



realtimepublishers.comtm

The Definitive Guidetm To

Building Highly Scalable Enterprise File Serving Solutions



Chris Wolf

Chapter 6: Managing High-Performance, Scalable, and Resilient Data Across the Enterprise ..	110
Challenges Facing Heterogeneous Networks	110
Inhibited Agility	111
Complexity	111
Integration Concerns	111
IT Risk and Compliance Considerations	111
Integrating Windows and Linux File-Serving Solutions	112
CIFS and NFS Integration	112
Managing ACLs	113
Integration with Existing Services	114
Backup and Recovery	114
Disaster Planning Essentials	115
Development	115
Disaster Planning Roles	116
Traditional Backup Methodologies	117
Snapshots	117
Server-Free Backups	117
Server-Less Backups	118
Archiving and Migration	120
Successful Backup Architectures	121
D2T	121
D2D	122
D2D2T	122
Benefits of Share-Data Approaches	124
Comparison: Consolidated vs. Distributed Backup Architectures	125
Data Recovery	127
The Advantages of Freedom	128
Benefits of Avoiding Proprietary Solutions	128
Uncapped Scalability and Performance	129
Architecture Flexibility	129
Freedom of Choice	130
Summary	130

Copyright Statement

© 2006 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 Realtime Nexus—The Digital Library. All leading technology guides from Realtimepublishers can be found at <http://nexus.realtimepublishers.com>.]

Chapter 6: Managing High-Performance, Scalable, and Resilient Data Across the Enterprise

The two previous chapters took an in-depth look at both Windows and Linux file-serving solutions and how these two operating systems (OSs) present individual challenges and advantages in the areas of performance, scalability, availability, and integration. As vendors have worked to manage and mitigate these challenges, varying new technologies have been developed to meet the need to manage data. As each new incarnation is adopted, heterogeneous networks have developed over time, presenting challenges in the areas of integration, backup and recovery, and freedom to manage your storage solution the way you see fit.

This chapter will examine the broader enterprise picture and leverage what previous chapters have discussed to develop a clearer understanding of the high-level responsibilities (and granular realities) of managing data across the enterprise. This chapter will assume a holistic vantage point to examine the challenges faced in the enterprise today and touch on key points to consider when defining the strategy behind building highly scalable enterprise file-serving solutions. Let's begin by examining a few of the challenges faced in heterogeneous networks and how those challenges come into play in the enterprise and progress into the areas of backup and recovery.

Challenges Facing Heterogeneous Networks

For reasons apparent to IT managers and administrators, it is appropriate to refer to large-scale implementations of networked computer systems as an enterprise—a word that not only means “an undertaking” but more specifically an undertaking of great scope, complication, and risk. The scope of an enterprise may vary, but generally, all enterprise environments are complicated landscapes comprised of one or more heterogeneous networks that come together to be defined as one enterprise Wide Area Network (WAN). As IT managers work to align solutions with their individual IT missions and goals, meet compliance requirements, increase their return on investment, and drive cost efficiency, decisions are made that very rarely permit each solution to align with a single OS or product offering—and what is decided upon gets added to the over-encompassing umbrella referred to as the enterprise. Heterogeneous networks comprised of many OSs and protocols present inherent challenges such as:

- Inhibited agility
- Complexity
- Integration concerns
- IT risk and compliance considerations

This section will focus on the numerous management concerns presented when managing a heterogeneous network environment.

Inhibited Agility

Although you will often find strength in diversity, diversity also has a price. Protocols that are not inherently compatible with one another complicate the environment, creating additional management overhead and inhibiting the capability of the enterprise to remain agile. Agility within the enterprise is the enterprise's ability to quickly reconfigure IT resources to meet changing business demands. An enterprise should strive to be flexible in its ability to quickly respond to changing business requirements to meet the growing needs of the enterprise.

Remaining agile in an environment populated by years of accumulated, different storage solutions is challenging, to say the least. Many data centers are faced with the realization that although in the business world “more” often equates to “better,” the same is not the case in the data center. Striving to meet shorter times to market for emerging business initiatives is often a challenge that is hindered by the capacity of both the physical and virtual enterprise IT infrastructure components required to meet the growing demands of the organization.

Complexity

Relying on multiple OSs, hardware platforms, and software platforms inhibits the ability of administrators to centrally manage the entire environment. As the complexity of the enterprise environment increases, so do the costs of managing the environment. In addition to the core competencies required of the engineering, architecture, and support staff, these costs include implementing and maintaining the systems and technologies required to support, maintain, and—when disaster strikes—recover the environment. Whenever possible, steps should be taken to simplify the enterprise architecture to minimize these costs and maximize the effectiveness of ongoing support efforts.

Integration Concerns

Enterprise IT managers are continually driven to seek harmony in their environments. Compatible systems reduce the total cost of ownership of the environment by allowing the enterprise to standardize and simplify. As new systems, protocols, and applications are developed by competing vendors, often with little internal desire to remain compatible with the competition, the challenge of integrating enterprise systems escalates and becomes burdensome. Focusing on integration will help you reduce cost by reducing complexity and simplify management efforts by reducing or eliminating incompatible systems.

IT Risk and Compliance Considerations

Each OS, firmware revision, and supporting software application presents unique security and compliance challenges. These range from the broad consideration of service packs and hotfixes to granular OS and application configuration. The amount of time associated with managing and mitigating risks and maintaining compliance across the enterprise will inherently impact server, system, or network availability. As the enterprise strives to remain secure and compliant, the impact to availability is often felt. These risks need to be analyzed for their potential to impact business processes and goals.

Integrating Windows and Linux File-Serving Solutions

Windows and Linux file-serving solutions represent different, yet often complementary, approaches to file serving. At their very core, these two OSs are dramatically different and distant in design; as such, the solutions developed for each OS have been accordingly disparate. Central to integrating these two platforms are the challenges presented by:

- Common Internet File System (CIFS) and Network File System (NFS) integration
- Managing Access Control Lists (ACLs)
- Integration with existing services

This section will focus on each of these challenges and how you can leverage the information covered in the previous chapters to effectively integrate a cross-platform file-serving solution.

CIFS and NFS Integration

Chapter 4 presented Microsoft's Shadow Copy and its reliance on CIFS as the entry point for users to recover files. Chapter 5 presented NFS, which has risen to become the unmatched standard file-sharing protocol on UNIX and Linux servers. In addition to the out-of-the-box packaged distributions of UNIX and Linux, many NAS devices, of which the vast majority are based on Linux, use NFS as well. Enterprise IT managers that are looking to maximize their return on investment by utilizing the best Windows- and Linux-based technologies are now faced with a bit of dilemma—incompatibility.

The gap that exists between these two protocols can be bridged, but doing so presents performance concerns. Software solutions exist that allow UNIX- and Linux-based servers to provide remote file access functionality to PCs without requiring NFS. The most widely used of these are SAMBA, Hummingbird NFS Maestro, and Microsoft Windows Services for UNIX (SFU). SAMBA is a server-side installation, while NFS Maestro and SFU are NFS redirectors that are installed on client workstations.

When considering the use of a redirector, it is important to keep in mind the scalability of the solution and the architecture's dependency upon it. Installing NFS Maestro on a few dozen workstations that require access to a UNIX system for a special purpose, such as a small accounting department with a need to access an NFS share to run a particular report, may be acceptable in the short term. Success, however, often equates to scale, so pay heed to the overall strategy. What may be communicated to you today as a special-purpose small-scale implementation could very well develop into a situation that is much larger in scale than anyone initially foresaw. Both Windows and Linux have support for the other's file system—Linux, by the use of SAMBA for SMB services and Windows through SFU. On a Linux server, one might encounter the native NFS stack running in combination with SAMBA to serve files to both systems; on Windows, the native CIFS stack can be used with SFU to provide the same services.

Installing and managing a centralized server to bridge the gap is often a better option not only to reduce complexity but also to help prevent application sprawl by preventing the need for redirector software to be installed on multiple systems across the enterprise.

The use of Linux as a mainstream file-serving solution is becoming a reality, and many IT managers are no longer being afforded the luxury of sticking with a single protocol for all their file serving needs. This reality adds to the complexity of the environment, as NFS and CIFS protocols may both be required. NFS vs. CIFS is not an all-or-nothing equation; both can exist simultaneously to further enhance the capabilities of an existing infrastructure, and as technologies are continually developed to further enable enterprise management, the inclusion and integration of both CIFS and NFS file-sharing protocols can dramatically simplify the inclusion of future systems into the multi-protocol file-sharing environment.


Managing ACLs

Security management is becoming an increasingly complex task. Adapting to industry and regulatory compliance needs puts the enterprise IT manager in a virtually continuous reactive state. As new compliance directives are developed and subsequently adopted by your organization, information security policies, standards, guidelines, and procedures are implemented to meet these needs. Legislative compliance standards such as the Health Insurance Portability and Accountability Act (HIPAA) of 1996 provide clear directives on what needs to be protected from disclosure as well as severe penalties for organizations who fail to meet those directives. The only piece this legislation doesn't supply is the means to implement the changes to meet the requirements—that decision is left to you.

Often the most challenging and certainly the most granular task an IT administrator may face is managing ACLs. This task can be further compounded when file-serving resources are spread across multiple platforms with dissimilar ACL architectures. Different ACLs supporting separate application systems or users make it more difficult to meet legislative compliance requirements because they lack the means to be centrally managed effectively.

Because the ACL security model in NTFS is more robust and fundamentally different than the file-security model used in Linux, no one-to-one mapping can be made between them. The fundamental problem occurs when a Windows client (which expects an NTFS ACL) accesses a Linux file, or a Linux client (which expects Linux file permissions) accesses a Windows file. In these cases, the file server must sometimes authorize the request using a user identity that has been mapped from one system to the other—or, in some cases, even a set of permissions that has been synthesized for one system based on the actual permissions for the file in the other system. This setup creates its own set of security concerns that also need to be managed.

There is some good news in that NFS version 4 standardizes the use and interpretation of ACLs across Posix and Windows environments. This standardization will make centralized management of ACLs between these two systems easier. It will also support named attributes. User and group information is stored in the form of strings, not as numeric values. ACLs, user names, group names, and named attributes are stored with UTF-8 encoding.

 For more information on NFSv4 refer to the NFSv4 home page at <http://www.nfsv4.org/>.

Integration with Existing Services

Transparency is an important concept of enterprise architecture that aides in security efforts by keeping authorized users relatively unaware of the inner-workings of the infrastructure and makes systems and resources easier to locate and use. In few places is transparency more vital to productivity than in the file-serving arena. The viability of a file-serving solution within an enterprise environment is, to a great extent, dependant upon that solution's capability to remain transparent to the end user. File services should appear to users to be a conventional, centralized file system. The number and location of servers and storage devices should be made invisible.

Backup and Recovery

The previous chapters spent time discussing the use of file servers as a centralized repository to simplify the process of backing up and recovering systems. It is important when designing and implementing file-serving solutions that due care be given to ensure that these solutions are also backed up. The following list highlights a few simple rules to remember when designing a backup strategy:

- A backup should be easy to do.
- A backup should be automated and rely on as little human interaction as possible.
- Backups should be made regularly.
- There should be at least two copies of the data stored on resilient media and kept at different locations.
- A backup should rely on standard, well-established formats.
- A backup should not use compression. Uncompressed data is easier to recover if the backup media is damaged or corrupted.
- A backup should be able to run without interrupting normal work.

A backup is simply one process in an overarching area of responsibility for an organization. Whether the focus is on business continuity planning (BCP), disaster recovery, or compliance, backup and recovery planning is going to be crucial to the success of an organization to recover. To effectively plan and implement a backup and recovery plan in support of business continuity or disaster recovery, several processes need to be examined for inclusion:

- Risk analysis—Risk is a subjective term and although storage administrators tend to treat all data as if it were mission critical, it is important to understand how much risk a line of business is willing to accept in terms of data loss.
- Scheduling—It is a generally accepted rule that backups should not affect production. Work with your business partners to understand their business requirements and minimize the impact of backup operations on their environment.

- Review of logs—Backups generate logs that need to be reviewed regularly in order to ensure that they're performing properly.
- Testing—Backups are often taken for granted. Until, of course, you need to recover the system and discover that a backup has malfunctioned during a real disaster recovery operation. Test backups often to ensure their quality and viability for recovery.
- Retention—All data is not created equal. Some can be discarded virtually immediately; other data must be archived and maintained for years to meet legal, regulatory, or compliance obligations. The amount of time data must be maintained is an important metric to determining the proper archiving solution.

Disaster Planning Essentials

An area that is all too often overlooked or given too little emphasis is that of disaster recovery planning. Systems that comprise an enterprise file-serving environment should be protected, ideally in a secure data center with sufficient resources to operate autonomously should a natural disaster or other emergency interrupt key services and utilities such as telecommunications and power.

For disaster planning to be effective, it must be put into the proper context. Disaster planning is often presented in a manner that is—for all intents and purposes—inaccurate. Enterprise managers often approach disaster recovery planning with a perspective that doesn't clearly define the benefits. Managers also often fail to communicate the need for disaster recovery planning effectively to the line of business operations that will depend upon it. The planning itself doesn't relate directly back to an immediate or pressing need, so it is therefore set aside.

When presented in context, the need for disaster recovery planning is clear. Although you hope to never need to use this parachute, there is peace of mind in knowing it is there to protect you. The upfront cost and effort required for disaster recovery planning may be uncomfortable, but when you clearly illustrate the risk and real potential for disaster, the response from the business should be one of appreciation rather than one of remorse. Once you understand why disaster recovery planning is essential and how to best communicate it, you can move forward to examine the stages of disaster recovery planning and how you can begin.

Development

The first stage of disaster recovery planning requires the development of a plan to document the procedures for responding to an emergency, providing extended backup operations during the interruption, and managing recovery and reclamation of data and processes afterwards (should an organization experience a loss of data access or processing capability). A disaster recovery plan is an enterprise document that should outline the roles and processes of senior management as well as IT management and other critical personnel in key areas such as security, facilities engineering, and finance.

Disaster Planning Roles

From the vantage point of the IT manager, there will be several key roles to be fulfilled that should be documented in the disaster recover plan:

- Identifying and prioritizing mission-critical applications
- Recovering and reconstructing all critical data, systems, and supporting infrastructure
- Continuously reassessing the recovery site's stability



Identifying and prioritizing mission-critical applications is only one step that requires a close and in-depth understanding of business needs. Throughout, this chapter will refer to the importance of maintaining open lines of communication with storage users to better understand their business and subsequently their current and future storage requirements.

Identify and Prioritize

Identifying and prioritizing mission-critical applications is a step that should be taken as part of disaster recovery planning and then periodically reevaluated to align with changing business needs. It is important that clear lines of communication be established early on between the various lines of business supported by the enterprise infrastructure and IT management so that IT managers can make informed decisions to protect their line of business partners. Reevaluating this step periodically is important, not only to keeping those communication lines open but also in providing visibility of the disaster recovery plan.

Recovery and Reconstruction

The speed, efficiency, and effectiveness of the recovery and reconstruction efforts should be the focus of a sound disaster recovery plan. Many times, organizations spend a great deal of effort and planning to ensure that the backup of systems themselves does not impact production and pay little attention to the recovery time associated with the process.

Reassess

When operating in a disaster from a recovery site, an organization is in its most vulnerable state of operations. Few organizations outside of government the financial services industry maintain multiple recovery sites that can be utilized during a disaster. IT managers must act vigilantly to monitor and protect the stability of the recovery site.

As part of an organization's disaster recovery planning efforts, a recovery team may be defined with the mandate to implement the recovery procedures once a disaster is declared. The recovery team's primary duty is to get critical business functions operating at the alternative or backup site.

Traditional Backup Methodologies

Backup methodologies vary with the size, scope, and criticality of the data they are intended to protect. In the most basic traditional model, critical data from a key system is copied to storage media such as a tape, or file server, separate from the client to provide protection in the event of client failure. This process depends heavily on the system resources of the client to perform either the tape backup or copy operation, which naturally affects the performance of the client (by way of CPU utilization, network utilization, and so on).

Tape storage can become quite costly when adopted as a standard methodology as the use of individual tape drives sprawls throughout the enterprise. In addition to the cost of personnel's time to handle the physical tapes and maintenance cost associated with the tape drive hardware and replacement of tapes as they become unviable, there is the cost of storage and shipping to consider when transporting the media off-site for safekeeping.

In the past, it has been generally accepted that a backup should be performed at least once every 24 hours. This number is arbitrary and should be reconciled against the actual business continuity and disaster recovery requirements of the system you intend to recover.

Snapshots

Although some backup solutions simply copy data directly to a tape or another disk, some solutions use another process that utilizes snapshots. A snapshot is a relative (or delta) copy of a data set. It is differentiated from a mirror in that there are links between the original (or source) and the copy (or mirror).

In the snapshot process, the backup software makes a copy of the pointers to the data, which indicate its location, then relies on data movers to pick up the pointers and transfer the data. Snapshot volumes are point-in-time copies of primary storage volumes. By creating snapshot volumes, the primary volumes continue to be available for production operations, while the snapshot volumes are used for offline operations such as backup, reporting, and testing. This setup results in improved backup operations, data reporting, application testing, and many other day-to-day operations.

Server-Free Backups

The biggest advantage to server-free backup is the reduction of workload on the target server. A server-free backup, as the name implies, frees target servers' CPU, memory, and I/O consumption during the backup process by decreasing the servers' involvement in the backup process. Essentially, the data being backed up will move from the target server's disk to a data mover (see Figure 6.1). In a server-free architecture, the data mover is another server that is dedicated to providing the actual transportation of the data. A data mover can also be a device, as we'll see in the next section, such as a Small Computer System Interface (SCSI) drive or router that reads the data from a network drive and writes it to the backup device. The data mover also manages the flow of data between the network drive and backup device to ensure that no information is lost.

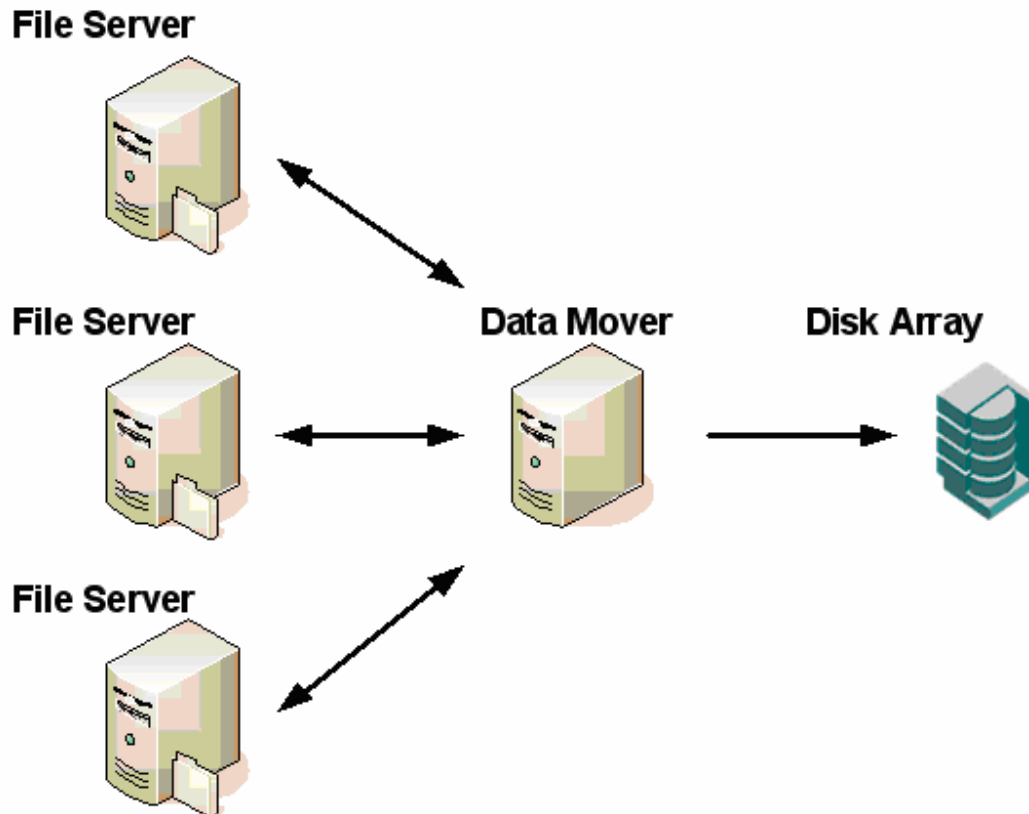


Figure 6.1: In a server-free design, a dedicated server acts as the data mover to free system resources on the target servers.

Server-Less Backups

Like server-free backups, server-less backups offer the advantages of efficiency, scalability, fault tolerance, and cost reduction, but are defined by a complete lack of dependence on a dedicated server to fulfill the role of a data mover. In a server-less environment, either the storage device or its supporting infrastructure are used to fulfill the role of the data mover. Technologies such as the SCSI-3 Extended Copy (XCOPY) command can be used to read and write data directly between a disk array and a secondary device and can take advantage of existing modules in backup applications to coordinate the backup process. This method of backup reduces total cost of ownership and operational costs by eliminating any need for additional servers and increases backup performance by eliminating the intermediary server from the backup process (see Figure 6.2).




Figure 6.2: A server-less architecture relies upon the storage solution, or the components of its supporting infrastructure, to transfer the data.

Some of the major advantages of using server-less backup include:

- Increased server efficiency
- Increased scalability
- Better fault tolerance
- Lower overall hardware costs

Utilizing a server-less architecture provides savings in the form of server elimination as there is no need for a dedicated server to perform the role of a data mover. In addition, this setup saves on network utilization, permitting the data to be transferred once directly from the server to the storage device and thus eliminating the additional step required with an additional server.

Archiving and Migration

There are times when data needs to be held onto for the long haul. Aside from traditional backup requirements, many organizations find that certain data may be required for legal, or compliance, reasons to be maintained for months or even years. Archiving refers to the processes supporting these needs and the storage requirements necessary to meet legal obligations.

When designing a file-serving storage solution, architects need to understand the archiving requirements and how they pertain to the data being stored so that the solution can be designed to meet those requirements. A tiered-storage approach provides insight into the value of data over time by classifying data early and progressively re-classifying the data as its value changes based upon differing motivators.

Take for example financial data whose value is immensely important during the time of and immediately following a transaction. Loss of a large-scale financial database, such as those used by credit card processors, during peak transaction hours could be catastrophic to the business. Once the transactions have cleared and have been reconciled, the data still remains important and the processor may need the data surrounding the transaction immediately on-hand for the next 30 days to facilitate refunds or for other internal business processes. As time passes, the data becomes less critical but still important as information about the transaction is used in tax reporting and, depending upon the industry and purpose of the transaction, regulatory compliance may require the data be kept for several years.

Understanding the need and scope of the long-term storage requirement will help drive storage decisions and the underlying financial motivations. As an understanding of the business need for the data matures, this understanding drives transformation. Enterprise storage architects, armed with an understanding of business needs, are compelled to consider storage architectures that meet those needs especially in the areas of performance, scalability, and resiliency—which leads many to consider migration to a consolidated storage architecture.

Unfortunately, storage consolidation isn't as simple as buying a large, enterprise-class storage array and migrating applications to the new platform. Migration takes time for planning if it's going to be successful. Essentially, there are four key areas that need to be assessed when planning for storage migration:

- Assess the current environment—During this phase, you'll need to identify your current storage capacity and gather metrics surrounding its utilization. You'll also need to re-establish your understanding of the business being supported by the storage environment and strive to understand the future requirements the business may be soon facing.
- Understanding the current costs—This phase is important in gathering cost justification for a storage consolidation effort. Work to gather the current storage hardware and software costs and understand how those costs impact the bottom line. Remember to focus attention on the cost of supporting the environment in time, personnel, support contracts, and licensing agreements.

- Assess storage management capabilities—It is important have a clear picture of the current management so that you can make informed decisions in the same context. In addition to administration, one must also consider the ability to monitor the environment for performance, availability, and security.
- Understand the future business and legal requirements—Although you might not have a crystal ball readily available to predict the future requirements that will be placed on your storage infrastructure, you can leverage the experience of your business colleagues whom often can provide a wealth of information about what the future may hold. New compliance issues and regulations that your partners may be working to meet may have significant impact on storage.

Once these steps have been completed, a clearer picture will have developed that will serve to guide you through the migration process. To aide in understanding a few of the architectures available, let's first examine what has worked in the past and what is now being adopted for use in enterprise backup architecture.

Successful Backup Architectures

There are many approaches to backup and recovery to be considered when developing an enterprise file-serving solution. Each will be dependant on the size, scope, business continuity, and disaster recovery needs of the file-serving solution you intend to develop. The three architectures most commonly found within an enterprise environment are

- Disk-to-Tape (D2T)
- Disk-to-Disk (D2D)
- Disk-to-Disk-to-Tape (D2D2T)


D2T

D2T has been, and to a great extent still is, the most common method used to create backups. D2T as a standalone backup solution still has many inroads within the enterprise, but its use as a standard for backup and recovery—as a standalone solution enterprise standard—is diminishing.

New storage solutions and architectures such as NAS and SAN, as discussed in Chapter 2, have opened avenues to consolidation. For small scale or systems isolated from the storage infrastructure, D2T is viable as an independent solution, and the relative size-to-cost ratio of tape storage makes it one of the most affordable solutions from a media standpoint. Within an enterprise file-storage solution, D2T is commonly found as a component of the more robust and resilient D2D2T architecture.

D2D

In D2D, as it is often represented, a computer hard disk is backed up to another hard disk rather than to removable media. D2D as a backup methodology enables both greater performance and higher capacity relative to tape or other removable media alternatives, which directly translates to shorter time to recovery. D2D can be used both to refer to a dedicated backup architecture in which one “disk” is used as a dedicated backup media and whose sole function is to serve as a backup device. D2D can also refer to contingency backup solutions in which one system is routinely backed up to a second identical system for recovery purposes.

 D2D is often confused with Virtual Tape Library (VTL) technology. A VTL is a data storage technology that employs the use of emulation that causes hard disks to behave virtually as if they were tape drives. However, D2D differs in that it enables multiple backup and recovery operations to simultaneously access the disk directly by using a true file system.

D2D has further advantages over D2T in that in midsized to large-scale implementations, D2D can lower the total cost of ownership of the backup and recovery solution due to increased automation of the process and lower hardware costs.

D2D2T

In D2D2T, data is initially copied to backup storage on a disk storage system and then periodically copied again to a tape or other removable media storage system. Traditionally, many businesses have done backup directly to relatively inexpensive tape systems. Many high-performance application systems, such as financial databases, however, have a production assurance or business continuity need to have their data immediately ready to be restored from secondary disk if and when the data on the primary disk becomes inaccessible.

As individual storage requirements have begun to be defined in terms of business criticality, rather than in terms of storage devices, organizations have adopted the concept of storage virtualization. In a storage virtualization system, IT managers can define an organization’s need for storage in terms of storage policies rather than physical devices. For example, if a financial database has a business requirement stating that no more than 15 minutes worth of data may be lost as the result of a technology failure, D2D2T makes a great deal of storage virtualization sense. Figure 6.3 demonstrates how D2D2T can be used as a backup architecture for such a production database.

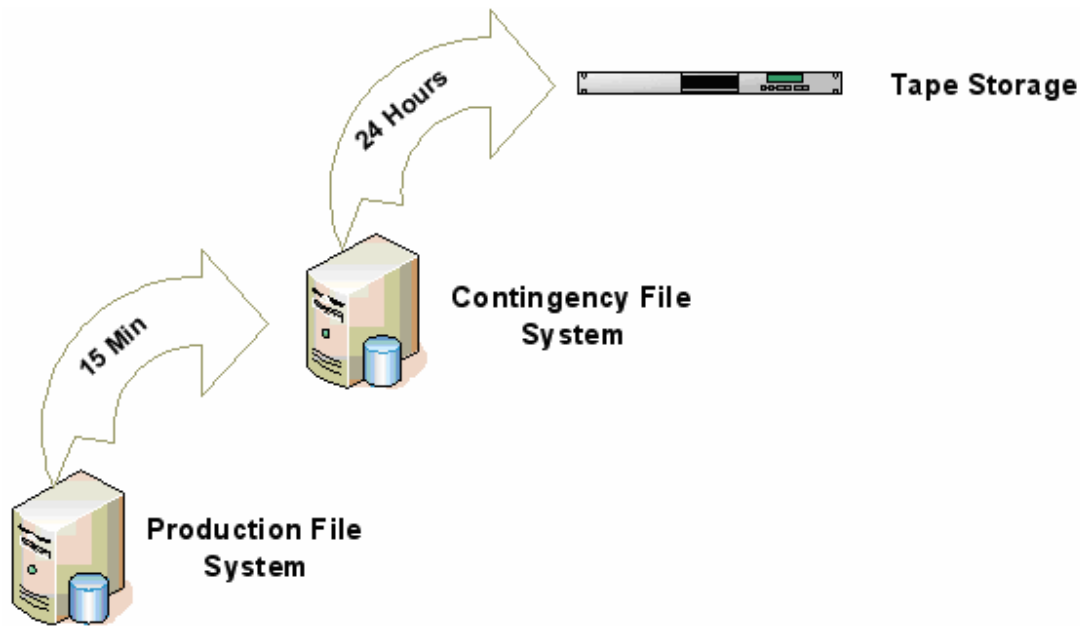


Figure 6.3: D2D2T and storage virtualization.

To meet a 15-minute goal, the supporting infrastructure would require a server to be dedicated to providing that contingency role. Every 15 minutes, data is backed up to the contingency server; then, at regular intervals, the contingency server can be backed up, which eliminates the demand for a full backup process from the production server. Other critical data, such as recent email, may also benefit from such a system. Email is considered by many to be mission-critical data in the short term, though dependency on this information eventually fades over time. D2D backup will enable email backups to be readily available for recovery in the short-term, then, as the organization's policy dictates, eventually moved for archival purposes.

D2D2T has further advantages in that once the data has been moved to a secondary device—either a dedicated server (as Figure 6.3 shows) or a disk dedicated to serving as backup media—the data can then be examined by other applications. For example, if one of the businesses being supported by the architecture desires reporting to be performed that is non-real-time and non-intrusive, the backup data present on the secondary disk may serve as a means to access that data without impacting production systems. Such methods should NEVER be used if the data itself is to be modified in any way, but for simple reporting based upon the data utilizing the backup copy, is a good way to reduce impact on production systems.

Benefits of Share-Data Approaches

Share-data approaches to storage architecture help to reduce the risk of data loss and increase productivity and collaboration as well as reduce backup and recovery processes and expenses, which is why they have become so popular. Over the years, as the ability to share data or “centralize” data storage has evolved, enterprises architects have battled to maintain availability, scalability, and manageability across their enterprises. However, the need for expansion and growth compounded by a sprawling file-serving architecture has crippled many storage solution initiatives. The benefits of share-data approaches are

- Reduction of complexity
- Increased performance
- Increased scalability
- High availability
- Consolidation of storage
- Simplification of management.

In many data centers, storage is an afterthought, something that is considered to be a byproduct of an application installation. Because of this mindset, storage is often dedicated to a server or collection of servers specific to the application systems they support. In a share-data approach, all servers can see all the data, storage is consolidated, and this architecture aggregates I/O performance and enables enterprise storage architects to greatly simplify storage management.

By approaching storage management with a share-data approach, the enterprise can consolidate existing storage and scale the environment as a whole without directly impacting any one application. This scalability means that a share-data approach is highly flexible, and by sharing the data throughout the storage infrastructure, the design is inherently fault tolerant allowing for failover with virtually no application disruption. Share-data approaches centralize, simplify, and holistically contain the storage infrastructure into a manageable solution.

Comparison: Consolidated vs. Distributed Backup Architectures

Consolidation and simplification will be the focus of countless IT projects this year as organizations strive to reduce their total cost of ownership and leverage new, higher-performing platforms available that result in increased storage density. Depending upon the project and subsequent storage application, consolidation can make a great deal of financial sense. The results of such a project can be felt directly on the bottom line by reducing hardware and software licensing costs, ongoing maintenance fees, and power and network requirements.

When comparing the architectures of consolidated and distributed backup architectures, several key points of contrast come to light. The first and broadest in scope is that there are many different hardware and software architectures and options to choose from—each supporting different or varying protocols and their own architectures that come together to develop the overall storage strategy. The following list highlights the key points to examine:

- Availability
- Scalability
- Interoperability
- Data protection

Distributed Approach

In a distributed environment, file servers and supporting infrastructure abound, and although traditional methodologies for providing high data availability have driven the enterprise storage infrastructure in this direction, such environments are more costly to maintain. High-availability has for years been the siren song of the distributed approach. In a distributed environment, there are more copies of the data on hand to meet the availability requirements of the business, but the concept of high-availability is not exclusive to distributed environments. Today, consolidated data storage can meet the demands of high availability as readily as many distributed approaches and in an architecture that lends itself more easily to centralized management and scalability.

Scaling a distributed architecture to meet the growing needs of the enterprise brings significant costs. In addition, new software, servers, drive arrays, and supporting infrastructure are brought online to meet growing needs—and the cost quickly adds up. Over time as new vendors present new options in storage, the sprawl of devices within your storage environment becomes increasingly difficult to manage. Storage administrators now face new challenges as a byproduct of years of distributing their environments and struggle to maintain interoperability of storage products within the enterprise.

Planning and providing for common data protection within the enterprise storage environment is a bit more difficult in the distributed approach. Consider, for an example, the process of moving large backup jobs in a distributed approach. With a dozen servers each being backed up independently to separate, distributed storage devices, the operational and subsequent disaster recovery of these systems can become quite complex. If experience has taught the enterprise storage community anything, it's that recovery needs to be as simple and swift a process as is technically and humanly possible.

Consolidated Approach

Server and storage simplification and consolidation are reaping great benefits to enterprise IT infrastructure. Technologies are continually being developed to increase the return on investment for costly mission-critical equipment. Servers, infrastructure, and storage devices aren't just costly to purchase, they're also expensive to maintain and support. Thus, leveraging these new advances can produce real financial savings.

Maintaining high-availability in a consolidated environment is no longer the exercise in futility it once may have seemed. However, many still compare the consolidated approach to storage architecture as putting too many eggs in one basket. The fact of the matter is that today consolidated storage networks can be just as highly available and resilient as their distributed counterparts. Redundancy methods and technologies—such as redundant host bus adapters (HBAs) to protect against cable failure, multi-pathing software, and resilient connectivity paths—have been developed to facilitate more highly available, robust, and resilient storage solution architectures.

A consolidated architecture excels in the scalability arena. Existing consolidated solutions can be scaled up much more easily than a distributed approach, which would often require not only the expense of the additional equipment but also the supporting architecture. Providing for common data protection is significantly simplified in a consolidated storage architecture. In its simplest form, a consolidated storage solution can offer a virtual one-to-one backup architecture. To illustrate the savings, consider an environment of 60 application servers geographically disbursed, each with their own individual storage requirements and supporting storage devices. To maintain common data protection throughout this environment, the backup and recovery efforts would need to be centered on those data storage devices and the subsequent supporting infrastructure; whether this results in individual tape drives for the servers or some other removable media, the cost to maintain the backups can become burdensome. Next consider the same 60 devices, but instead of a tape drive, they're all linked to a common SAN. The backup effort has now been reduced by a factor of 60 because the data has been consolidated to a single, albeit virtual, location. Figure 6.4 illustrates a potential layout for a dedicated SAN contingency environment.

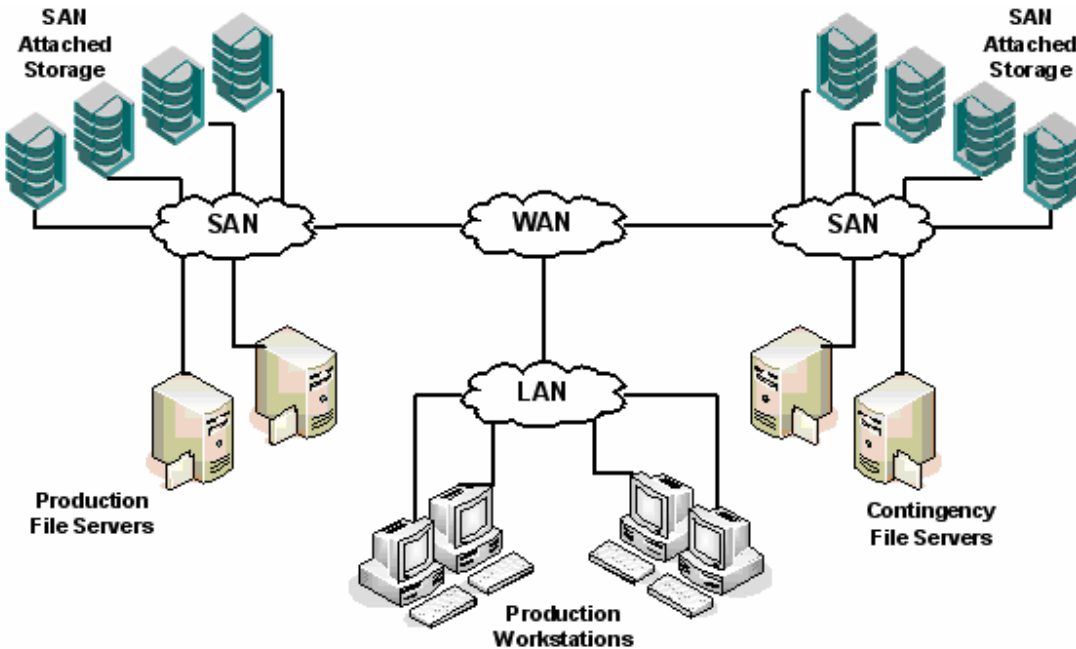


Figure 6.4: An example dedicated SAN contingency environment.

Data Recovery

The only reason to make a backup copy of any data is to be able to restore that data after it has either been lost or damaged. Data recovery is the process of recovering data from primary storage media when it cannot be accessed normally as a result of physical damage to the storage device or logical damage to the file system that prevents it from being mounted by the host OS.

Physical damage of a storage device can occur for a multitude of reasons ranging from malfunction of the physical inner workings of the device—such as a magnetic I/O head of a drive making physical contact with one of the drive platters—to disasters that impact the equipment directly as in the case of a fire or a flood. Most organizations lack the facilities, tools, and experience to recover physically damaged media in-house and must rely on external data recovery centers that specialize in the recovery of data from physically damaged media. This is a costly undertaking. Not only is data unavailable and the impact of the absence burdening the lines of business that depend upon it, but the entire process can be extremely expensive. The impacts and implications of physical damage are strong motivators to design and implement recovery solutions that reduce or completely eliminate dependency upon any one physical device.

Logical damage is by far the most common data recovery focus; fortunately, however, despite the ease with which logical damage can be inflicted, it is, to a great extent, offset by the relative ease (in comparison to physical damage) to which the damage can be recovered. Logical damage is often the byproduct of a sudden loss of power to a file storage device that prevents the file system structures from being completely written to the storage medium; however, problems with hardware, drivers, supporting infrastructure, and system crashes can have the same effect. The result is that the file system is left in an inconsistent state. This situation can cause a variety of problems, such as strange behavior (for example, infinite directory recursion, drives reporting negative amounts of free space), system crashes, or an actual loss of data.

Various programs exist to correct these inconsistencies and most OSs come with at least a rudimentary repair tool for their native file systems. Linux, for instance, comes with the *fsck* utility, and Microsoft Windows provides *chkdsk*. Third-party utilities are also available, and some can produce superior results by recovering data even when the disk cannot be recognized by the OS's repair utility.

The Advantages of Freedom

Throughout, this chapter has discussed many pitfalls to data storage integrity and availability. There are an overwhelming abundance of storage solutions vendors, each with their own architecture and agenda to fulfill. Although they may be working to align their interests with the needs of enterprise consumers, to date, there is no one single universally accepted miracle solution to meet all the needs of all consumers. The question then becomes which solutions most closely align with the needs of your enterprise and how can you leverage these solutions to meet those needs? Central to this concern is the ability for today's solutions to thrive in future environments. Freedom from proprietary solutions and standards, as defined by vendors, is critical in maintaining the flexibility and scalability of a storage architecture.

Benefits of Avoiding Proprietary Solutions

The drive for innovation is intoxicating. Vendors continually work to develop solutions and market those solutions based upon a foundation of their own technology to further their stake in the enterprise market; foundation being the operative word. When building an enterprise storage architecture, enterprise architects should approach their task with the same wisdom as a wise man who would build his house upon a rock rather than the sand.

Building an enterprise foundation that includes proprietary solutions is akin to building a house upon sand, which shifts over time and provides little stability against the elements—which, in storage architecture, are the battering of change and how that change impacts the ability of the solutions to be scalable, highly available, and resilient. To avoid the pitfalls of this approach, build a storage solution that embraces standardization and openness.

Proprietary solutions come with a price in that they can be difficult to manage and these difficulties often increase in direct correlation to the size of the implementation. Although a small organization can often withstand the year-to-year changes in proprietary solutions and standards, larger organizations have a much more difficult time weathering the storms. Proprietary solutions are, by definition, difficult, if not impossible, to integrate with other solutions, so although a proprietary solution that is being touted as highly scalable today may seem to meet the immediate needs of the business, a wise storage architect will weigh in other factors not the least of which is the solution's ability to be integrated with other key architectures.

Uncapped Scalability and Performance

If an enterprise file storage system is to be measured by any means it is in scalability and performance. For a solution to be viable, it must perform to meet the needs of the business, and to stay viable, it must scale to meet the growing needs of the organization.

In a distributed environment, scalability has been historically hindered by the dedication of servers to fulfill a specific storage role. A consolidated share-data approach is unhindered by those chains of bondage as new storage capacity can be brought online without directly impacting any of the applications supported by the environment. New servers and storage can be added as needed with no service disruption, enabling the storage environment to grow without directly impacting the business it is designed to support.

Some third-party solution vendors, such as PolyServe, Symantec (which acquired VERITAS), and IBM offer solutions that enable uncapped scalability and performance. As your enterprise storage requirements grow, such solutions' architecture facilitates the growth of the storage environment by providing flexibility in architecture and freedom of choice.

Architecture Flexibility

Some third-party solution architectures enable storage growth and the ability to "scale on demand," which means that your storage architecture can remain agile and flexible to meet the growing needs of your business. By providing for a centralized, consolidated yet highly available and flexible storage architecture, these solutions enable the enterprise to better serve changes in business requirements as increased demands for capacity and performance, whose cost at one time may have been prohibitive, are now being realized.

Freedom of Choice

The storage solutions offered by third parties such as PolyServe are not constrained by hardware platform or OS to the same degree other solutions may be and enables multiple low-cost Linux- or Windows-based servers to function as a single, easy-to-use, highly available system. For example, PolyServe's Matrix Server includes a fully symmetric cluster file system that enables scalable data sharing, high-availability services that increase system uptime and utilization, and cluster and storage management capabilities for managing servers and storage as one solution independent of hardware platform or supported application system. All of this equates to freedom of choice, and because storage administrators are unburdened by a dependency on any one particular hardware or software platform, they are free to reuse existing infrastructure or more closely align their storage infrastructure with enterprise IT hardware and software roadmaps easing administration, management, and compliance efforts.

Summary

Throughout this chapter, you have seen how storage solutions are often the critical pivot point for management and business concerns. The abilities of an enterprise storage environment to remain agile, reduce complexity, seamlessly integrate, and reduce risk are all concerns that have a direct financial impact on business operations. Disaster recovery planning has been underscored as a central theme to storage management and one that deserves due care throughout the life of your enterprise environment. Remaining flexible and not allowing your environment to become constrained by proprietary solutions will provide the freedom you require to provide agility and manage the environment as a whole.

As you have seen throughout the course of this guide, there are many disk, server, performance, and availability choices at your disposal—each with their own benefits and limitations. As you continue to work to bring harmony to your enterprise storage solution environment, battle current file-serving storage growth problems, provide for data path optimization, and set out to build high-performance, scalable, and resilient Windows and Linux file-serving solutions, remember to keep an eye on the big picture. Storage solutions are about more than just providing storage; they're about enabling the business to succeed by providing quick and easy access to data whenever and wherever it is required in a manner that is cost effective to implement, manage, and maintain.

Download Additional eBooks from Realtime Nexus!

Realtime Nexus—The Digital Library provides world-class expert resources that IT professionals depend on to learn about the newest technologies. If you found this eBook to be informative, we encourage you to download more of our industry-leading technology eBooks and video guides at Realtime Nexus. Please visit <http://nexus.realtimepublishers.com>.