



HP Scalable NAS for Windows File Server Consolidation

Reference architecture and example customer configurations

Table of contents

Executive summary	2
Windows file servers: Challenges of the traditional approach	2
Drawbacks	3
Large management burden	3
Expensive use of assets	3
Sub-optimal availability	3
HP Scalable NAS solution	4
Scalable	4
Manageable	4
Available	4
Standards-based Windows solution fully supported by Microsoft and HP	5
Key solution building blocks	5
Servers	5
Software	5
Storage	5
Configuration options	5
Best-practice configuration guidelines	6
Servers	6
Storage/SAN	6
Failover and disaster recovery	6
Dedicated development and test cluster	6
Reducing TCO	7
Large business	7
Small to medium-sized business	8
Example customer configurations	9
Large consolidation: Media Company	9
Problem scenario	9
Solution overview	9
Benefits	9
Solution details	10
Solution components	10
Medium-sized consolidation: Professional Services Organization	11
Problem scenario	11
Solution overview	11
Benefits	11
Solution details	12
Solution components	12
Small consolidation: Oil and Gas Company	13
Problem scenario	13
Solution overview	14
Benefits	14
Solution details	14
Solution components	15
Conclusion	16
For more information	16

Executive summary

Organizations deploy Microsoft Windows® file servers to provide many useful services. However, traditional file servers lack scalability. As a result, organizations typically add more file servers piecemeal as their needs grow. This leads to server sprawl and associated problems such as an excessive management burden, low utilization of file servers and storage, and sub-optimal availability.

This document describes, with examples, how organizations can consolidate their Windows file serving environments using HP Scalable NAS, a highly scalable, manageable and available solution. Consolidation eliminates server sprawl, delivering benefits including lower management costs, improved performance, greater utilization, high availability, and the ability to easily scale the solution to meet future needs.

The document includes example customer configurations that represent large, medium and small consolidations. Each example explains the customer's previous environment and challenges, and describes the HP Scalable NAS solution. It illustrates the consolidated environment and lists the HP products required to implement the solution.

The document also describes:

- Challenges of traditional file serving approaches
- HP Scalable NAS benefits and solution building blocks
- Best-practice configuration guidelines
- Example TCO reductions that can be achieved through consolidation

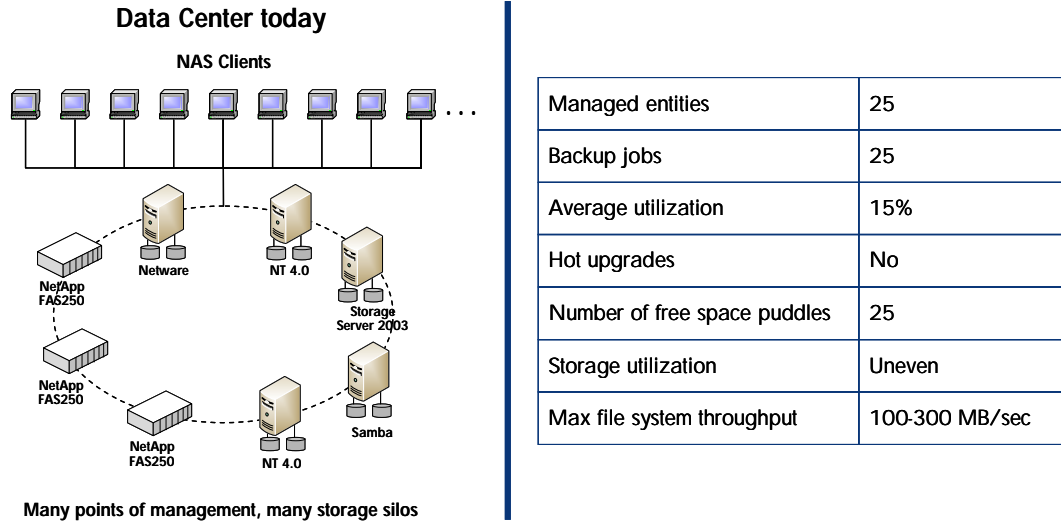
Windows file servers: Challenges of the traditional approach

Organizations of all sizes have deployed Windows file servers to provide many useful services, such as:

- File sharing by groups or departments
- Users' home directories
- User profiles
 - Microsoft Active Directory group profiles
 - Citrix XenApp (formerly Citrix Presentation Server) user profiles for thin-client environments
 - VMware user profiles for virtualized desktops
- Application-specific network storage, including:
 - Microsoft SharePoint Server files
 - Microsoft Systems Management Server image repositories
 - Storage for streaming technologies, such as Microsoft SoftGrid encapsulated application images
 - Client storage and synchronization products such as iFolder

Traditionally, these file servers have typically consisted of servers with direct attached storage (DAS), or traditional dual-head network attached storage (NAS) appliances. However, both of these solutions offer limited scalability. As a result, most organizations have added file servers piecemeal to support their growing needs. This leads to the typical situation shown in Figure 1.

Figure 1: Typical file server environment—Many storage silos and individually managed servers



Drawbacks

Large management burden

- Many separate “puddles” of storage exist; each is a storage silo requiring individual management
- Individual backup, antivirus scans, and patching jobs are needed for each server
- Increasing capacity is complex and difficult, because individual servers have limited scalability
- Replicating data across multiple servers is time-consuming, and replication errors can result in inconsistent business information on different servers

Expensive use of assets

- Server utilization is typically less than 15 percent
- Excess storage capacity is needed to handle peaks and growth, because servers are dedicated to specific purposes
- Dual-head NAS servers include passive servers. This means that organizations are paying for unused processing capacity

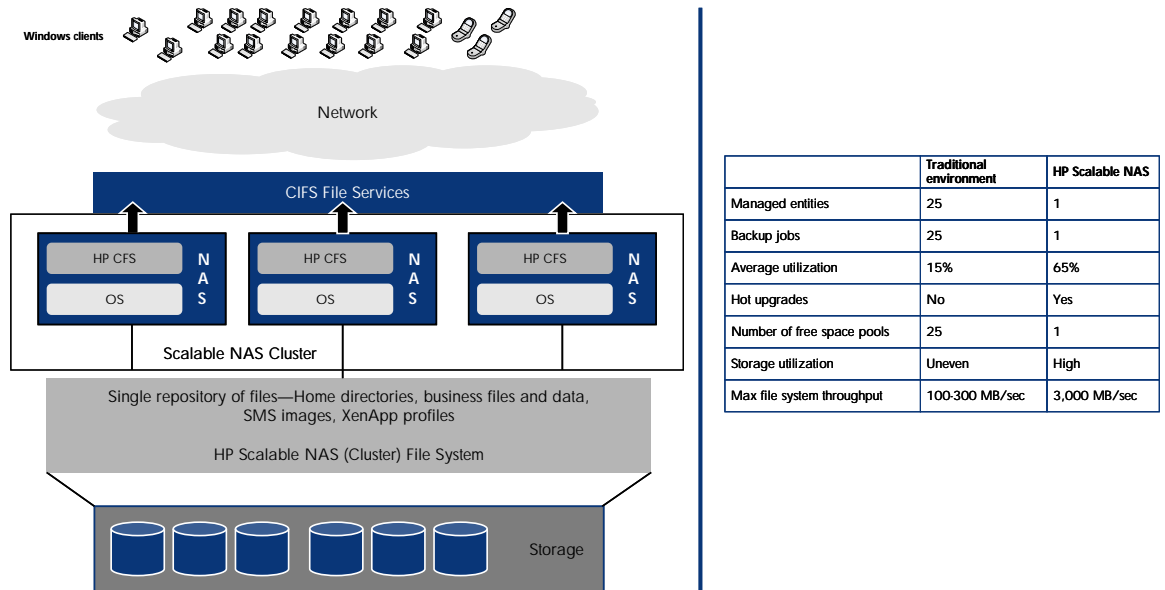
Sub-optimal availability

- Most servers are not fault tolerant, increasing business risk
- Considerable planned downtime is needed for maintenance, which increases overall cost
- Traditional failover approach using active/passive clustered pairs is complex to manage, relatively slow, and may require administrator intervention

HP Scalable NAS solution

HP Scalable NAS addresses the challenges of traditional approaches by providing a highly scalable, manageable, available solution that can be used to consolidate existing environments, reducing server count by up to 75 percent.

Figure 2: HP Scalable NAS deployment model addresses traditional NAS file server limitations



Scalable

HP Scalable NAS uses a shared data architecture that enables organizations to combine up to 16 standard servers into a cluster. This provides file serving performance that scales in a near linear fashion to more than 3 GB per second. HP Scalable NAS (PolyServe) software binds these servers together so that they operate as a single highly-scalable NAS system accessing a shared pool of up to 512 TB of storage. All nodes can cache and serve all files; clients can be load-balanced across all nodes. This greatly increases utilization, eliminating wasted storage puddles and unused processing capacity.

Manageable

Organizations can manage the entire cluster as a single entity from a single console. They can run a single backup and antivirus scan for the entire storage pool, rather than separate backup jobs and antivirus scans on many individual servers. HP Scalable NAS can flexibly allocate storage from a single pool of free space, eliminating the restrictions of fixed data partitions.

Available

HP Scalable NAS provides high availability—often 99.999 percent, equivalent to less than five minutes per month of planned downtime—without wasted processing capacity. Storage capacity can be added on-the-fly. If deployed with the recommended storage area network (SAN) architecture, there is no single point of failure (SPOF). All nodes are active. A single extra server node can provide failover capabilities for the entire cluster. This failover capability can also be used to perform server maintenance without downtime.

Standards-based Windows solution fully supported by Microsoft and HP

HP Scalable NAS includes a symmetric cluster file system that runs on Windows Server 2003 R2. It provides native Common Internet File System (CIFS) support for high performance, and is integrated with Microsoft Active Directory and Microsoft Distributed File System. HP worked directly with Microsoft, licensing the Microsoft Installable File System (IFS) kit, in order to develop a Windows-compatible cluster file system that would support common enterprise-class Windows applications.

HP Scalable NAS works with industry-standard antivirus, backup and management products, enabling seamless integration into existing environments. A list of backup and antivirus products is at <http://h18000.www1.hp.com/products/storageworks/ebs/index.html>

Integration with Microsoft System Center Operations Manager 2007 means organizations can use the OpsMgr 2007 console to monitor HP Scalable NAS along with the rest of their Windows environment. (Please visit www.hp.com/go/polyserve/scom)

Key solution building blocks

HP Scalable NAS solutions can include a variety of servers, software, and storage options, depending on customer requirements. Representative information is included here; details of specific current options can be found at www.hp.com/go/scalablenas

Servers

HP ProLiant DL380/DL385 Servers: Standards-based rack-mounted servers with one or two Intel or AMD multi-core processors and up to 64 GB RAM. Each server can be configured with multiple gigabit Ethernet (GbE) and Fibre Channel adapters.

HP BladeSystem: An integrated system of enclosures and modular blades, including HP BladeSystem c-Class Server Blades with one or two Intel or AMD multi-core processors and up to 128 GB RAM. Each server can be configured with multiple GbE and Fibre Channel adapters.

Software

HP Scalable NAS (PolyServe) Software: Software that enables up to 16 servers to act as a scalable, easily managed, high-availability cluster accessing a shared pool of storage.

Storage

HP StorageWorks XP Disk Arrays: Enterprise-class arrays with maximum internal capacity ranging from 236 TB to 1.13 PB.

HP StorageWorks Enterprise Virtual Arrays (EVA): High-availability, high-performance SAN arrays that are easy to deploy and manage for midsized and enterprise customers with a maximum internal capacity ranging from 96 TB to 240 TB.

HP StorageWorks Modular Smart Arrays (MSA): Affordable entry-level SAN arrays with up to 48 TB capacity.

Configuration options

Organizations can buy HP Scalable NAS in different packaged configurations, depending on their requirements and existing environment:

HP StorageWorks 4400 Scalable NAS File Services: An integrated storage solution based on an initial three-node server cluster, expandable to 16 nodes, with HP StorageWorks EVA4400 storage.

HP StorageWorks Enterprise File Services (EFS) Clustered Gateway: A cluster of up to 16 server nodes with HP Scalable NAS (PolyServe) software. This is configured as an initial cluster with two nodes, plus additional nodes as required. HP Clustered Gateway nodes can be used with common Fibre Channel based storage arrays that the customer may already own.

HP StorageWorks Scalable NAS (PolyServe) Software: Scalable NAS software that organizations can use to create a cluster based on their existing servers and disk storage.

Best-practice configuration guidelines

To ensure high performance and availability, best practices include the following configuration guidelines:

Servers

- At least four Network Interface Card (NIC) ports per server or blade. This provides primary and fallback networks for user data and for communications within the cluster.
- NIC teaming for user data. This allows automated failover of TCP/IP traffic or throughput aggregation for maximum performance.
- Dual-port Host Bus Adapter (HBA) in file servers. This makes sure that each cluster node has at least two paths to the SAN fabric for high availability and throughput aggregation.

Storage/SAN

- Dual Fibre Channel fabric. Together with Dual-port HBAs and Dual Controller Storage Arrays this makes sure that the SAN infrastructure has no Single Points of Failure (SPOF).
- SAN Fabric Zoning. Placing each cluster node, along with its storage, in a SAN zone optimizes performance and provides operational isolation.
- Dual Controller Storage Arrays. Making sure that no SPOF within the storage subsystem requires an enterprise class array with dual controllers; active:active array architecture; and support for selectable RAID levels, disk sparing, and multiple power supplies and cooling fans.
- Physical Storage Design. Define a small set of standardized Logical Unit Number (LUN) sizes appropriate for the cluster's operational usage. This allows LUNs to be easily re-used and allocated on demand.
 - Avoid creating LUNs that share physical devices that are used by multiple workloads with conflicting resource requirements.
 - Avoid creating a single large LUN. This is inflexible and limits performance.
- Logical Storage Design. Divide users into groups so that each group uses a different file system volume. This increases performance and isolates any performance issues so that they do not affect multiple groups of users.

Failover and disaster recovery

To make sure failover can occur if a server node fails:

- For each highly available File Serving Service, define a primary and at least one failover node.
- Configure the cluster so that there will be enough surviving nodes for File Serving Service instances to failover to in the event of a server failure.
- Make sure each server has enough resources to run its designated failover instances.

Dedicated development and test cluster

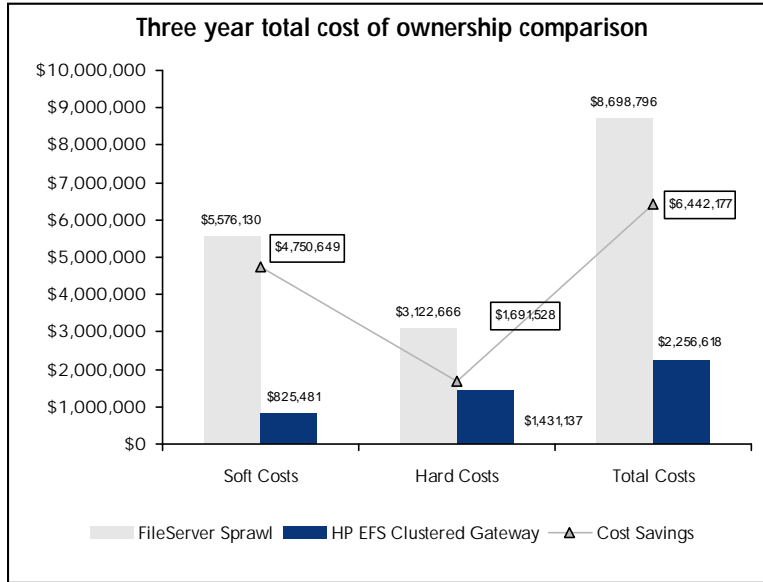
Use a dedicated development and test cluster. This provides an environment for testing configuration changes and updates before deployment, substantially reducing risks to the production environment.

Reducing TCO

Consolidation with HP Scalable NAS can substantially reduce TCO¹ due to benefits including reduced management costs, improved availability, and lower hardware and software costs.

Large business

Figure 3: Consolidation from 100 file servers with SAN and DAS to 10-node HP EFS Clustered Gateway



Summary of Consolidation Results


Summary	File Server Sprawl	HP EFS Clustered Gateway	Improvement
Server count	100	10	90%
Server utilization rate	12%	50%	38%
Total storage managed (TB)	36.0	18.4	49%
Percentage of data highly available	8%	100%	92%
Power usage (3 year cost)	\$122,149	\$13,633	89%
Net present value savings (\$M)	N/A	\$5.29	N/A
Annualized ROI (%)	N/A	95%	N/A
3 Year TCO reduction (%)	N/A	74%	N/A

Assumptions: Storage needs growing at 25% per year; original setup included 92 standalone file servers and four active:passive pairs for a total of 100 servers; 20 of the servers with 500 GB SAN storage each, 80 with 400 GB DAS each; cost of SAN storage \$1.80 per GB; cost of downtime \$2,000 per hour; soft costs based on fully burdened administrator cost of \$100,000 per year.

¹ Savings may differ because they are based on each customer's environment, and figures are provided "as is" and without warranty of any kind.

Small to medium-sized business

Figure 4: Consolidation from 18 file servers to three-node HP EFS Clustered Gateway



Summary of Consolidation Results

Summary	File Server Sprawl	HP EFS Clustered Gateway	Improvement
Server count	18	3	83%
Server utilization rate	15%	50%	35%
Total storage managed (TB)	4.5	2.2	51%
Percentage of data highly available	0%	100%	100%
Power usage (3 year cost)	\$21,987	\$4,090	81%
Net present value savings (\$M)	N/A	\$0.66	N/A
Annualized ROI (%)	N/A	61%	N/A
3 Year TCO reduction (%)	N/A	64%	N/A

Assumptions: Storage needs growing at 20% per year; 250 GB DAS per server before consolidation; SAN storage costs of \$2 per GB; cost of downtime \$1,000 per hour; soft costs based on fully burdened administrator cost of \$100,000 per year.

Example customer configurations

Large consolidation: Media Company

Problem scenario

IT was being consolidated due to corporate mergers and the centralization of departmental IT. As part of this, the company was consolidating six data centers into two. The environment consisted of 96 file servers supporting 16,000 desktop clients, with a pure Windows infrastructure.

File server uses included:

- User Profiles (Citrix Presentation Server and Windows desktops)
- Home Directories
- Group Shares
- Remote synchronized backup (iFolder)
- Software distribution repository for Microsoft SMS
- Encapsulated application images for Microsoft SoftGrid

Key challenges included:

- Management of 96 individual servers was extremely resource-intensive
- The servers were not scalable and offered limited performance

Solution overview

The company consolidated 96 file servers into two HP EFS Clustered Gateway clusters, each serving 50 percent of the users. This provides a native CIFS solution that is highly available, scalable, and offers better performance.


Benefits

The solution has reduced infrastructure and operating costs by more than 50 percent, and improved peak load performance by more than 200 percent.

The company now manages two clusters rather than 96 individual servers, and has achieved a huge reduction in the management burden. For example, only two backup jobs are needed instead of 96. The company has seen similar reductions in antivirus scans, patching, hardware upgrades, and storage management effort.

Solution details

Figure 5: One of two eight-node HP EFS Clustered Gateway clusters based on an HP BladeSystem with HP EVA8100-class storage.



The full deployment was two eight-node HP Scalable NAS clusters, each based on an HP BladeSystem with HP EVA8100-class storage, plus a four-node cluster used for development and testing.

Solution components

- 3 HP BladeSystem c7000 Enclosure
- 20 HP ProLiant BL460c Server Blades, each with
 - 2 Quad-Core Intel® Xeon® Processors
 - 12 GB RAM
 - 4 HP Dual Port Gigabit PCI-E Server Adapters (NIC)
 - 1 HP Dual Channel 4 GB PCI-E HBA
- Windows Server 2003 Enterprise Service Pack 2 x64 R2 edition
- 40 CPU license for HP Scalable NAS software, plus HP Scalable NAS Software Installation and Startup
- 2 HP EVA8100, each with 3.4 TB storage
- 1 HP EVA4400 dedicated development and test cluster with 3.4 TB storage

Necessary SAN switches and other networking components

Medium-sized consolidation: Professional Services Organization

Problem scenario

The company wanted to achieve the benefits of consolidation while supporting a full disaster recovery capability. The previous environment was based on Microsoft Cluster Server (MSCS). As the company had outgrown the capabilities of a single cluster, it had implemented three active clusters to support its 5000 desktop clients. Disaster recovery capabilities were provided by three passive clusters located in a separate disaster recovery data center about two miles away. HP StorageWorks Continuous Access Software replicated data between the two sites; HP Data Protector Software provides automated backup.

File server uses included:

- User profiles
- Home directories
- Group shares

Key challenges included:

- Managing three separate clusters and storage pools was complex and resource-intensive
- Upgrading was a difficult administrative operation
- Users needed to be divided between servers, resulting in unbalanced server loads and uneven performance
- Maintenance required downtime
- Storage was expected to grow rapidly to 4 TB

Solution overview

The company consolidated the environment using a single HP Scalable NAS eight-node stretched cluster spanning the two sites. The native CIFS solution is highly available and scalable, and offers better performance.

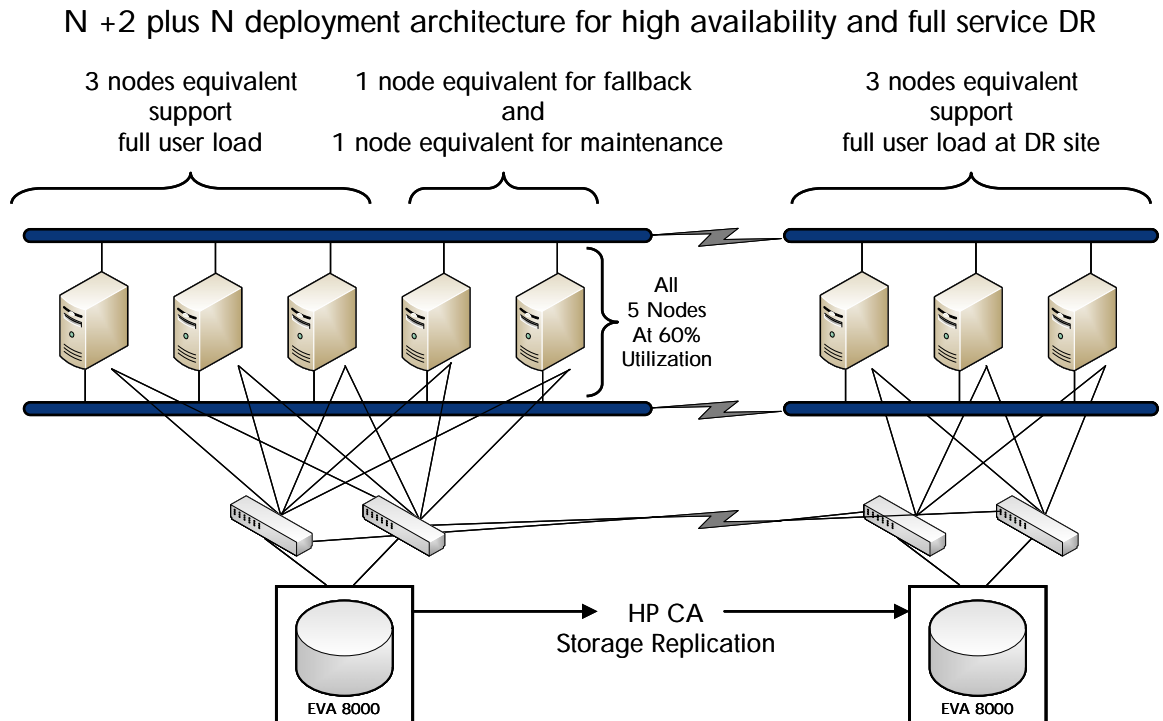
Benefits

The company can grow storage more easily because their volumes are no longer limited by Microsoft Clustering Service constraints. They have simplified disaster recovery management using a single cluster instead of the previous three clusters. There is no downtime required for maintenance because workloads can easily be moved between nodes in the cluster.

They now have user workloads automatically balanced across available servers within the cluster. They enjoy much simpler storage management, including having a single backup job instead of the three previously required. Managing growth is much easier as they can incrementally add storage capacity and/or server (node) capacity.

Solution details

Figure 6: An eight-node stretched HP EFS Clustered Gateway (Initial 2 node) cluster spread across the two sites



- Five of these nodes are at the production site; three handle the full user load, one is used for backup and one for support and maintenance
- Three nodes at the disaster-recovery site are capable of handling the full user load
- Two HP StorageWorks EVA 8000 arrays, one at each site
- HP StorageWorks Continuous Access EVA Software provides synchronous storage replication from the production site to the disaster recovery site

Solution components

1 HP EFS Clustered Gateway Initial 2 node Cluster

Includes Installation and Startup Service and two HP EFS Clustered Gateway nodes.

Note: In addition to basic setup and configuration, customer training is provided on basic operations. This can be configured to include a rack, FC switches, network switch, and all necessary cables.

- 6 HP EFS Clustered Gateway additional nodes (DL380 G5-WSS), for a total of 8 nodes.
Each cluster node includes an HP DL380 G5 ProLiant server, Windows Storage Server x64 2003 R3 OS, HP Scalable NAS software with the CIFS File Serving option, Clustered Volume Manager option.
Each HP DL380 G5 ProLiant Server includes:
 - 2 Quad-Core Intel Xeon processors
 - 12 GB RAM
 - HP Dual Channel 4 GB PCI-E HBA
 - HP Dual Port Gigabit PCI-E Server Adapter (NIC)
 - Embedded Dual Multifunction Gigabit NICs
 - Smart Array Controller with 256 MB cache (RAID 0/1/5)
 - 72GB 3G SAS 10K SFF hard drive
- 2 HP StorageWorks EVA 8000, each with more than 4 TB storage using 146 GB 15K hard drives
- 1 HP StorageWorks Continuous Access EVA8000 Series Unlimited Software License

Small consolidation: Oil and Gas Company

Problem scenario

The company operated a Citrix XenApp server (Presentation server) environment, using more than 60 servers running the Citrix software to support the company's 1600 users. The Citrix XenApp Servers were connected to four active:passive pairs of Windows file servers where user profiles and user data is stored. Each file server pair supported 400 users, providing 20 GB personal data storage for each user—a total of 800 GB per file server. Storage administrators also over-allocated 10 percent disk capacity for user data growth on each server. In addition, each file server contained 3 TB of time-sensitive pricing and other corporate information, which was replicated across the servers to make sure that all users have access to the same data. Total storage is 20.8 TB. The company was experiencing rapid growth, both organically and through acquisitions.

File server uses included:

- User home directories
- Citrix profiles
- Replicated corporate and pricing data

Key challenges included:

- Performance limitations: 10-second response times during peak periods such as morning user base logins and end of day saving of user's desktops, applications, and data.
- File servers cannot scale to overcome performance problems or accommodate growth.
- Solution is inflexible and hard to manage.
- Users needed to be divided between servers, resulting in unbalanced server loads and uneven performance.
- Replicating data is time-consuming, and replication failures result in out-of-date pricing information on some servers, causing serious business issues.
- Inefficient, expensive use of assets: Only 3 TB of corporate data needs to be shared among all users—but the data is replicated across four servers, so a total of 12 TB is needed. This represents an additional consumption of 9 TB with initial acquisition cost over \$10,000 per TB.

Solution overview

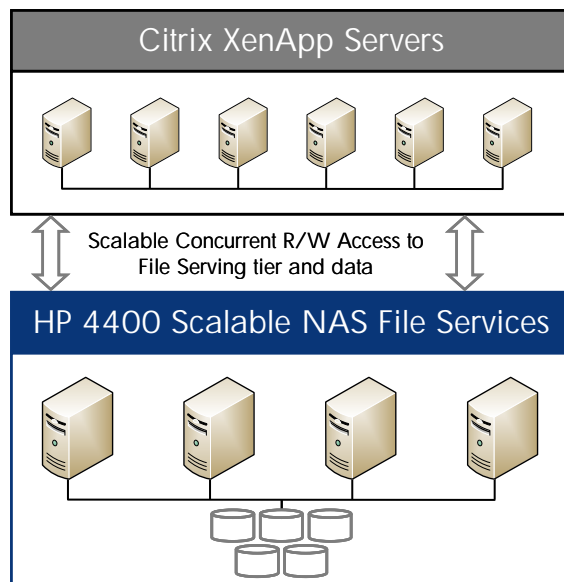
The company initially implemented HP StorageWorks 4400 Scalable NAS File Services, a fully configured packaged solution including three server nodes and storage. This improved performance, removed the need for data replication. It reduced the storage footprint from 20.8 TB to 11.4 TB. Later, the company needed additional capacity after it acquired another company; it was able to easily accommodate the growth by adding a fourth cluster node.

Benefits

- More than 300 percent performance increase. User desktop response times cut down to less than 3 seconds during peak periods
- Solution easily accommodated growth due to acquisition; company can independently scale application and file-serving capacity as needed. Additional users are automatically load-balanced among servers
- Simpler administration
- Improved availability; controlled failover allows maintenance without downtime
- Ensured users have “one view of the truth.” Data replication is no longer necessary, eliminating the risk of inconsistent business data
- More efficient use of storage—no need for replicated data or over allocation
 - Uses only 3 TB rather than 12 TB for corporate data
 - Storage allocated for user data growth reduced by 50 percent to 400 MB

Solution details

Figure 7: HP StorageWorks 4400 Scalable NAS File Services deployment



Solution components

HP StorageWorks 4400 Scalable NAS File Services, a pre-configured solution including:

- 3 DL380 G5 Clustered Gateway server nodes
- DL320 G5p Management Server
- Windows Storage Server x64 2003 R2
- HP StorageWorks Business Copy EVA for local data replication
- HP StorageWorks Continuous Access EVA for remote data replication
- Installation and startup

Solution includes 4.8 TB of storage as standard; the company added 15x 450 GB 15K hard drives to provide a total of 11.4 TB of storage.

- 1 HP DL380 G5-WSS Clustered Gateway additional node, for a total of 4 nodes.

Each cluster node includes HP DL380 G5 ProLiant server, Windows Storage Server x64 2003 R3 OS, HP Scalable NAS software with the CIFS File Serving option, Clustered Volume Manager option.

Each HP DL380 G5 ProLiant Server includes:

- 2 Quad-Core Intel Xeon processors
- 12 GB RAM
- HP Dual Channel 4 GB PCI-E HBA
- HP Dual Port Gigabit PCI-E Server Adapter (NIC)
- Embedded Dual Multifunction Gigabit NICs
- Smart Array Controller with 256 MB cache (RAID 0/1/5)
- 72GB 3G SAS 10K SFF hard drive

Conclusion

HP Scalable NAS solutions can be used to consolidate large, medium, and small Windows file server environments. Using HP Scalable NAS solutions, organizations can implement highly scalable, manageable, and available solutions to reduce administrative burden, and hardware and software costs, while realizing other benefits including higher performance and the ability to more easily accommodate business growth.

For more information

<http://www.hp.com/go/scalablenas>

http://h18006.www1.hp.com/products/storageworks/scnas_4400/index.html

Technology for better business outcomes

© Copyright 2008 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.

Linux is a U.S. registered trademark of Linus Torvalds. Microsoft and Windows are U.S. registered trademarks of Microsoft Corporation. UNIX is a registered trademark of The Open Group.

4AA2-3471ENW, November 2008

