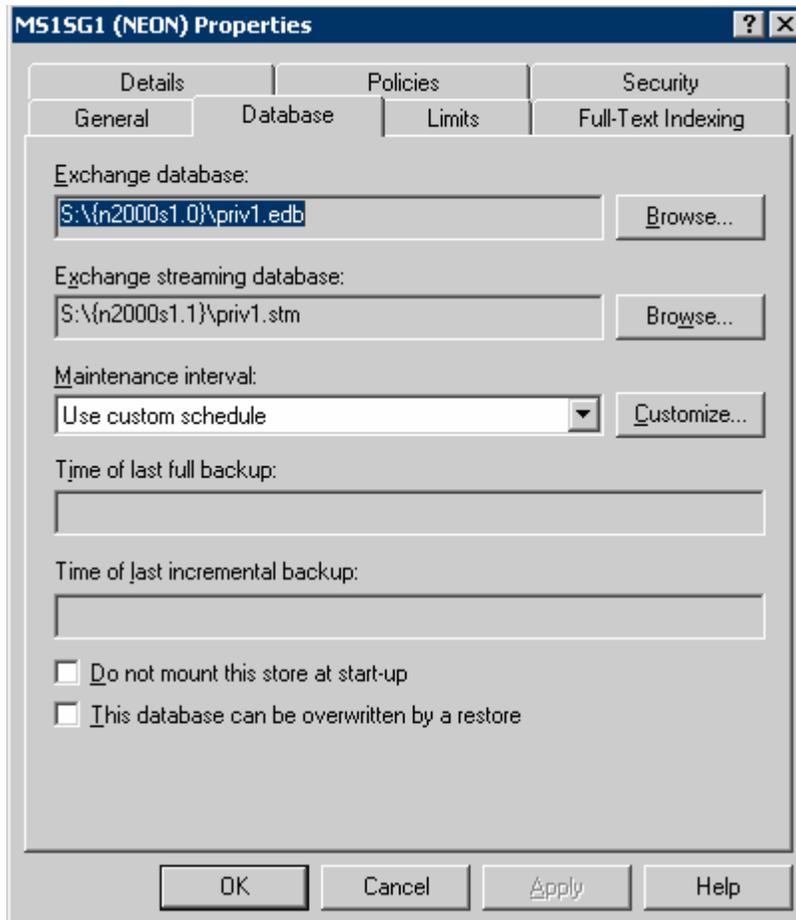


Instead of using the Remote Storage wizard you may want to use the command-line tool, WSSEchMove.exe. To use this tool, see the help provided for remote storage tools.

Note that when relocating transaction logs, the System path location (containing the checkpoint file) is not relocated. It is not possible to edit these properties directly, except in the circumstance that you are moving the transaction logs from the NAS back to local (for example, direct attached) storage.

Figure 10 shows the Exchange System Manager properties for a database, illustrating the DFS location.

Figure 10. Database DFS location as shown in Exchange System Manager



Managing Exchange 2003 on NAS

This section highlights some of the differences in managing Exchange on NAS. The Windows Storage Server devices from HP support HP Systems Insight Manager and HP OpenView for management and troubleshooting, integrating well into your existing HP ProLiant server management infrastructure.

Monitoring

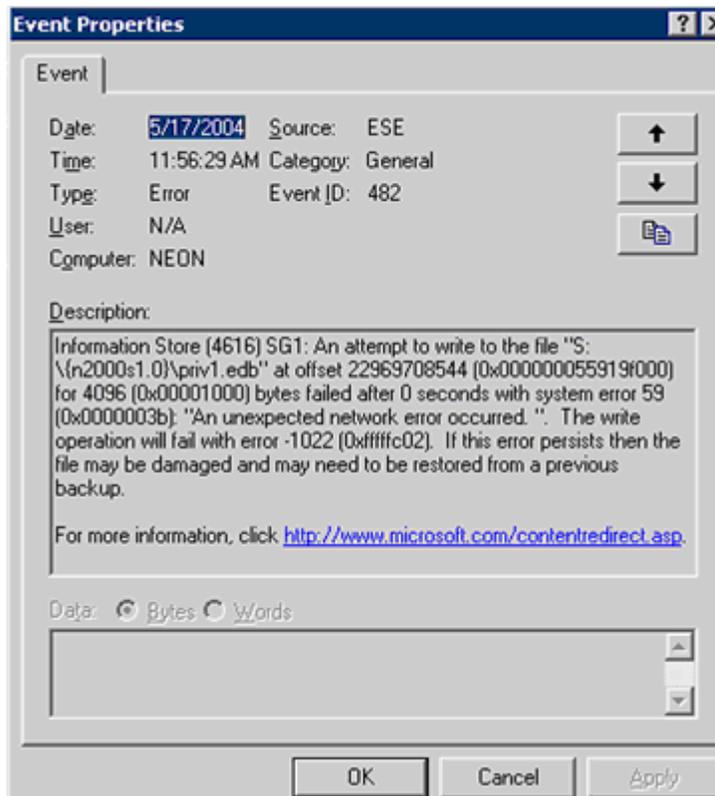
During stress testing it was found that if the recommended guidelines on sizing were exceeded (see earlier sections), there was a risk that the Exchange server would disconnect the Information Stores from the NAS. Figure 11 shows the resulting event 482 logging the Store shutdown. The events preceding this failure are listed below, so that you may monitor them (for example, in HP OpenView) and react quickly to this outage, bringing the Stores back online. If this occurs repeatedly, it might be necessary to move mailboxes to another server or increase the spindle count of the database array if it does not have enough spindles to support the user load.

481—Information Store (4616) SG1: An attempt to read from the file "S:\{n2000s1.0}\priv1.edb" at offset 17538949120 (0x0000000415672000) for 4096 (0x00001000) bytes failed after 0 seconds with system error 1398 (0x00000576): "There is a time and/or date difference between the client and server." The read operation will fail with error -1022 (0xffffc02). If this error persists then the file may be damaged and may need to be restored from a previous backup.

1159—Database error Disk IO error occurred in function JTAB_BASE::EcRetrieveColumns while accessing the database "SG1\MS1SG1"

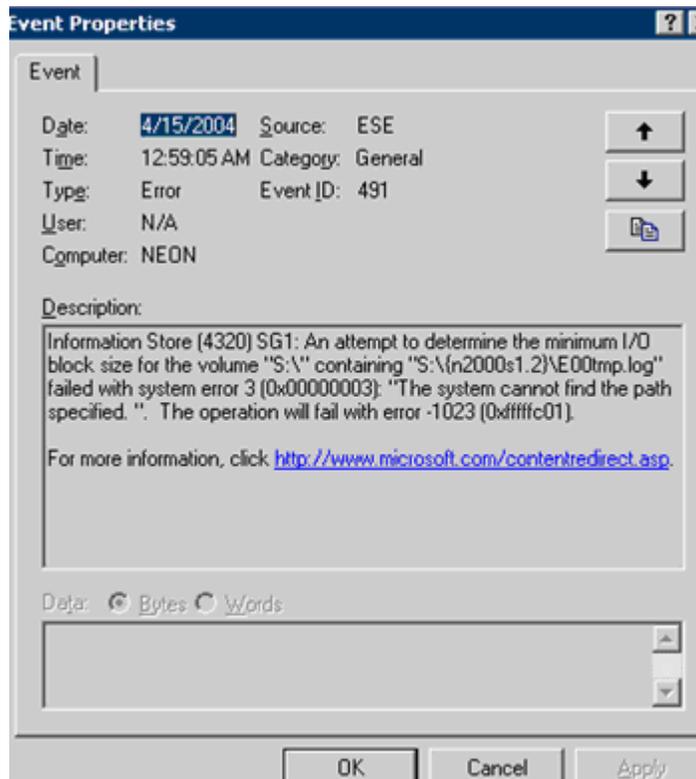
439—Information Store (4616) SG1: Unable to write a shadowed header for file S:\{n2000s1.0}\priv1.edb. Error -1022.

Figure 11. Event ID 482 resulting from Store shutdown (disconnect from NAS)



Similarly, if the log volume was placed on the NAS device (instead of local to the Exchange server), and the system stressed, there was a risk that the log drive would disconnect, as Figure 12 shows.

Figure 12. Event ID 491 resulting from log failure and Store shutdown (disconnect from NAS)



The key to preventing disconnection between Exchange and the Store device (NAS) is keeping the load under a certain level of latency. The PerfMon counters available for monitoring the level of latency have been made available in Exchange Server 2003 Service Pack 1 and are illustrated in Figure 13. These counters are only visible when enabled by manually adding the registry DWord value "Show Advanced Counters" set to 0x1 in the registry key HKLM\System\CurrentControlSet\Services\ESE\Performance.

For more information, see Microsoft KB 839687, "Microsoft support policy on the use of network-attached storage devices with Exchange Server 2003."

Figure 13. PerfMon counters available for monitoring the level of latency

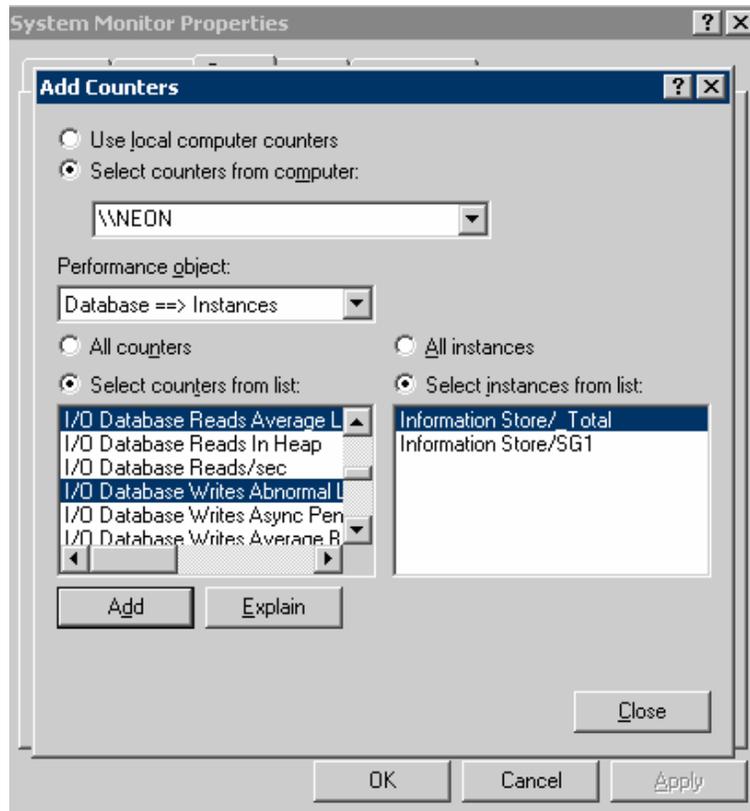
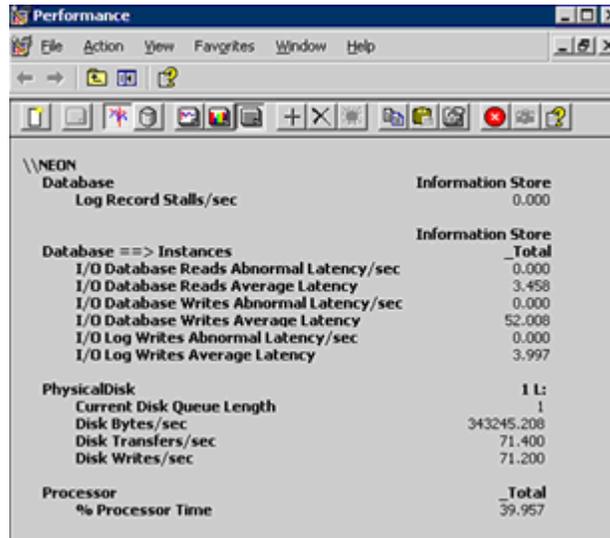


Figure 14 shows live monitoring of latency for the Exchange database drives located on NAS. The Windows System Monitor view in the figure is for a single point in time, but a longer-term capture for a heavily loaded Exchange server during the configuration tests found latencies averaging 200 milliseconds and spiking over 400 milliseconds. (Over time, the average value decreased to around 170 milliseconds for this heavily loaded server.)

Figure 14. Windows System Monitor showing the counters for monitoring latency of NAS drives



Managing Exchange files and services

To manage the Exchange files as you would on a typical Exchange server with local or storage area network (SAN) attached drives, you might want to map a drive to the Windows Storage Server shares. Note that these mapped drives can be used to view files and run Exchange utilities as necessary (for example, ESEUTIL, although it will run fastest directly on the NAS), but the mapped drives cannot be used directly when creating new Stores—they must be created locally and relocated, as explained earlier.

Before restarting the Exchange server, it is recommended that you shut down the Exchange Information Store service on the Exchange server, and then shut down the server. This procedure helps to ensure that all transaction logs are committed to the database and the checkpoint is advanced.

To restart both servers, restart the Windows Storage Server computer first, and then power off and restart the Exchange server. This ensures that all storage provided by the Windows Storage Server computer is available when Exchange Server 2003 starts. If a server operator or administrator is not available to power the Exchange server back on (for example, in a remote-server scenario), then a best practice is to set the startup delay in the boot.ini on the Exchange server substantially longer than the startup delay in the boot.ini of the Windows Storage Server computer.

Antivirus

Note that administrators who are more familiar with deploying NAS for file and print servers must understand how to deploy both file-based and e-mail antivirus scanning. File-based scanning products are often used on PC desktops. The antivirus products used in the Feature Pack test environment are not file-based scanning, as is most often deployed on file servers. Instead, these antivirus products run with Exchange Server 2003 and scan e-mail attachments. A list of products can be found at <http://www.microsoft.com/exchange/partners/antivirus.asp>.

If file-based scanning is deployed to protect the file shares on the Windows Storage Server NAS, the administrator must be careful to exclude the Exchange files from scanning, to prevent inadvertent false positives (that is, matching a viral signature against the binary data stored in Exchange databases, logs, or other files, which can result in a server crash if the file is quarantined by the anti-virus software). This list of file exclusions should include all Exchange databases, logs, SMTP mailroot, and other files such as message-tracking logs. For more details, see Microsoft knowledge base article 823166, "Overview of Exchange Server 2003 and antivirus software" at <http://support.microsoft.com/default.aspx?scid=kb;EN-US;823166#3>.

Full-text indexing

Full-text indexing (FTI) enables the Microsoft Outlook and Microsoft Outlook Web Access e-mail clients to perform full-text searches against both messages and attachments. The index of Exchange server content makes it possible to locate words in attachments (including the document properties) that might otherwise be missed using the basic search capability of Microsoft Outlook. Without FTI enabled, the Exchange server will perform ordinary, character-based searches. FTI is a beneficial function of Exchange, but proper configuration and tuning of the deployment is crucial to avoid risks that may outweigh the benefits. The FTI feature places an additional burden on the performance of the disk subsystem, and the FTI files must be located on the disks attached directly to the Exchange server (rather than on the NAS). Also, note that if you are deploying Microsoft Outlook 2003 in cached mode, it does not benefit from FTI.

For full details, see the HP white paper "Optimizing Full-Text Indexing on Microsoft Exchange Server 2000 and 2003" at <http://h71028.www7.hp.com/ActiveAnswers/Render/1,1027,6637-6-100-225-1,00,00.htm>.

Disaster recovery

For a more comprehensive overview, see the HP white paper "Microsoft Exchange Server 2000 and 2003 Backup and Restore using HP Technology" at <http://h71028.www7.hp.com/ActiveAnswers/Render/1,1027,6535-6-100-225-1,00,00.htm>.

Backup and restore testing

As part of the configuration validation testing, backup and restore tests were performed to measure throughput rates for data restoration. Table 6 shows the measured restore rates for a 24-GB Exchange Store with the SDLT tape device connected to different locations. For this test suite, all servers were attached to a dedicated Gigabit switch and all servers used the HP network teaming driver (in Transmit Load Balancing mode to create a 2-Gbps link).

The conclusion from the tests is that the tape device can be attached directly to the Windows Storage Server or a dedicated backup server and achieve nearly the same performance as if attached to the Exchange server. This can save costs in a multiple server environment by allowing a single dedicated tape device to back up many servers over the network.

Table 6. Measured restore rates for 24-GB Exchange Store with tape devices at different locations

Tape device location	SDLT throughput MB/s	SDLT throughput GB/hr
Dedicated backup server	19.7	69
Exchange server	19.6	69
Windows storage server	17.5	61.5

Recovery storage groups

Recovery storage groups (RSG) is a new feature in Exchange Server 2003. For more detail, see the HP white paper "Exchange Server 2003 Recovery Storage Group and HP ProLiant Server Technologies" at <http://h71028.www7.hp.com/ActiveAnswers/Render/1,1027,6183-6-100-225-1,00,00.htm>.

Working with RSG is slightly different using a Feature Pack environment. When you create the RSG, it will be on a local drive and then you must relocate it. To relocate an RSG to a Windows Storage Server computer, you must first create the RSG on the Exchange server, and then use the Remote Storage wizard to move the storage group, before mounting its Stores or performing any restore operations.

Windows Storage Server future

The next generation of Windows Server (codenamed Longhorn) promises to deliver a new file system, which will enhance new Windows Storage Servers in the areas of finding information (content addressability) and security. Customers adopting Windows Storage Server 2003 now should face a smoother migration path to the new Windows-powered storage devices when Longhorn ships.

For "WinFS Overview—Significant New Storage Innovation for Longhorn," see <http://msdn.microsoft.com/msdntv/episode.aspx?xml=episodes/en/20031211WinFSQC/manifest.xml>.

For more information

HP storage website

The HP website has the latest information on this product, as well as the latest drivers. Access storage at <http://www.hp.com/storage/>. From this website, select the appropriate product or solution.

HP authorized reseller

For the name of your nearest HP authorized reseller:

- In the United States, call 1-800-345-1518.
- In Canada, call 1-800-263-5868.
- Elsewhere, see the HP website for locations and telephone numbers: <http://www.hp.com/>.

HP technical support

Telephone numbers for worldwide technical support are listed on the following HP website: <http://www.hp.com/support/>. From this website, select the country of origin.

Note: For continuous quality improvement, calls may be recorded or monitored.

Be sure to have the following information available before calling:

- Technical support registration number (if applicable)
- Product serial numbers
- Product model names and numbers
- Applicable error messages
- Operating system type and revision level
- Detailed, specific questions

© 2004 Hewlett-Packard Development Company, L.P. The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein. Reproduction, adaptation or translation without prior written permission is prohibited, except as allowed under the copyright laws.

Intel, Pentium, and Xeon are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries. Microsoft and Windows are U.S. registered trademarks of Microsoft Corporation.

5982-7017EN, 11/2004

