

# A Guide to successfully consolidating storage to HP StorageWorks NAS using Quest Consolidator

Matt Alsip – Hewlett-Packard Company  
Brenda Steinsky – Quest Software



Abstract.....	2
Introduction.....	2
A data migration scenario (or, “How hard can it possibly be?”).....	3
Setting the stage.....	3
Addressing the challenge.....	4
If at first you don’t succeed... ..	4
What went wrong? .....	6
A new approach—implementing a strategy for success.....	6
Quest Consolidator—the whole solution approach.....	7
Selecting a storage solution .....	9
Storage 101: matching storage needs with storage solutions.....	9
Minnetonka’s decisions.....	9
Sizing the system.....	10
Sizing the storage .....	12
Configuring the storage solution.....	14
The coup de grâce—migration made easy.....	15
Summary .....	17
Additional resources from HP and Quest Software.....	17
Appendix A—storage assessment inventory .....	18
Appendix B—consolidation project planning checklist.....	21

## Abstract

Customers planning on migrating from a legacy Microsoft® Windows NT® infrastructure to Directory Services, File Services, and Print Services based on Microsoft Windows® 2003 Server technology face a daunting task of designing the new infrastructure; determining hardware and software to upgrade, retire, or consolidate; and planning the migration. To ensure a smooth and seamless migration, customers should thoroughly understand their current environment and develop a well-thought design that will meet current and future needs. Customers can benefit from guidance that outlines the solutions and services available that provide the best value, performance, reliability, and scalability.

HP now offers a compelling range of storage solutions for file serving and print serving, bringing together the most recent technology advances from Microsoft and HP StorageWorks. Quest Software offers a full suite of tools to aid the customer in planning, executing, and managing a migration effort from Windows NT to Microsoft Windows Server 2003. Together, HP and Quest have partnered to provide comprehensive solutions and services for customers looking to upgrade their infrastructure.

This document outlines the primary issues that must be addressed by customers planning to migrate and consolidate their file and print services from Windows NT to Microsoft Windows Server 2003. The discussion includes references to recommended software solutions from Quest Software to aid in the planning and execution of a successful migration, as well as references to recommended storage solutions from HP StorageWorks that are part of a well-planned networking and storage architecture.

## Introduction

*It's 7:00 pm Sunday evening in the server room of Minnetonka Healthcare, Inc. (MHI), and the two administrators working through a server consolidation are helping themselves to another portion of regret. A simple job of moving data from one server to another has turned sour, and with many more servers still waiting to be migrated, it appears that the task is turning out to be much more complex than anyone anticipated.*

As newer generations of computer hardware and software become available to replace the older, obsolete versions, network administrators face the task of migrating critical data and services from one platform to another. At first glance, the prospect of moving data from one computer to another appears to be ridiculously simple. With a network connection, it should require nothing more than a COPY command and a coffee break. However, as any Information Technology (IT) administrator who has lived through the experience can tell you, data migration can be a very complex and time-consuming project. The potential for cost overruns, blown schedules, disruptions to productivity, or data loss is very real in any migration. Because of that, most organizations will often opt to patch and prop up outdated systems, rather than migrate to new ones.

At the same time, the need to migrate is a stark reality in many IT organizations today. Recently, Microsoft announced its upcoming plans to retire support for the Microsoft Windows NT 4.0 operating system. Companies with an investment in Windows NT Server now face the very real prospect of losing phone support, security updates, and bug fixes for the servers that comprise the backbone of their network. In addition, the fragmented storage infrastructure caused by directly attaching dedicated drives to each server (direct attached storage [DAS]) is driving up management costs as each server needs individual attention to ensure the storage is available, secure, backed-up, and safe from malicious viruses and worms.

This dilemma is exacerbated in small- and medium-sized businesses, where budgets and expertise in the IT department may be limited. Few resources can be spared to investigate, analyze, plan, and execute a data migration project. What is needed is expert guidance and tools that simplify the process, without incurring the extra overhead from an army of consultants.

Companies can leverage a great deal of new features and functionality from the latest storage systems. New systems offer native integration to Active Directory, snapshots, quota management, dynamic volume growth, and improved performance. There are also benefits to be gained in new hardware, with faster network interfaces, faster processors, and larger, faster disk arrays. HP and Quest Software provide the migration and storage solutions that companies need.

This white paper outlines the issues involved in a data migration project and discusses best practices available to plan and execute an easy-to-manage migration that is comprehensive and budget friendly and causes minimal impact on end users. In addition, it recommends specific storage solutions to help solve the problems caused by fractured storage and server bloat.

The topic of data migration can include a wide variety of platforms to migrate from and platforms to migrate to. For the purposes of this discussion, it is assumed that the customer wants to upgrade or migrate away from file and print servers that host Microsoft Windows NT 4.0 Server or Microsoft Windows 2000 Server. It is important to note that this discussion is limited to file and print serving only. Server applications, such as SQL Server or Microsoft Exchange, are beyond the scope of this document.

## A data migration scenario (or, "How hard can it possibly be?")

To illustrate the challenges and potential pitfalls involved in migrating and consolidating file and print servers from Windows NT to Microsoft Windows Server 2003, the experience of a fictitious company, MHI, is focused on. A medium-sized corporation, MHI employs approximately 950 people in five separate facilities. Based in Minneapolis, Minnesota, MHI is in the beginning stages of the migration and consolidation process.

This white paper begins by presenting an understanding of the storage dilemma faced by the IT department at MHI and the initial course of action taken toward solving its problems.

### Setting the stage

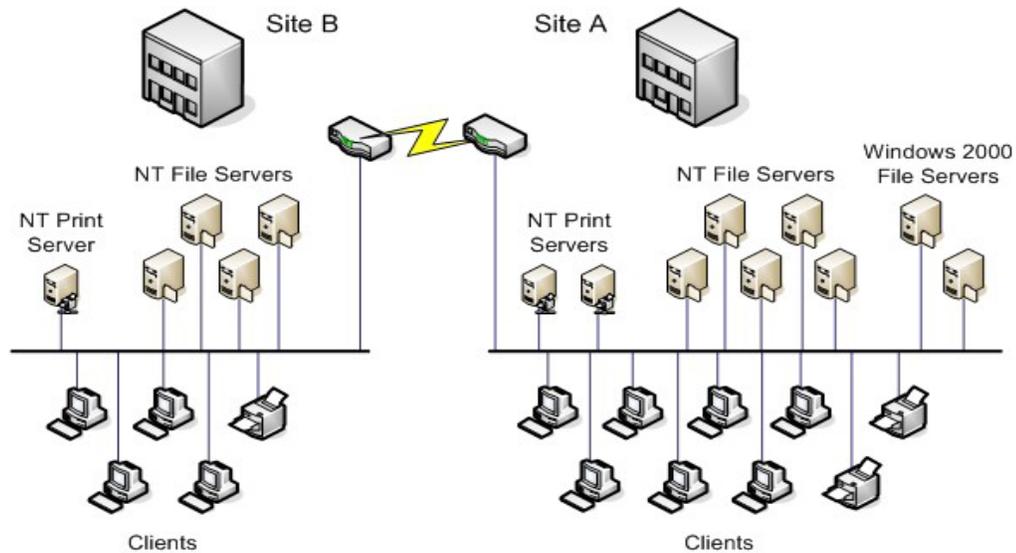
MHI employs five people in their IT department. The network infrastructure is rather modest, even though the company has grown rapidly in the past few years, including a merger with another healthcare provider in the Minneapolis area. Because of MHI's rapid growth and merger activities, the IT team had to accept a less desirable and inefficient legacy structure.

MHI's network is composed of the following components:

- Two Active Directory domains that were recently upgraded from NT domains
  - Domain Controllers also serve as DNS servers, WINS servers (legacy), and DHCP servers
- Exchange servers for e-mail
- Web servers
- Servers for hosting databases for patient records, billing, payroll, and HR records
- 11 file servers, total
  - Nine file servers on Windows NT 4.0 Server
  - Two file servers on Windows 2000 Server systems
- Three print servers on Windows NT 4.0 Server

All of this equipment is physically distributed between two locations (another legacy from the merger). For a graphical view of the file and print-serving environment, see Figure 1.

**Figure 1.** Legacy file and print serving environment



## Addressing the challenge

MHI management and IT staff are aware of just how inefficient this arrangement is, but have been too busy and budget-constrained to do anything about it. As it is, the time spent backing up all those servers is the equivalent of a full-time job. Now, three of the NT servers are getting close to their capacity limits, forcing MHI to make a decision: invest in a new storage platform or patch the existing servers with larger disk drives.

To conserve resources, MHI decides to upgrade one of the newer file servers based on Windows 2000, and then consolidate the storage from the other NT servers on the upgraded server. Since the Windows 2000 server hosted six drive bays for SCSI drives, MHI can fill up all the bays with 146-GB drives amounting to about 730 GB of usable space (using RAID 5 for redundancy). To accomplish this, IT staff must upgrade the following:

- Operating system—Make use of the features in Windows Server 2003
- Memory—Provide adequate memory
- Disk Drives—Purchase six 146-GB SCSI disk drives
- Ethernet NIC Interface—Gain a faster network connection with a 1-Gb interface

## If at first you don't succeed...

Before they invest in these upgrades, they decide to proceed with a small test migration to work out the difficulties of moving over the data from the NT servers. They want to focus on one NT server in their production environment and move 100 user home directories. They design a quick checklist that is similar to the following:

1. Create a tape backup from one of the NT servers.
2. Use the tape backup to restore the data on the Windows 2000 server.
3. Manually re-create the shares on the Windows 2000 server.
4. Manually reconfigure the DNS server to redirect users attempting to connect to the old NT server to the Windows 2000 server.

They decide to work over the weekend during a scheduled maintenance window to minimize the downtime the company would experience.

Friday Night	After multiple announcements to end users by e-mail, the NT server is taken offline and backed up on tape. Simultaneously, the Windows 2000 server is backed up to tape so that a recent image is available in case something goes wrong. As this process is time consuming, the job is allowed to complete overnight.
Saturday Morning	With a fresh copy of the NT server complete, a restore is started on the Windows 2000 server. As this process is also time consuming, the team decides to go to lunch.
Saturday Afternoon	The team returns from lunch and discovers the restore has completed. However, they notice that thousands of files did not get restored to the Windows 2000 server because they were duplicate files. This is unexpected. The team is left to wonder if the duplicate files are different in some way. If indeed they are different files, which files are more important or are they both needed? Since there is no way to tell, the team decides to restore the Windows 2000 server to its original state, clearing all the data from the Windows 2000 server. As they do so, they realize that they must keep track of the shares that are set up on the system so that they can manually restore them later. Finally, they begin to restore from the tape made on Friday night.
Saturday Night	Late in the night, the restore is finally complete. They now have to manually restore all the original shares on the system and their permissions. Next, the IT team decides to restore the Windows NT data to the Windows 2000 server in a separate directory tree to eliminate the potential for having duplicate files in the restore.
Sunday Noon	The restore of the NT data on to the Windows 2000 server is complete. The IT team now works to re-create the shares and permissions that were on the NT server. Since the NT server hosted over 100 home directories of end users, this process is very time consuming.
Sunday Evening	The IT team discovers that a number of the share names are identical between the shares on the Windows 2000 server and the shares on the original Windows NT server. Since all share names on a given server must be unique, some of the shares must be renamed. Renaming the shares defeats the purpose of having the Domain Name Server (DNS) redirect the server name to the new location, because with different share names, the references maintained by the clients still will not work. Instead, the IT team opts to notify all users by e-mail as to the location of the new shares.
Monday Morning	The IT team spends a fair part of the morning assisting end users to manually re-connect their share drives to the new locations. This is because many users on the network do not have enough experience on computers to perform this task by themselves.

## What went wrong?

The migration of data from one server to another turned out to be a very time-consuming, error-prone process that resulted in a lot of downtime. The IT team was exhausted after having worked around the clock for the entire weekend. End users were impacted at every point: Friday afternoon they had to stop working with specific data while a server was backed up; all weekend they could not connect to some network servers; Monday morning they were not able to work until IT staff visited their workstations to “fix” their desktop shortcuts, drive mappings, and links in critical business documents and spreadsheets.

The exercise also exposed some interesting issues that the IT team did not anticipate.

How much data is duplicated on the network? Is the duplicate data relevant? Examining these issues led to more questions such as how much storage is actually wasted with duplicate files, unused files, files that no longer have an owner, and files that should not be stored on network servers (such as music files)?

The team met later in the week to evaluate the process and determine the steps necessary to make improvements. Obviously, more planning would be required not only to nail the correct process but also to determine how much storage is really needed, how to plan for future storage needs, and how to migrate data efficiently.

Beginning their evaluation, the team lists their primary errors with regards to the failed migration:

1. Failed to anticipate the prolonged unavailability of data on two servers
2. Under-estimated the amount of time and effort necessary to manually re-create the shares
3. Forgot about the need to re-create the share security, which in turn, caused “access denied” for many end-user clients
4. Failed to establish a process to effectively handle duplicate shares or folders
5. Introduced the risk of overwriting duplicate folders and files because of the tape restore approach, resulting in data loss
6. Failed to handle the cutover process efficiently, in turn, affecting end-user access to new shares
7. Failed to implement policies and procedures to eliminate the junk data on the network

The IT team had incorrectly assumed that the process was reasonably straightforward, and instead discovered many troubling and time-consuming issues that caused excessive loss of productivity. They also failed to maximize the opportunity at hand by using their consolidation efforts as a means of rationalizing their storage.

## A new approach—implementing a strategy for success

Armed with new resolve and backed by lessons learned, the team now recognizes the need to properly evaluate the migration and storage needs of MHI. They begin by establishing a storage migration project that outlines clear deliverables, milestones, and budget. Next, the team outlines the objectives of the project:

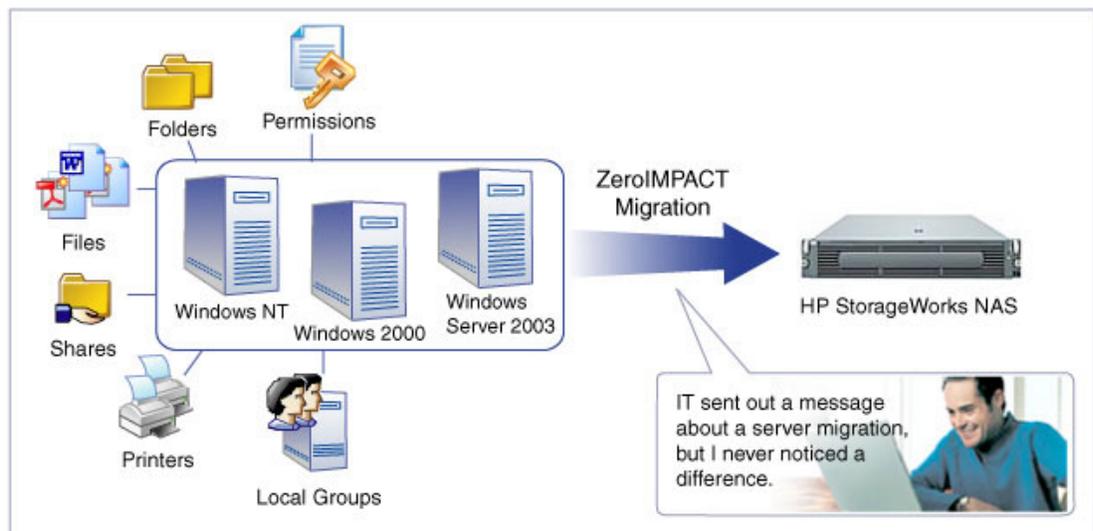
1. Determine the file serving storage needs of the organization for the present and the foreseeable future.
2. Evaluate and procure a storage solution that meets the needs identified in step 1.
3. Establish a data migration process that meets the following criteria:
  - a) Zero impact to end clients
  - b) Minimal impact to IT staff (both in development and execution)
  - c) Reasonable timeframe to execute (one week per server)
4. Establish a budget and schedule.

Before the IT team tackles the next migration, it must do some upfront investigation. Two immediate tasks are identified: (1) determine the storage needs of the organization, and (2) research and evaluate consolidation tools that will meet the criteria of the migration process. The team divides the assignments between its members and goes to work.

## Quest Consolidator—the whole solution approach

In the course of its investigation of data migration tools, the IT team discovers a product that not only offers a comprehensive set of migration tools for Windows, but also offers something else that they had not expected: a storage analysis tool for their network. This combination of features will help satisfy the team's need to evaluate the existing storage infrastructure and plan and execute the migration. The product is Quest® Consolidator from Quest Software.

Quest Consolidator was designed to tackle the unique requirements of moving Windows-based data across computers, domains, and forests. It even handles moving data across systems of differing operating system versions. Quest Consolidator migrates files, folders, permissions, shares, share security (including the local groups and their memberships), and printers. The data and subsequent synchronizations occur online. Users can work with their network files during the entire migration process. Only during a final synchronization, usually performed during a scheduled maintenance window, are users unable to access their data.



The product also contains several advanced features that help administration teams create automated, unattended migration processes that enables them to return to completing everyday tasks. For example, Quest Consolidator:

- Performs intelligent synchronization, transmitting only the delta changes
- Performs online printer migration with built-in support to handle updating drivers no longer supported in Active Directory
- Supports the ability to launch, sync, and pause/resume jobs automatically—according to a pre-defined schedule
- Supports bandwidth throttling on a per project basis—helping reduce the impact on network performance when migrating during business hours or over slow links

- Prechecks to ensure sufficient disks space exists on the new storage to complete the consolidation project
- Prevents unwanted files, such as MP3s or unused data, from being migrated to new storage
- Can apply Windows Inheritance to data coming across from NT servers

After the files are migrated, Quest Consolidator can be used to update end-user workstations. This includes updating desktop shortcuts, drive mappings, printer connections, OLE links and hyperlinks in Microsoft Office documents, spreadsheets, and presentations. This removes the need for manual updates, whether performed by the project staff, the help desk staff, or the end users.

In addition to a comprehensive list of tools to migrate data and update clients, Quest Consolidator contains a built-in storage analysis component called Storage Analyzer. Storage Analyzer is an agentless technology that gathers and reports accurate storage metrics. This information helps administrators perform several storage management tasks such as identifying and deleting unwanted data, compressing unused data, trending server and storage growth over time, identifying server volumes reaching disk full status, and planning consolidation projects. The Storage Analyzer component will help administrators obtain the answers they need to critical questions:

- How much data do I have in my network?
- Where is this data located?
- Who owns this data?
- When was this data last accessed?
- How much data on the network is superfluous?
  - Duplicate data
  - Orphaned data (files whose owner no longer exists)
  - Rarely accessed data
  - Inappropriate or illegal data (for example, MP3 music files, movies, games)

To more effectively evaluate Quest Consolidator, the team acquires a trial version of the software from Quest's website (<http://www.quest.com/microsoft>). The team discovers that the product provides all the functionality it needs to evaluate MHI's storage needs and perform the migration with almost no downtime. In addition, the per-server licenses are reasonably priced, helping to keep the team on budget. Based on the results of its investigation, the team moves forward and purchases the necessary licenses.

Since Quest Consolidator must leverage the functionality of the database engine, the consolidation team installs the Quest Consolidator product on an existing SQL server. After becoming familiar with Consolidator's Storage Analyzer and Data Migration components, the team decides to scan all of the file servers (both NT and Windows 2000). The scan will enable the team to identify and remove unwanted data such as duplicate, orphaned, and prohibited files (such as MP3 files) before the migration is performed. Additionally, the team wants to determine which files are rarely used to archive those files if it ever needs to conserve space.

The initial scan takes only a few hours to run. For this process, the team runs the Enterprise Storage Cleanup report to see how much unwanted data had been identified on the network servers. This report provides them with metrics on the duplicates and "against policy" files as well as orphaned and unused data. The results surprise the team. It turns out that about 35% of the files on their servers are "junk." They quickly identify the location of junk files and use the built-in "delete" script to remove the files on each server to be consolidated. After this task is complete, they take an inventory of the remaining storage and determine that the cumulative space used by all the file servers is about 630 GB. Much to their surprise, they discover that even though three of the file servers are almost full, almost half of the cumulative storage on all the file servers is unused. With this information, the team now has a starting point from which to design their new storage solution.

## Selecting a storage solution

Based on the disastrous results of its first migration attempt, MHI's IT team knows that upgrading the current file and print servers to Windows Server 2003 is an inefficient solution. Having the storage distributed to each individual server requires significant management and maintenance overhead. Plus, the time and cost of performing regular backups, installing security updates, and balancing the storage load among many servers is eating away at the IT budget. To appropriately address these issues, the team knows that it needs a centrally managed storage repository.

### Storage 101: matching storage needs with storage solutions

In the early days of networking, the online storage used by a workstation or server was provided by one or more physical drives that were directly attached to the computer. These solutions came in two types: integrated development environment (IDE) and SCSI. Although SCSI offered more flexibility and scalability than IDE, both solutions ultimately were dedicated only to the systems to which they were attached. These solutions are commonly known as DAS. The three primary drawbacks of DAS solutions are (1) they do not scale very well, (2) storage is inefficiently utilized, and (3) they are costly to manage.

Today, two new storage paradigms have emerged to address the problems of DAS: network attached storage (NAS) and storage area networks (SAN). In the simplest terms, NAS can be thought of as a replacement for the file and print servers currently used on the network. These devices are designed and tuned for dedicated file serving, with features that allow the customer to replace several NT file servers with a single NAS device. This is a very attractive solution for smaller businesses.

On the other hand, SAN is best described as a network architecture for storage. In this arrangement, storage systems (often referred to as storage arrays) stand apart from the servers, and a network infrastructure is used to connect the storage systems to the servers. The SAN paradigm offers a great degree of flexibility and scalability for the customer. However, since SAN usually requires a fair amount of investment in network infrastructure, equipment, and associated expertise, it is normally found in larger businesses.

### Minnetonka's decisions

As it plans to migrate away from the NT servers, the team evaluates the different storage technologies available. The initial investigation reveals two primary options: NAS and SAN. Based on the size of MHI and the configuration of NAS, the team focuses on NAS as their storage solution and begins a more in-depth investigation of the line of NAS devices offered by HP.

The investigation yields promising information. HP NAS solutions are especially attractive to MHI because HP NAS is based on Windows Storage Server 2003. MHI already has a great deal of experience and expertise with Windows-based products. Thus, a Windows-based NAS is guaranteed to be compatible with all Windows clients and is easily managed with the existing skill set. MHI has a rather conservative IT department and is not enthusiastic about experimenting with NAS solutions based on proprietary operating systems that emulate Windows. In addition, HP NAS leverages other HP technologies, such as the HP ProLiant server platform and HP StorageWorks storage arrays. NAS solutions, based on industry-standard components, ensure that the platform is stable, reliable, compatible, and a good value.

HP offers four different NAS solutions, based on the needs of the customer.

HP StorageWorks NAS 1200s	A low-end solution using four hot-swap IDE drives. This solution is great for the budget-conscious customer, as it can provide up to 1 TB of raw storage at a very attractive price. The only drawbacks are that the solution does not scale up and performance is not the best because the system uses software RAID for redundancy.
HP StorageWorks NAS 2000s	A mid-range solution using SCSI drives either internally or externally using the HP StorageWorks Modular Smart Array 30 (MSA30) drive enclosure. Scaling to 27 TB and using hardware RAID for higher performance, this option is more attractive to customers looking for scalability and performance in a NAS solution.
HP StorageWorks NAS 4000s	An upper-range solution that uses storage from a SAN. Since the HP NAS 4000s server does not come configured with any storage, it is often referred to as a "NAS head" or a "NAS gateway." This solution is compatible with all HP SAN arrays, including the HP StorageWorks XP, the HP StorageWorks Enterprise Virtual Array (EVA), and the HP StorageWorks Modular Smart Array 1000 (MSA1000).
HP StorageWorks NAS 9000s	A high-end NAS solution. This product is similar to the HP NAS 4000s, except that it uses the ProLiant DL580 hardware platform instead of the ProLiant DL380 hardware platform. Both the HP NAS 4000s and the HP NAS 9000s may be clustered for high availability.

After analyzing the NAS devices offered by HP, the team quickly focuses its attention on the two products it considers most appropriate for its specific needs: the HP NAS 1200s and the HP NAS 2000s. Although the HP NAS 2000s will cost almost twice as much as an equivalently configured HP NAS 1200s, there are compelling reasons to seriously consider the HP NAS 2000s. The HP NAS 1200s offers up to 700 GB of usable storage, adequately serving the immediate needs of the company. However, storage consumption projections indicate that MHI storage needs will grow to 1 TB within two years, indicating that the HP NAS 2000s may be the more appropriate, long-term selection. With this in mind, the MHI team selects the HP NAS 2000s as the appropriate strategic storage investment.

With an appropriate platform selected, the team must answer two critical questions before making a purchase:

1. What size system is needed to handle the required storage throughput?
2. How many disk drives are needed to meet the current storage needs?

### **Sizing the system**

System sizing is the process of determining how much horsepower the system needs to do the required job. In this process, the following questions must be answered:

1. How many NAS systems are needed?
2. How many processors per NAS system are needed?
3. How much memory per NAS system is needed?
4. How many Ethernet interfaces per NAS system are needed?

System sizing is something of a black art, as many different parameters affect performance in different ways. When sizing a solution for file and print serving, it is important to consider the following factors:

1. How many clients will concurrently connect to the solution for file serving?

This is a difficult question to answer, because even though clients may have the share mapped on their systems, connections are not actually established until the user accesses the data on the share.

2. How many concurrent open files will the solution have to support?

Keep in mind that any one client can have multiple open files at any one time. The number of open files on a server directly affects the memory usage on the system.

3. How many network shares will the solution have?
4. What are the throughput requirements per user?

This will determine how fast a user can read or write a file. This factor may be very important if clients are often reading large files.

5. What proportion of file serving activity will be read/write operations?

The cache in the disk arrays and/or RAID controllers may be adjusted to optimize performance for one type of operation (reading or writing).

6. How many printers are to be supported for print serving?
7. How many concurrent print jobs must be supported?
8. What are the throughput requirements per print job?
9. Will real-time scanning of files be performed with anti-virus software?
10. Will snapshots be used to generate temporary archives?

This can have an impact on the backend storage as the snapshots are taking place. On HP StorageWorks NAS, snapshots are performed by Microsoft's Volume Shadow Copy Service (VSS).

11. Will replication be used?

Replication software consumes processor cycles, impacting performance.

12. What is the average file size being read from and written to the server?

The larger the average files size is for each user, the less likely that caching will be effective in staging data in memory.

The factors that most directly impact performance are the throughput requirements. Since throughput is normally gated by network bandwidth and the number of disk spindles the data is distributed on, performance is usually not improved by adding memory or processors to a NAS server. When performance requirements exceed the specifications of a NAS system, it is generally recommended to use another NAS system. If network bandwidth appears to be the bottleneck, then this limitation may be mitigated with additional network interfaces on the NAS server. The NIC teaming protocol, Fast Ether Channel (FEC), may be used with a qualified switch to effectively add network bandwidth to the system.

The other factors that affect system sizing may be more directly impacted by the amount of memory or the number of processors on a system. User connections and network shares each consume memory. System performance will suffer if the server runs low on memory and has to resort to virtual memory. As long as sufficient memory is allocated to the system, these factors should not impact the system performance. In addition, more processors will improve performance on systems that utilize anti-virus or replication software.

MHI utilizes home directories for each of their users, thus generating a network share for each employee on the network. To ensure that the NAS system has enough resources to accommodate all these shares and concurrent connections, the team plans to upgrade the NAS servers with 2 GB of memory.

While researching a storage solution, the team learns to appreciate the direct relationship between performance and throughput requirements. As a result, its next goal is to calculate the throughput requirements of the solution. Although the bandwidth requirements of the users do not appear to be a major factor, there is some concern that the storage solution might become a bottleneck. Therefore, the team estimates the average utilization and compares that figure against the theoretical throughput of the NAS server. To create an accurate calculation, the team not only considers the users, it also considers printers. For simplicity's sake, the team treats each printer as a client on the network, assigning the same throughput requirements of any user. As a result, the 30 printers on the network are added to the client list.

The IT team's calculation assumes simultaneous burst activity by 5% of the users on the network. In addition, it assumes the bandwidth of each user at approximately 10 Mbps. Although many users are connected to the network with 100 Mbps interfaces, most users will be accessing the storage system from one of the remote sites, over slower wide-area network (WAN) links. Throughput is calculated using the following formula:

$$\text{Throughput} = ((\# \text{ users} + \# \text{ printers}) \times 5 \text{ percent}) \times 10 \text{ Mbps}$$

Accordingly, the team's calculation shows MHI's throughput requirement is 490 Mbps.

$$T = ((950 + 30) \times .05) \times 10 \text{ Mbps}$$

$$T = 490 \text{ Mbps}$$

This estimate is approximately half the theoretical maximum throughput of the Ethernet connection to the NAS server (1 Gb). While not exact, this calculation provides MHI with some confidence that one NAS server with a single gigabit Internet interface will adequately serve the needs of all the clients on the network.

---

**NOTE:**

The calculations and assumptions made in this example are intended for illustration only and do not reflect any standard industry practice for determining system size. System sizing is a complex issue for most environments, and requires special consideration of many factors that may influence the performance of a storage system. The customer is encouraged to review sizing literature from the respective software and hardware vendors before investing in a solution.

---

### **Sizing the storage**

Now that the system has been sized, the IT team must determine how many disk drives to purchase to meet the needs of usable space. In this process, the following questions must be answered:

1. What is the drive capacity that you want to use?
2. What level of data redundancy will be applied to this storage (RAID)?

To meet the needs of usable storage, some calculations must be done to determine how much physical capacity to order. The calculations are not complicated, but they do depend on the method by which the customer plans on protecting the data.

The Smart Array Controller in the HP NAS 2000s includes several options for storage allocation and protection. Each method has a different impact on the amount of storage needed to provide redundancy in the system. The equations associated with each option use variables defined as follows: *space* is the usable space in the RAID set, usually in gigabytes; *n* is the number of disk drives in the RAID set; and *drvcap* is the physical space in each disk drive, usually in gigabytes.

RAID 0 Data striping. This option provides no data redundancy. Therefore, the aggregate storage from all the drives is usable space. However, if one drive fails, then all data volumes in the set are lost. This option is best suited for systems where data loss is not an issue (such as caching servers). Usable storage is calculated as:

$$\text{space} = (n) * \text{drvcap}$$

RAID 1 Data mirroring. This option provides complete data redundancy as every block of storage is duplicated on a separate disk drive. This option requires twice as much storage than what would otherwise be needed for the data volumes. Usable storage is calculated as:

$$\text{space} = (n / 2) * \text{drvcap}$$

RAID 5 Data parity. This option includes an extra drive in the set so that parity information can be calculated and stored for every set of blocks stored on the drives. This method allows any single drive in the set to fail, and yet be able to completely recover all the data. This is a very popular arrangement as it offers the most efficient use of storage and sufficient data protection. Usable storage is calculated as:

$$\text{space} = (n - 1) * \text{drvcap}$$

RAID ADG Advanced Data Guarding. This option is similar to RAID 5 except that two drives are dedicated to parity storage. This method allows any two drives to fail in the set without data loss. Usable storage is calculated as:

$$\text{space} = (n - 2) * \text{drvcap}$$

In addition to these features, the Smart Array Controller allows the customer to specify one or more disk drives to be dedicated as online spares (hot spares). These drives will be used to automatically replace failed drives in a RAID set that includes redundant drives (RAID 1, 5, and ADG).

Another important feature of the Smart Array Controller is online Array Capacity Expansion. This allows more disk drives to be included into an existing RAID set without taking the system offline. Storage management is greatly simplified when existing data volumes are easily extended to include storage from newly added disk drives.

Since MHI's system is sized as a single NAS server with one Smart Array Controller, the team determines that the best approach to storage management on the HP NAS 2000s is to create a single RAID group with sufficient redundancy to protect against a double failure. Based on that determination, the team selects RAID ADG as its storage option, thus requiring two extra drives to account for parity data. Since it is electing to use double protection, the team chooses not to employ an online spare.

To determine how much physical capacity to order for the RAID ADG storage option, the team uses the following formula:

$$\text{Space} = (n) * \text{drvcap}$$

In this formula, the variables are defined as follows: *space* is the usable space in the RAID set, usually in gigabytes; *n* is the number of disk drives in the RAID set; and *drvcap* is the physical space in each disk drive, usually in gigabytes.

Based on the storage analysis it performed earlier, the team knows that MHI needs 630 GB for file serving and 30 GB for print serving. Beyond this, a certain amount is factored in for growth for the next two years (40%) and for buffer space (15%). Given this information, the team specifies 1063 GB

as the absolute minimum storage requirement. It also selects the largest SCSI disk drives available for the HP NAS 2000s: 146-GB drives. The team hopes that this selection will help to ensure a longer support life. With this information, the team plugs in the variables to calculate how many drives MHI needs for the HP NAS 2000s:

$$1063 \text{ GB} = (n - 2) * 146 \text{ GB}$$

$$n = 9.28 \text{ (the team elects to round up to 10.0)}$$

With this information, MHI is now ready to make the investment in new storage. They purchase an HP NAS 2000s (external storage), configured with 10 SCSI 146-GB disk drives.

### Configuring the storage solution

The physical hardware that comprises the HP NAS 2000s is two rack mount components—a NAS server, based on the ProLiant DL 380, and an MSA 30 storage enclosure, which will host the disk drives.



HP NAS 2000s



MSA 30 Storage Enclosure

---

After the hardware is mounted in a rack and physical connections are made, the system can be easily configured for network management using a direct console connection (monitor, keyboard, and mouse). This portion of the configuration assigns the appropriate IP information so that the web management interface can be accessed from the network. After networking is configured, the console components may be disconnected, as all other management may be performed from remote client systems on the network.

Remote management of HP NAS may be handled by a variety of standard Microsoft management tools, like the Microsoft Management Console and the Remote Desktop Connection. However, the recommended tool for managing the HP NAS 2000s is the web-based NAS management interface. This tool offers virtually all the management functions for a file and print server in one place. The HP NAS 2000s also has the specific management tools to manage the disk array built into the NAS management interface. The disk array management tool, known as the Array Configuration Utility (ACU), allows the user to configure the physical disk drives into a logical drive (based on the desired RAID level) that is then presented to the Windows Disk Manager as a new disk to be partitioned.

MHI uses the NAS disk array management tool, ACU, to create a single logical disk from the 10 drives using RAID ADG. From this logical, disk two partitions are created. One small (50-GB) partition for print serving and a separate partition for file serving (using the remaining space).

The MHI IT team is now ready to put Quest Consolidator to work and migrate the file and print services from the legacy servers to the new HP NAS 2000s.

## The coup de grâce—migration made easy

The IT team plans to execute the migration in stages. It schedules each server for migration over a weeklong period to reduce the load on the administrators and to minimize the impact on the organization if something goes wrong. The team expects Quest Consolidator to shoulder the load of moving the data and data security to the NAS server without impacting the network performance. It should also deliver a ZeroIMPACT™ end-user experience—the built-in utilities should update drive mappings, desktop shortcuts, OLE and hyperlinks, printer connections, and home directory profile paths without the need for end-user manual intervention. MHI end users should experience no downtime.

The following describes how their week progressed:

- Tuesday Morning:** The team started the day by creating four data migration projects. Each project contained 25-user home directories. They chose most of the default consolidation options, but added in Active Directory inheritance, placed a filter to prevent unused data from being migrated, and used a “rename” option should a duplicate local group or share be encountered. They also configured a launch and synchronization schedule that would see the data migration begin at midnight and then sync every night at midnight. The entire process takes less than an hour.
- Tuesday Midnight:** The data consolidation projects launches on schedule—automatically. No administrators are onsite to initiate this process.
- Wednesday Morning:** The team observes that the migrations completed overnight. They run a detailed migration activity report to validate how much data was migrated, how long it took to complete the migration, the average throughput rate, and what data, permissions, local groups, and shares were migrated—or not. Everything appeared to be tracking well.
- Wednesday Midnight:** The first synchronization launches automatically according to the pre-configured schedule.
- Thursday Morning:** The team runs a migration “history” report. It verifies that the synchronizations completed successfully and in very little time. The administrators decide to schedule the cutover to the NAS server this coming weekend. The users whose home directories are in the process of being moved are notified by e-mail, and the team asks them to keep their computers turned on over the weekend so a network update can be completed.
- Sunday Afternoon:** The administrators report for work to perform the final sync and cutover. They change the root access control lists (ACLs) on the server, preventing users from gaining access to their network files and launch a synchronization manually. When the sync has completed, they run a migration activity report and confirm that all files were successfully migrated. They check the new storage server and verify the desired directory structure, shares, and permissions are as they intended. This provides them with a green light to perform the cutover. The administrators launch a series of update actions that will ensure end user desktop drive mappings, shortcuts, OLE links, and profile paths are updated using Consolidator’s built-in utilities. An hour later the cutover was completed.

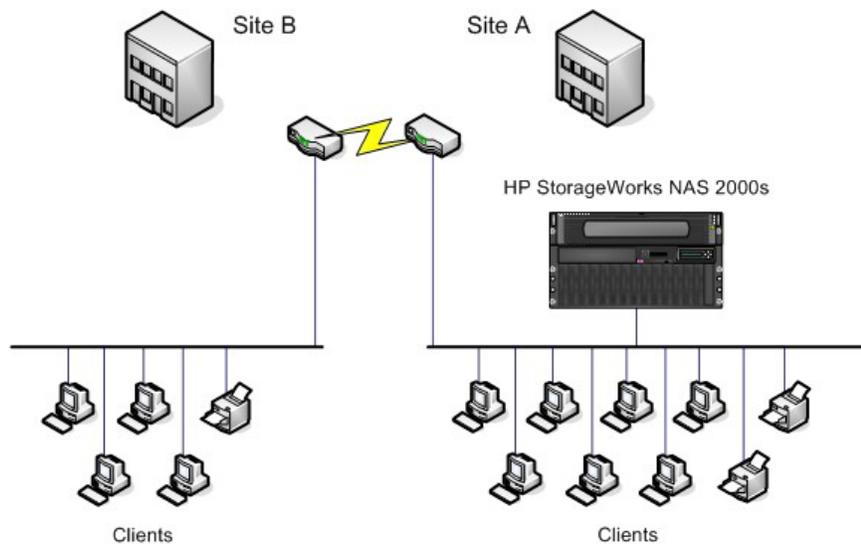
Monday Morning: The administrators report for work a bit earlier than usual to assist with any Help Desk support calls attributed to the server consolidation. As it turns out, the migration appears to have covered all the bases. No migration-related calls are logged.

By the conclusion of the consolidation process, MHI has transitioned to a high-performance, efficient, scalable, and manageable storage solution. Several key factors contributed to the success of the consolidation project:

- Decreased time spent on consolidation efforts—Total hands on time by the administrators to install the product; identify and delete the unwanted files; migrate and sync the data; and perform a successful cutover took only a small number of hours—spread out over a normal work week.
- Supported by a powerful solution—The IT team made the decisions leaving Quest Consolidator to do the work.
- Huge reduction in overtime costs—Very little weekend work by the IT staff was required.
- Significantly decreased downtime and no loss of productivity—Virtually no downtime occurred since the cutover took only a short period of time during a normal maintenance window.
- Eliminated potential for denied access—The business-relevant data and its associated security was migrated to the new HP NAS storage server.
- Eliminated “junk data”—Junk data was successfully identified and totally removed from the network.
- Minimized impact on help desk—The help desk did not have to field a high-volume of calls as a result of the consolidation process.
- Zero impact to the end user—End users were oblivious that a major network change had even occurred.

The following figure represents the change to the storage infrastructure for file and print serving.

Figure2. New file and print serving environment.



## Summary

Migrating file and print services to new platforms is a more complex process than most people realize. Aside from moving the data and associated metadata (security settings) intact, all the dependencies that shared file systems have to a particular server must be addressed when moving the storage to a new location. Even in smaller environments, these issues can cause a severe impact on data availability and end-user productivity. However, with proper planning and investment in appropriate tools, the migration process can be easily managed by the IT administration team.

Migration efforts often expose a lack of management and planning in the legacy storage infrastructure in a company. When transitioning to new storage hardware, a new management paradigm is needed to capitalize on the investment. Companies can use the same tools to not only plan and execute a successful consolidation but also to track storage usage on a continuing basis.

With the help of our fictitious company, Minnetonka Healthcare, Inc., this white paper demonstrated the value of using Quest Consolidator from Quest Software to manage the complexities of a file server migration effort. In addition, this white paper illustrated how Quest Consolidator can be used to effectively evaluate the storage resources on a network as an aid to migration and as a method for post-migration management.

The scenario also highlighted the significant importance of proper planning when making major changes to the storage infrastructure on a network. New NAS and SAN storage technologies offer a great deal more flexibility, scalability, performance, and manageability over the DAS solutions of yesterday. The specific storage solution depends on the application, and, in some cases, may include components of both NAS and SAN.

From the enterprise to the small business, HP StorageWorks NAS offers the most compelling solution because of its foundation of industry-standard ProLiant, HP StorageWorks, and Microsoft components. In addition, the HP StorageWorks NAS solution offers a broad spectrum of supported software and hardware options to intelligently protect and manage the storage solution.

Companies anticipating a file and print migration effort will benefit from the collective storage solutions offered by HP StorageWorks and Quest Software.

## Additional resources from HP and Quest Software

[Storage Optimization Offering](#)—Register today for free tools to help you discover the potential return on investment your company can achieve by consolidating to Windows Server 2003 using Quest Consolidator.

For more information on Quest Consolidator, visit: <http://wm.quest.com/>

For more information on HP StorageWorks NAS, visit: [www.hp.com/go/nas](http://www.hp.com/go/nas)

## Appendix A—storage assessment inventory

Use this template to help you record the information collected during your storage assessment. Each section has a grid for one server. Copy and paste to insert the number you require.

Information from this template will help you:

- Plan and scope your consolidation project
- Understand your hardware and capacity requirements
- Spot potential areas for wasted space cleanup

**Note:**

If you are using **Quest Consolidator's Storage Analyzer**, the Enterprise Storage Overview, Enterprise Storage Cleanup, Server Summary, Server Details, Top Users, Share Summary, Printer Summary, Share Details, and Printer Details reports will provide you with this information.

File and Print Servers	TOTAL # Servers:	
Specifications	SERVER 1	
	OS Version	
	SP #	
	CPU	
	BIOS	
	RAM	
	Physical Disks	
Current Capacity	SERVER 1	
	# Volumes	
	Total space	
	Used space	
	Free space	
	% Used space	
	% Free space	
Volumes Approaching Capacity	SERVER 1	
	Volume 1	
	Volume 2	
	Volume 3	
	Volume 4	

<b>Printers</b>	<b>TOTAL # Printers:</b>	
Specifications	SERVER 1	
	# Printers	
	Printer Name	
	Share Name	
	Port Name	
	Driver	
	Location	
	Comments	
<b>Shares</b>	<b>TOTAL # Shares:</b>	
Specifications	SERVER 1	
	# of Shares	
	Total share size	
	Individual share list	
	Share location	
	Share comments	
<b>Home Directories</b>	<b>TOTAL # Home Directories:</b>	
Specifications	SERVER 1	
	# of Home Directories	
	User	
	Path	
	Size	
	# Files	
	# Folders	
<b>Top Users</b>		
Specifications	SERVER 1	
Top 10 Users/Server	User name	
	Amount of stored data	

Data	<b>TOTAL Amount of Data:</b>	
Enterprise	Potential junk	
	Recoverable space	
	Duplicate files	
	Orphaned files	
	Unused files	
	Wasted space	
	Against policy	
Server	<b>TOTAL Amount of Data:</b>	
	Potential junk	
	Recoverable space	
	Duplicate files	
	Orphaned files	
	Unused files	
	Wasted space	
	Against policy	

## Appendix B—consolidation project planning checklist

Use this checklist to help consider your consolidation project and migration options.

Item	Response	Comments / Notes
<b>Network / Business Restrictions</b>		
Will consolidation activities be restricted to specified hours during the day?	Yes / No	If yes, consider using the “blackout windows” option in Consolidator to automatically pause/resume jobs.
Are there slow links in your LAN/WAN?	Yes / No	If yes, consider using the “tape restore” option in Consolidator.
How fast is your current network?	100/Mb 1/Gb	If Gb, consider increasing the number of copying threads in Consolidator from the default 5 to 6-7.
When is your scheduled maintenance window?		
How much lead time is required for your Change Control Board?		
<b>Data Optimization</b>		
Delete duplicate, orphaned, and inappropriate files from the network before consolidation?	Yes / No	A built-in delete script in Storage Analyzer can handle this.
Remove unresolved SIDs?	Yes / No	Migration options in Consolidator handle this.
Compress unused data as it is migrated?	Yes / No	Migration options in Consolidator handle this.
Make unused file “read only” when migrated?	Yes / No	Migration options in Consolidator handle this.
Prevent redundant folders from being migrated? (WINNT)	Yes / No	Folder filter options in Consolidator handle this.
Prevent unused files from being migrated?	Yes / No	File filter options in Consolidator handle this.
Prevent unwanted files from being migrated? (.MP3)	Yes / No	File filter options in Consolidator handle this.
Implement W2K inheritance when migrating from NT 4?	Yes / No	The W2K Inheritance option in Consolidator handles this.
Maintain “last accessed date” on data?	Yes / No	Migration options in Consolidator handle this.
<b>Directory Optimization</b>		
Do you want to impose/re-impose standard naming conventions on groups, shares, or folders as part of the consolidation?		The “rename” option in Consolidator can handle this for local groups, shares, or folders.
<b>Network Optimization</b>		
Limit the bandwidth used by consolidation activities?	Yes / No	Bandwidth throttling option in Consolidator can handle this.
Migrate group memberships along with references?	Yes / No	Migration options in Consolidator handle this.

<b>Project Optimization</b>		
Prepare migration activity reports?	Yes / No	Run the migration reports in Consolidator.
Receive e-mail notification when migration jobs complete?	Yes / No	E-mail notification option in Consolidator handles this.
Automate the updating of end-user desktop drive mappings, OLE links, printer connections, desktop shortcuts, and profile paths?	Yes / No	Use Consolidator's built-in update utilities to handle this.

© 2004 Hewlett-Packard Development Company, L.P., and Quest Software, Inc.

The information contained herein is furnished for information use only, does not constitute a commitment from HP or Quest Software, and is subject to change without notice. The only warranties for HP or Quest 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. Neither HP, nor Quest Software, shall be liable for technical or editorial errors or omissions contained herein.

Microsoft, Windows, and Windows NT are registered trademarks of Microsoft Corporation.

5982-6723EN, 06/2004

