



realtimepublishers.comtm

The Definitive Guidetm To

Building Highly Scalable Enterprise File Serving Solutions



Chris Wolf

Chapter 2: Taming Storage Growth—A Modern Perspective	22
Current Storage Problems	22
Availability	22
Growth	23
Management.....	23
Expanding Backup Windows.....	23
Existing Storage Solutions	24
SAN.....	24
NAS Filers	24
DFS	25
Virtualization	25
Storage Virtualization	25
Server Virtualization.....	32
Virtual Machines.....	32
Shared Data Clusters.....	33
Comparing Virtual Machines and Shared Data Clusters	34
Examining Unappliance vs. Appliance Solutions.....	35
Proprietary vs. Open Solutions	36
Volume Economics.....	37
Integration with Existing Infrastructure and Investments.....	37
The Scalability Dilemma	37
Backup Challenges.....	38
Taming Server and Storage Growth—the Non-Proprietary Approach.....	40
Storage Consolidation via SAN	40
Server Consolidation via Clustering	41
Planning for Growth While Maintaining Freedom.....	41
Summary	42

Copyright Statement

© 2005 Realtimepublishers.com, Inc. All rights reserved. This site contains materials that have been created, developed, or commissioned by, and published with the permission of, Realtimepublishers.com, Inc. (the "Materials") and this site and any such Materials are protected by international copyright and trademark laws.

THE MATERIALS ARE PROVIDED "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, TITLE AND NON-INFRINGEMENT. The Materials are subject to change without notice and do not represent a commitment on the part of Realtimepublishers.com, Inc or its web site sponsors. In no event shall Realtimepublishers.com, Inc. or its web site sponsors be held liable for technical or editorial errors or omissions contained in the Materials, including without limitation, for any direct, indirect, incidental, special, exemplary or consequential damages whatsoever resulting from the use of any information contained in the Materials.

The Materials (including but not limited to the text, images, audio, and/or video) may not be copied, reproduced, republished, uploaded, posted, transmitted, or distributed in any way, in whole or in part, except that one copy may be downloaded for your personal, non-commercial use on a single computer. In connection with such use, you may not modify or obscure any copyright or other proprietary notice.

The Materials may contain trademarks, services marks and logos that are the property of third parties. You are not permitted to use these trademarks, services marks or logos without prior written consent of such third parties.

Realtimepublishers.com and the Realtimepublishers logo are registered in the US Patent & Trademark Office. All other product or service names are the property of their respective owners.

If you have any questions about these terms, or if you would like information about licensing materials from Realtimepublishers.com, please contact us via e-mail at info@realtimepublishers.com.

[**Editor's Note:** This eBook was downloaded from Content Central. To download other eBooks on this topic, please visit <http://www.realtimepublishers.com/contentcentral/>.]

Chapter 2: Taming Storage Growth—A Modern Perspective

Chapter 1 introduced the vast array of storage and file serving solutions available today. It is important to understand how each file serving and storage technology works, and equally crucial to know which technology is right for your organization's needs. This chapter provides a detailed examination of the current storage and growth problems facing IT and explores what server, storage, and application vendors are doing to address these problems.

Current Storage Problems

Most organizations face several problems with their storage infrastructures, most notably:

- Availability
- Growth
- Management
- Backup window expansion

This section will look at the root causes for each of these problems, setting the foundation for a later discussion of how vendors are using technology to address these issues.


Availability

Availability refers to data being obtainable when a user or application needs it. Because *need* is a relative term, the definition of availability can vary from one organization to the next. For a small medical office, availability probably means access to resources from 8:00 AM to 6:00 PM Monday through Friday. For an ecommerce Web site, availability means 24 × 7 access to data. For most organizations, availability likely falls in the middle of the two previous examples. Regardless of an organization's definition of availability, the performance of IT staff is often measured by the availability of data.

Several problems can derail data availability:

- System hardware failure
- System software failure
- Power failure
- Network failure
- Disk failure

Fortunately, many of the single points of failure that prevent data availability can be overcome with redundancy. An organization can overcome power failure through the use of UPSs and backup generators. Reliance on additional switches, routers, and redundant links can help prevent network failure; RAID can help overcome data availability problems that result from disk failure.

 Chapter 3 will focus on how to overcome these issues so as to ensure data availability.

System hardware or software failure is often more difficult to rebound from. To overcome this challenge, several solutions are now available, including NAS and shared data clusters. These technologies will be examined later in this chapter.

Growth

Another problem facing administrators today is growth. As storage requirements and client demand increases, how do you accommodate the increase? To put growth into perspective, according to IDC's Worldwide Disk Storage Systems Tracker, disk factory revenue grew 6.7 percent year over year, as reported in the first quarter of 2005. The report also noted that the 6.7 percent growth value represented 8 quarters of consecutive growth. The increase in reviews has occurred despite the fact that the cost per gigabyte of disk storage continues to fall. For example, the 2005 Worldwide Disk Storage Systems Tracker also reported that capacity continues to grow at an exponential rate, with 2005 year-on-year capacity growing 58.6 percent.

Market trends have shown that nearly all administrators face growth. One of the major problems is how to effectively manage it. Is the ideal solution to continue to add capacity or is a better solution to rebuild the network infrastructure so that it can effectively scale to meet the needs of future growth? Many administrators are deciding that now is the time to look at new ways of managing data, as earlier architectures are not well-suited to meet the continued year-on-year demands of growth.

Management

With growth comes additional headaches—the first of which is management. As online storage increases, what is the best way to effectively manage the increase? If each server on the LAN is using local DAS storage, this situation creates several potential bottlenecks, data islands, and independently managed systems. If client demand is also increasing, is the best option to add servers to the LAN to deal with the heavier load? Ultimately, the problem that is hurting data management today is that many administrators are trying to use traditional architectures to deal with modern problems.

Expanding Backup Windows

As the amount of data grows, so do the backup windows for many organizations. With traditional servers with DAS storage and LAN-based backups, it has become almost impossible to back up servers over a LAN within the time of a backup window. This challenge has resulted in many organizations altering their backup schedules and doing fewer full backups in order to have backups complete before business starts each morning. To deal with the issue of expanding backup windows, many IT shops have considered or have already deployed solutions such as NAS, SAN, DFS, and virtualization. The following sections explore the part that each of these technologies plays within the network.

Existing Storage Solutions

There are many vendors in the market arena that offer products to solve the current storage problems of today. Although these solutions can ease the management burden of an administrator, the solutions' one-size-fits-all approach doesn't guarantee success. It's the responsibility of the organization and the IT staff to understand each of the available storage and consolidations solutions, then select the solution that best fits the company's mission.

SAN

For many organizations, a SAN is often the answer for consolidating, pooling, and centrally managing storage resources. The SAN can provide a single and possibly redundant network for access to all storage devices. Depending on the supporting products purchased by an organization, a SAN might also provide the ability to run backups independent of the LAN. The advantage to LAN-free backups is almost always increased throughput, thus making it easier to complete backups on time.

As SANs have become an industry standard for consolidating storage resources, hundreds of application vendors now offer products that help support storage management and availability in a SAN. In addition, as SANs are assembled using industry standard parts and protocols such as FCP and iSCSI, an administrator can design a SAN using off-the-shelf parts.


NAS Filers

NAS filers have been at the heart of the file server consolidation boom. Organizations that face the challenge of scaling out file servers can simply purchase a single NAS filer with more than a terabyte of available storage.

One of the greatest selling points to NAS filers has been that they're plug-and-play in nature, allowing them to be deployed within minutes. At the same time, however, most NAS solutions are vendor-centric, meaning that they don't always easily integrate with other network management products. NAS vendors such as EMC and NetApp offer support for a common protocol known as Network Data Management Protocol (NDMP), which allows third-party backup products to back up data on EMC and NetApp appliances. The benefit of NDMP is that it is intended to be vendor neutral, meaning that if a backup product supports NDMP, it can back up any NDMP-enabled NAS appliance.

Microsoft NAS appliances that run the Windows Storage Server OS, however, do not support NDMP. Backing up a Windows Storage Server NAS will require the installation of backup agent software on the NAS itself.

Traditional NAS vendor offerings do not allow administrators to install backup software on their filers. The reason for this restriction is to guarantee the availability of the NAS; however, it significantly ties the hands of administrators when they're looking for flexibility.

 Later, this chapter will spend more time contrasting the role of NAS in server and storage consolidation with that of other products.

DFS

DFS has been seen as an easy way to combat server sprawl, at least from a user or application perspective. As servers are added to a network to accommodate growth and demand, the new servers can be referenced under a single domain-based DFS root. This feature allows the addition of the new servers to be transparent to users.

DFS can equally support server consolidation. For organizations that are consolidating and removing servers from the LAN, DFS can add a layer of transparency to the consolidation process. If user workstations and applications are set up to access file shares via the DFS root, administrators are free to move and relocate shares in the background and simply update the links that exist at the DFS root once the migration is complete. Thus, the way users access file systems will be the same both before and after the migration.

Virtualization

Virtualization technologies have recently jumped to the forefront of organizations' efforts to consolidate and simplify data access and management. This section will look at how storage virtualization has aided in storage consolidation efforts of SANs and how server virtualization has enabled companies to reduce the number of physical servers on their LANs by as much as 75 percent.

Storage Virtualization

As the number of managed storage resources on a LAN grows, so does the time and cost of managing those resources. Implementing a SAN provides an excellent first step toward consolidation and easing an administrator's storage management burden, however, the SAN alone may not be enough. This is where storage virtualization comes into the picture. There are plenty of ways to define *storage virtualization*, but to keep it simple, consider storage virtualization to be the logical abstraction of physical storage resources. In other words, a logical access layer is placed in front of physical storage resources.


 Storage virtualization is often a confusing topic due to the fact that several storage vendors have their own definition of the term. Competing vendors—most of which claim to have invented storage virtualization—may offer differing definitions. A common voice for storage virtualization can be found at the Storage Network Industry Association (<http://www.snia.org>).

Figure 2.1 provides a simple illustration of storage virtualization. The primary point of storage virtualization is to logically present physical storage resources to servers. This setup often results in better storage utilization and simplified management. As Figure 2.1 illustrates, storage virtualization starts by adding a data access layer between systems and their storage devices.




Figure 2.1: Virtualization access layer for physical storage resources.

The actual virtualization layer can be comprised of several different technologies. Among the virtualization technologies that may exist between servers and storage are:

- In-band virtualization
- Out-of-band virtualization
- Hierarchical Storage Management (HSM)
- Policy-based storage virtualization

The next four sections explore how each of these virtualization architectures aids in storage consolidation.

In-Band Virtualization

With storage virtualization, the term *in-band* implies that the virtualization device lies in the data path. The purpose of the device is to control the SAN storage resources seen by each server attached to the SAN. This level of virtualization goes far beyond traditional SAN segmentation practices such as zoning by allowing an administrator to allocate storage resources at the partition level, instead of at the volume level. Figure 2.2 shows an example of in-band virtualization.




Figure 2.2: In-band storage virtualization.

Notice that there is a *virtualization appliance* in the data path. The role of the appliance is to logically present storage resources to physical servers connected to the SAN. Also, as it resides directly in the data path, the appliance will provide more control over physical separation of SAN resources. For simplicity, the virtualization appliance is shown as an independent device, but such doesn't have to be the case. Although established NAS vendors such as Network Appliance and EMC are now offering virtualization appliances that fall in line with their general NAS philosophy, other fabric switch vendors such as Cisco Systems and Brocade are integrating virtualization intelligence into their fabric switches. Thus, the virtualization appliance does not have to be a standalone box and instead can seamlessly integrate into a SAN fabric.

☞ Oftentimes multi-function SAN devices appear to initially have a higher cost than their single function counterparts. However, every device introduced to the data path in a SAN can result in a single point of failure. This shortcoming is often overcome by adding redundant components. Thus, for fault tolerance, an organization will need two of each single function device. If the devices across the SAN can't be cohesively managed, separate management utilities will be required for each. Comparing this option with solutions such as running VERITAS Storage Foundation on top of a Cisco MDS 9000 switched fabric will reveal a significantly lower cost of ownership.

To get the most out of in-band virtualization, many organizations deploy a software storage virtualization controller such as IBM SAN Volume Controller or VERITAS Storage Foundation. With the virtualization component residing in fabric switches as opposed to living on standalone appliances, an organization will have fewer potential single points of failure in the SAN.

Many organizations have been wary about deploying in-band virtualization because of the overhead of the virtualization appliance. In being inside the data path and having to make logical decisions as data passes through the appliance, the virtualization appliance will induce at least marginal latency to the SAN. Vendors such as Cisco have worked to overcome this issue by adding a data cache to the appliance or switch. Although adding a cache to the appliance can improve latency, it will still likely be noticeable in performance-intensive deployments.

Out-of-Band Virtualization

Out-of-band storage virtualization differs from in-band virtualization in the location of the virtualization device or software controlling the virtualization. With out-of-band virtualization, the virtualization device resides outside of the data path (see Figure 2.3). Thus, two separate paths exist for the data path and control path. (With in-band virtualization, both data and control signals use the same data path.)




Figure 2.3: Out-of-band storage virtualization.

With control separated from the data path, out-of-band virtualization deployments don't share the same latency problems as in-band virtualization. Also, as out-of-band deployments don't reside directly in the data path, they can be deployed without major changes to the SAN topology.

Out-of-band solutions can be hardware or software based. For example, a DFS root server can provide out-of-band virtualization. DFS clients locate data on the DFS root server and are then redirected to the location of a DFS link. Data transfer occurs directly between the server hosting the DFS link and the client accessing the data. This setup causes out-of-band deployments to have less data path overhead than in-band virtualization.

Another advantage of out-of-band virtualization is that it's not vendor or storage centric. For example, IBM and Cisco sell a bundled in-band virtualization solution that requires specific hardware from Cisco and software from IBM. Although both vendors' solutions are effective, some administrators don't like feeling that an investment in technology will equal a marriage to a particular vendor. However, purchasing a fabric switch such as the Cisco MDS 9000 series that offers in-band virtualization as opposed to a dedicated in-band appliance will still offer some degree of flexibility. If an organization decides to move to an out-of-band solution later, the company will still be able to use the switch on the SAN.

There are several options available for pooling and sharing disk resources on a SAN. Although both in-band and out-of-band virtualization differ in their approach to storage virtualization, both options offer the ability to make the most out of a storage investment. Consolidating storage resources on a SAN is often the first step in moving toward a more scalable storage growth model. Adding virtualization to complement the shared SAN storage will provide greater control of shared resources and likely allow even more savings in terms of storage utilization and management with a SAN investment.

HSM

HSM is a management concept that has been around for several years and is finally starting to gain traction as a method for controlling storage growth. Think of HSM as an automated archival tool. As files exceed a predetermined last access age, they are moved to slower, less expensive media such as tape. When the HSM tool archives the file, it leaves behind a stub file, which is usually a few kilobytes in size and contains a pointer to the actual physical location of the file's data. The use of stub files is significant because it provides a layer of transparency to users and applications accessing the file. If a file has been migrated off of a file server, leaving a stub file allows users to access the file as they usually do without being aware of the file's new location. The only noticeable difference for the users will be in the time it takes for the file to be retrieved by the HSM application.


 Some HSM tools have moved away from the use of stub files and work at the directory level instead, thus allowing the contents of an entire folder to be archived. NuView's StorageX performs HSM with this approach. StorageX is able to leverage the existing features of DFS and NFS to archive folders, while adding a layer of transparency to the end user.

Figure 2.4 shows a simple HSM deployment. In this example, files are migrated from a disk array on the SAN to a tape library. The migration job would be facilitated by a file server attached to the SAN that is running an HSM application.




Figure 2.4: Migrating files older than 6 months to tape.

HSM tools typically set file migration criteria based on:

- Last access time
- Minimum size
- Type

When a migration job is run, files that meet the migration criteria are moved to a storage device such as a tape and a stub file is left in its place. The advantage of HSM for storage consolidation is that it allows the control of the growth of online storage resources. By incorporating near-line storage into a storage infrastructure, an organization can continue to meet the needs of online data demands while minimizing the demand of online storage devices as well as the size of backups.

To further understand HSM, consider the example of a law office. Many legal organizations maintain electronic copies of contracts; however, it may be months or even years before a contract document may need to be viewed. In circumstances such as this, HSM is ideal. HSM allows all contract documents to remain available while controlling the amount of online disk consumption and needed full backup space.


Policy-Based Storage Virtualization

Several backup software vendors, including VERITAS and CommVault, currently offer policy-based storage virtualization. This approach to virtualization has simplified how backup and restore operations are run. By using a logical container known as a *storage policy*, backup administrators no longer need to know the physical location of data on backup media. Instead, the physical location of data is managed by the policy. Consider a storage policy to be a logical placeholder that defines the following:

- Backup target device (library, disk array, and so on)
- Backup medium (tape, disk)
- Backup data retention days

When a server is defined to be backed up using enterprise backup software, an administrator does not need to select a backup target. This selection is automated through the use of the storage policy. Although backups are simplified through the use of storage policies, restores are where an organization will see the most benefit. When an administrator attempts to restore data, he or she does not need to know which tapes the necessary backups were stored on. Instead, the administrator simply selects the system and the files to restore. If any exported tapes are needed for the restore job, the administrator will be prompted. If the tapes are already available in the library, the restore will simply run and complete.

The advantage to this approach is that administrators don't need to scan through a series of reports to find a particular backup media in order to run a restore. Instead, they can simply select a server, file, and date from which to restore. This level of storage virtualization is especially useful in large and enterprise-class environments with terabytes to petabytes of managed storage. As storage continues to grow, it is becoming increasingly difficult to manage. Adding a storage policy management layer to data access alleviates many of these problems.

 For more information about storage virtualization, download a copy of SNIA's "Shared Storage Model: A Framework for Describing Storage Architectures" at http://www.snia.org/tech_activities/shared_storage_model. This document describes the SNIA standards for a layered storage architecture.

Storage virtualization is often seen as a first step in consolidating network resources. With storage resources centrally pooled and managed, the constraints of running tens to hundreds (or even thousands) of independent data islands on the LAN are eliminated. By consolidating to a SAN, a storage investment will now be closer to storage needs and organizations will have an easier time backing up servers within an allocated backup window. With storage firmly under control, the next logical step in consolidating network resources is server virtualization.

Server Virtualization

Server virtualization involves freeing servers on the network from their normal hardware dependencies. The result of server virtualization is often additional space in the server room. This benefit results from the fact that most likely several servers on the network are not using nearly all of their physical resources. For example, an organization might have one file server that averages 10 percent CPU utilization per day. Suppose that peak utilization hits 30 percent. Ultimately, the majority of the server's running CPU resources are doing nothing.

One way to solve the problem of under-utilized server hardware resources is to run multiple logical servers on one physical box. Today, there are two fundamental approaches to achieving this:

- Virtual Machines
- Shared data clusters

The next two sections show how each of these approaches allows for a reduction in the amount of physical resources on the LAN.

Virtual Machines

The use of virtual machines enables the running of multiple independent logical servers on a single physical server. With virtual machines, the needed hardware for the virtual machine to run is emulated by the virtual machine-hosting application. The use of virtual machines offers several advantages:

- Allows for the reduction in the number of physical systems on the network
- Provides hardware independence—a virtual machine can be migrated from one physical host to another without significant driver updates
- Enables an organization to run legacy application servers on newer, more reliable hardware

Simplified Recovery

When restoring a Windows server backup on a system containing different hardware, it can be difficult to get the restored backup to boot on the hardware. This difficulty is usually attributed to the characteristics of the Windows System State. When a System State backup is run on a system, the system's registry, device drivers, boot files, and all other system files are collectively backed up. When the System State is restored, all of these files come back as well. This behavior can cause problems if an administrator is restoring to different hardware, which is usually the case when a restore is needed to recover from complete system failure or a disaster. Once the restore operation writes all of the old registry, boot, and drive settings to the new system, odds are that the system will blue screen the first time it boots. Sometimes recovery from these problems lasts only a few minutes. However, for some administrators, this process may take several hours.

In virtualizing system hardware, restoring a virtual machine to another virtual machine running on a separate host system should be relatively problem free because the hardware seen by the OS on each system will be nearly identical. Thus, portability is another benefit provided by virtual machines in production environments.

Major Drawbacks

As a result of the advantages of virtual machines, some administrators rushed to fully virtualize their complete production environments, only to later return some virtualized servers back to physical systems. The reason was the major drawbacks to running virtual machines in production:

- Additional latency
- No reduction in the number of managed systems

Virtual machine host applications emulate system resources for their hosted OSs, so an additional access layer exists between virtual machines and the resources they use. Some of the latency encountered by virtual machines can be reduced by having them connect directly to physical disks instead of using virtual disk files. However, on CPU-intensive applications, the latency is still noticeable.

The other hidden drawback to consolidation through the use of virtual machines is that it does not reduce the number of managed systems on the network—instead it can increase that number. For example, if an administrator plans to consolidate 24 servers to virtual machines running on three hosts, there will be 27 servers to manage—the 24 original systems as well as each of the three virtual machine host servers. Therefore, although virtual machines enable fewer physical resources in the server room, there will still be the same number or even more servers that will require software, OS, and security updates.

Shared Data Clusters

Shared data clusters have emerged as a way to escape from the boundaries of server consolidation through virtual machines. Unlike with virtual machines, deploying shared data clusters can provide the ability to reduce both the number of physical systems and managed OSs on the network.

A major difference between shared data clusters and virtual machine applications is in their approach to consolidation. Shared data clusters are application centric, meaning that the clustered applications drive the access to resources. Each application, whether a supported database application or file server, can directly address physical resources. Also, in being application centric, the consolidation ratio doesn't need to be 1-to-1. Thus, 24 production file servers could perhaps run on four cluster nodes. As the virtualized servers exist as part of the clustering application, they are not true managed systems. Instead, there would be only four true servers to update and maintain. This option offers the benefit of physical server consolidation at a highly reduced cost of ownership.

Shared data clusters also offer the following advantages:

- Failover support—If a cluster node fails, an application can move to another node in the cluster, thus maintaining data availability after a system failure
- Shared data support—Shared data clustering provides the ability for data sharing between hosted file serving and database applications
- Load-balancing support—Client connections can be distributed among multiple cluster nodes; with traditional failover clustering, each virtual server entity is hosted by a single system
- Simplified backup—With shared data support, backups can be driven through a single cluster node, thus simplifying backups and restores as well as reducing the number of required licenses for backup software

Comparing Virtual Machines and Shared Data Clusters

To make sense of these two approaches, Table 2.1 lists the major differences between server consolidation via virtual machines and by deploying shared data clusters.

Administrative Task	Virtual Machines	Shared Data Clusters
Reduce the number of managed systems	Each virtual machine must still be independently updated and patched; the number of managed systems will likely increase as virtual machine hosts are added	The number of managed systems is significantly reduced, with the cluster nodes representing the total number of managed systems
Consolidate legacy file servers to a single physical system	Supported	Supported
Failover	Yes, via installed OS or third-party application	Yes
Virtualization software overhead	Up to 25 percent	None
Single point of failure	Potentially each virtual machine	No
Ability to share data	No	Yes
Backup and recovery	Each virtual machine must be backed up independently	Shared storage resources in the SAN can be backed up through a single node attached to the SAN
Support for legacy applications	Yes	No

Table 2.1: Virtual machine vs. shared data clusters.

As this table illustrates, server consolidation via shared data clusters offers several advantages over consolidation with virtual machines. With shared data clusters, you wind up with fewer managed systems, no additional CPU overhead, the ability to share data between applications, no single point of failure, and far fewer managed systems. However, application support is limited to what is offered by the shared data cluster vendor. Virtual machines can host nearly any x86 OS and thus are well-suited for consolidating legacy application servers, such as older NetWare, Windows NT, or even DOS servers that have had to remain in production in order to support a single application.

When looking to use virtual machines or shared data clusters to support server consolidation, an organization might not need to choose between one and the other. Instead, an organization should look to use each application where it's best suited—virtual machines for consolidating legacy application servers, and shared data clusters for supporting file and database server consolidation for mission-critical applications.



In consolidating to a shared data cluster, organizations not only see the benefit of fewer managed systems but also realize the benefits of high availability and improved performance.

Examining Unappliance vs. Appliance Solutions

At the heart of server consolidation are two fundamentally different points of view—appliance and unappliance. Each differs in its approach to both the hardware and software used in the consolidation effort:

- Appliance—Vendor-proprietary hardware and/or software solution that provides for server and/or storage consolidation
- Unappliance—Vendor-neutral x86-based hardware and software solution

These approaches to consolidation are significantly different, with both short-term and long-term consequences. This section will look at the specific differentiators between unappliance and appliance solutions:

- Proprietary vs. open solutions
- Volume economics
- Integration with existing infrastructure and investments
- Scalability
- Backup challenges

The section will begin with a look at the key differences between proprietary and open solutions.

Proprietary vs. Open Solutions

When retooling a network, there are convincing arguments for both proprietary and open solutions. Proprietary solutions are usually packaged from a single vendor or group of vendors and are often deployed with a set of well-defined guidelines or by a team of engineers that work for the company offering the solution.


A long-term benefit of a proprietary solution is that an organization ends up with a tested system that has predictable performance and results. A drawback to this approach, however, lies in cost. The bundled solution often costs more than comparable open solutions.

Another difference with proprietary solutions is that they may be managed by a proprietary OS. For example, many NAS filers run a proprietary OS, so management of the filer after it is deployed will require user training or at least a few calls to the Help desk.

Aside from the initial cost, buying a proprietary solution could cause additional higher costs in the long run. Upgrading a vendor-specific NAS, for example, will likely require the purchase of hardware through the same vendor. Software updates will also need to come through the vendor. Finally, if the network outgrows the proprietary solution, an organization might find that it must start all over again with either another proprietary solution or an open-standard solution.

The most obvious difference between open solutions and proprietary solutions is usually in cost. However, the differences extend much further. Open solutions today are based on industry-standard Intel platforms and enable the ability to run an x86 class OS, such as Windows, Linux, and NetWare. For hardware support, there are several vendors selling the same products, thus lowering the overall cost. Also, with open-standard hardware being able to support a variety of applications and OSs, as servers are replaced, they can be moved to other roles within the organization.

A drawback to open systems is often support. Proprietary solution vendors will frequently argue that they provide end-to-end support for their entire solution. In many cases, several vendors involved with a problem existing on a network using open architecture may point fingers when a problem occurs. For example, a storage application vendor may say a problem is the result of a defective SAN switch. The SAN vendor may go back and say that the problem is with a driver on the application server, or that the application is untested and thus not supported.

 Although finger pointing often occurs in the deployment and troubleshooting of hardware and software on open architecture, this is often the result of a lack of knowledge of most of the parties involved. Consider Ethernet as an example. Today, just about anyone will help assist problems with Ethernet network troubleshooting, and can do so because this open standard has been around for several years. The same can be said about SANs. As SANs have matured, the number of skilled IT professionals that understand SANs has grown too. This growth leads to more effective troubleshooting, less finger pointing, and often less fear of migrating to a SAN.

Thus, although open systems don't always offer the same peace of mind as proprietary solutions, their price is often enough to sway IT decision makers in their direction.

Volume Economics

In IT circles, proprietary is almost always equated to expense. This association is perhaps the simplest argument for going with a non-proprietary solution. With non-proprietary hardware, an organization can choose preferred servers and storage infrastructure. Also, use of industry-standard equipment allows for the free use any of the existing management applications on the consolidated server solution.

In March of 2004, PolyServe studied the price and performance differences between a proprietary NAS filer and a shared data cluster. In the study, the company found that going with a shared data cluster over NAS resulted in 83 percent savings. To arrive at the savings, PolyServe priced the hardware and software necessary to build a 2-node PolyServe shared data cluster. The cost of 12.6TB storage in a SAN and two industry-standard servers with two CPUs, 2GB RAM and running Windows Server 2003 (WS2K3) and PolyServe Matrix Server was \$79,242. In comparison, two NAS files with 12.6TB of storage with the CIFS file serving option and cluster failover option cost \$476,000. If you look at the solution in terms of cost per terabyte, the proprietary NAS appliance-based solution cost \$38,000 per terabyte. The unappliance-based solution cost \$6300 per terabyte.

Integration with Existing Infrastructure and Investments

Integration is another key differentiator between appliance and unappliance philosophies. Proprietary appliances are often limited to management tools provided by the appliance vendor. Installation of management software on an appliance is often taboo. Many NAS appliances can connect to and integrate with SANs to some degree, but the level of interoperability is not that of an open unappliance-based system.

The Scalability Dilemma

With an initial investment in a NAS appliance, an organization's needs will likely be satisfied for the next 12 to 18 months. Many proprietary NAS appliances offer some scalability in terms of storage growth by allowing for the attachment of additional external SCSI arrays or connectivity to fibre channel storage via a SAN.

Scaling to meet performance demand is much more difficult for proprietary NAS. Unlike shared data clusters on industry-standard architecture, proprietary NAS solutions may offer failover but do not offer load balancing of file serving data. As two or more NAS appliances cannot simultaneously share the same data in a SAN, they cannot offer true load balancing for access to a common data store. Instead, as client demand grows, the NAS head often becomes a bottleneck for data access. To address scalability, organizations often have to deploy multiple NAS heads and divide data equally among them. Tools such as DFS can allow the addition of the NAS heads to be transparent to the end users.

Unappliance-based shared data clusters do not run into the same scalability issues as proprietary NAS appliances. With open hardware and architecture, additional nodes can be added to the cluster and attached to the SAN as client load increases. Furthermore, with the ability to load balance client access across multiple nodes simultaneously with a common data store, shared data clusters can seamlessly scale to meet client demand as well.

Backup Challenges

Many proprietary NAS appliances have not been able to work well with current open backup practices. Instead, each vendor typically offers its own method of advanced backup functionality. For example, both EMC and NetApp provide proprietary snapshot solutions for performing block-level data backups. To perform a traditional backup, backup products that support NDMP can issue NDMP backup commands to the NAS appliance. Figure 2.5 illustrates an example of an NDMP-enabled backup.




Figure 2.5: NDMP-enabled backup.

As additional NAS appliances are added, they too could be independently backed up via NDMP. To keep up with newer industry standard backup methods such as server-free and server-less backups (discussed in Chapter 6), NAS appliance vendors have worked to develop their own methods of backing up their appliances without the need of CPU cycles on the NAS head.

Not all NAS appliances can directly manage robot arms of a tape library, so often a media server must be involved in the backup process to load a tape for the NAS appliance. Once the tape is loaded, the NAS can then back up its data.

Figure 2.6 shows how backups can be configured on a shared data cluster. Notice in this scenario that one of the cluster nodes is handling the role of the backup server. With two other nodes in the cluster actively serving client requests, the dedicated failover node is free to run backup and restore jobs behind the scenes.

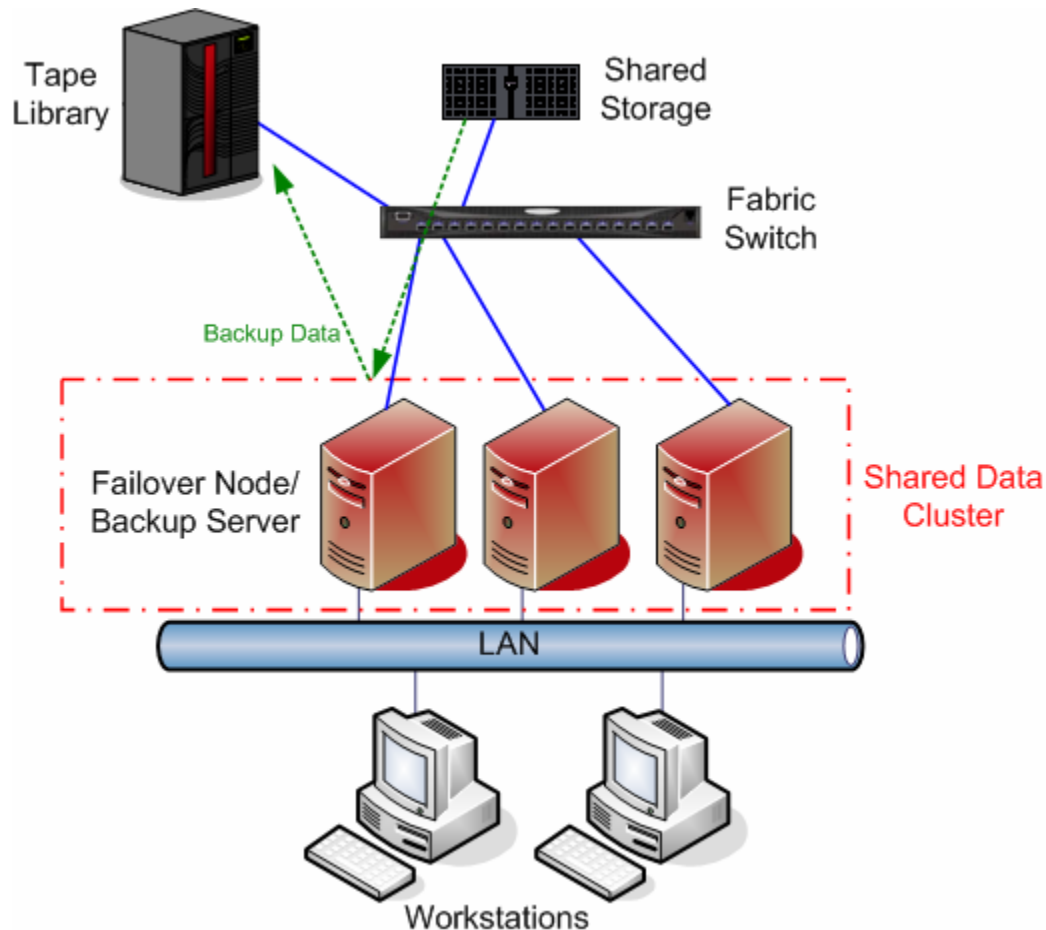



Figure 2.6: Unappliance-based shared data cluster backup.

The key to making this all work lies in the fact that all nodes in the shared data cluster can access the shared storage simultaneously. This functionality allows the passive node to access shared data for the purpose of backup and restore. Also, as the passive node is doing the backup “work,” the two active nodes do not incur any CPU overhead while the backup is running.

 Notice that the backup data path appears to reach the failover node before heading to the library. The behavior of the backup data will be ultimately determined by the features of the backup software and SAN hardware. For example, if the backup software and SAN hardware support SCSI-3 X-Copy, backup data will be able to go directly from the storage array to the tape library without having to touch a server.

Another advantage to backing up data in an unappliance shared data cluster is the cost of licenses. As all nodes in the cluster could potentially back up or restore the shared data, only one node needs to have backup software installed on it, and consequently a backup license. As a cluster continues to scale, an even higher cost savings for backup licensing will become apparent.

Taming Server and Storage Growth—the Non-Proprietary Approach

With most of the industry leaning in the direction of managing growth and scalability through non-proprietary hardware, the remainder of this chapter will focus on building a scalable file serving infrastructure on non-proprietary hardware. Taming server and storage growth requires consolidation with an eye on scalability and management of both file servers and storage resources.


Storage Consolidation via SAN

For several years, SANs have been a logical choice to support storage consolidation. As SANs share data using industry-standard protocols such as FCP and iSCSI, there are products from several hardware vendors available to build a SAN. Because the products adhere to industry standards, they allow for the flexibility to mix and match products from different vendors.

One of the primary goals of consolidating via a SAN is to remove the numerous independent data islands on the network. The SAN can potentially give any server with access to the SAN the ability to reach shared resources on the SAN. This feature can provide additional flexibility for both data management and backups. In addition to sharing disks, an administrator can share storage resources such as tape libraries, making it easier to back up data within the constraints of a backup window. With SAN-based LAN-free backups, backup data does not have to traverse a LAN in order for it to reach its backup target. To ensure that an administrator is able to perform LAN-free backups on the SAN, the administrator will need to ensure that the backup vendor supports the planned SAN deployment.

The use of SAN routers enables the continued use of SCSI-based storage resources on the SAN. This may mean that to get started with a SAN, an organization will need to purchase only the following:

- 1 to 2 HBAs for each server—Two HBAs are required for multipath support and allows for fault tolerance in the complete SAN data path
- 1 to 2 fabric switches—Two switches are required in redundant fabrics)
- 1 router—The quantity of routers will vary depending on the number of SCSI resources being moved to the SAN

 The number of routers required for the initial SAN deployment may be increased to address concerns such as the creation of a bottleneck at the fibre channel port of the router. Chapter 3 covers these issues in detail.

With several servers potentially accessing the same data on the SAN, care will need to be taken to ensure that corruption does not occur. The choice of virtualization engine or configuration of zoning or LUN masking can protect shared resources from corruption.

Server Consolidation via Clustering

Server consolidation via clustering offers a true benefit of reducing the total number of managed systems on the network. In addition, consolidating to a shared data cluster provides greater flexibility with backups, the ability to load balance client requests, and failover support.

Shared data clusters offer the benefit of scaling to meet both changes in client load and storage. Because they run on industry standard hardware, scaling is an inexpensive option. Furthermore, consolidating to shared data clusters offers significant savings in cost of ownership and software licensing. Consolidating a large number of servers to a shared data cluster results in the need to maintain fewer OS, backup, and antivirus licenses. With fewer managed systems on the network, administrators will have fewer hardware and software resources to maintain on a daily basis. With substantial cost savings over proprietary NAS appliance solutions, shared data clustering has emerged as an easy fit in many organizations.

Planning for Growth While Maintaining Freedom

When planning for growth while consolidating, there are several best practices to keep in mind. When planning to consolidate the network, consider the following guidelines:

- Design for scalability
- Design for availability
- Use mature products
- Avoid proprietary hardware solutions

One of the most over-used words in the IT vocabulary is *scalability*. However, it is often one of the most important. Scalability means that all elements of the network infrastructure restructure should support company growth, both expected and unexpected. For example, if current requirements warrant the purchase of an 8-port fibre channel switch, consider purchasing a 16-port switch. Ensure that the planned SAN switches offer expansion ports so that the option is available to scale out further in the event that growth surpasses projections.

Scalability can also be greatly aided by the use of centralized management software. Many of the solutions previously mentioned in this chapter can help with centrally managing storage resources and backups across the enterprise. If performance is an issue, another major scalability concern should be in technologies that support load balancing of client access. Such technologies alleviate single server bottlenecks that result from unexpected client growth.

In designing for availability, redundancy is always crucial. Redundancy often starts with shared storage on the SAN via a RAID implementation. However, adding redundancy to the physical disks is most valuable when the disks are behind a fully redundant data path. Ensuring true redundancy means not only protecting disks but also

- Using cluster architecture to prevent data access loss due to a system failure
- Using redundant switches in the SAN fabric and redundant HBAs
- Using multipath-compliant HBA drivers on servers to ensure that they can realize the benefits of redundancy in the SAN
- Planning for redundant power sources to prevent data loss or corruption from a power failure
- Adding redundancy to the LAN to prevent switch failure from interrupting data access

Using mature products refers to products that have established reputations in IT. Regardless of how great the sales pitch, bringing an unknown product into the network is always a risk. Products with the backing of OS vendors such as Microsoft are more likely to do what they promise, and their vendors are more likely to be in business at both the beginning and end of the project.

With server and storage consolidation, it's often tempting to look to purchase a proprietary solution for a single hardware vendor. With a SAN, there will be little problem integrating solutions from vendors such as Brocade, McData, and Cisco Systems. However, several vendors offer end-to-end solutions. If the proposed solution involves proprietary hardware, it's unlikely that the solution will work as well with the rest of the systems on the network. For example, managing and reporting can't be done by the tools currently be used and instead will need to be achieved through an add-on utility from the product vendor or by using custom scripts. Most vendors have no problem offering professional services and can write a script to handle most tasks. However, each time support is required for the script, an organization might end up spending additional dollars on more professional services.

As an organization begins to piece together a planned network, it must pay close attention to the product support matrixes listed on each vendor Web site. Doing so ensures that the proposed pieces have been tested and will work well together. For planned products not on a vendor's support matrix, a good practice is to negotiate a pilot period to ensure the product works with new hardware and software.

Summary

Chapters 1 painted a picture of the current file serving landscape. This chapter looked deeper into the picture and examined the available server and storage consolidation alternatives. Building on this framework for the technologies that may be involved in consolidation, the next chapter will look at how to piece these technologies together while emphasizing how to maintain high availability and high performance.

Content Central

[Content Central](#) is your complete source for IT learning. Whether you need the most current information for managing your Windows enterprise, implementing security measures on your network, learning about new development tools for Windows and Linux, or deploying new enterprise software solutions, [Content Central](#) offers the latest instruction on the topics that are most important to the IT professional. Browse our extensive collection of eBooks and video guides and start building your own personal IT library today!

Download Additional eBooks!

If you found this eBook to be informative, then please visit Content Central and download other eBooks on this topic. If you are not already a registered user of Content Central, please take a moment to register in order to gain free access to other great IT eBooks and video guides. Please visit: <http://www.realtimepublishers.com/contentcentral/>.