



Building Better Protected Storage using Windows® Storage Server 2003

White Paper

Published: *January 2004*

Table of Contents

Introduction	1
Tape Backups.....	1
Replication-based Technologies	1
Causes of Data Loss	2
Replication Methods	3
Whole-file Replication.....	3
Application Replication	4
Hardware Replication	5
Software Replication	6
Clustering	6
Choosing a Replication Technology.....	7
Scenarios for Data Protection.....	9
Providing High Availability	9
Providing Effective Disaster Recovery	10
Enhancing Backup and Restore.....	11
Migration Projects.....	13
The Best Solution	15
Summary.....	16
For More Information	16

Introduction

Storage usage is growing at an unprecedented rate within companies today. Businesses are becoming increasingly dependent on continuous access to stored data. As the number of mission-critical servers and storage resources grow, so does the importance of protecting against service interruptions, disasters, and other problems that may threaten the organization's ability to provide access to its key data when needed. The expansion in the scope, size, and number of critical servers means that storage availability is growing in priority.

There are a number of strategies that can be employed to protect important data. Each strategy has its place, because each has particular characteristics that make it well- or ill-suited for particular tasks. This paper will examine four separate data protection strategies, then compare their merits for the most common business continuance scenarios.

Tape Backups

The most common method of storage protection is also the oldest: backup and restore to magnetic tape, which has been around for almost forty years and is still the bedrock of most recovery strategies. The cost per megabyte for tape storage is fairly low; it's easy to move tapes to secure offsite storage, and the technology continues to scale well for many applications. However, tape backups have limitations; namely, the amount of time required to back up and restore large volumes of data, and the accompanying latency between when the data was protected and when the loss occurs. Accordingly, much attention is being focused on replication-based technologies.

Replication-based Technologies

Replication-based technologies offer the promise of capturing a data set at a particular point in time, with minimal overhead required to capture the data or to restore it later. There are four main methods of interest in today's storage environments:

- **Whole-file replication** copies files in their entirety. This is normally done as part of a scheduled or batch process, since files copied while their owning applications are open will not be copied properly. The most prevalent use of this technology is for login scripts or other files that don't change frequently.
- **Application replication** copies a specific application's data; the implementation method (and general usefulness) of this method varies dramatically, based on the feature set of the application, the demands of the application and the way in which replication is implemented. This model is almost exclusively implemented for database-type applications.
- **Hardware replication** copies data from one logical volume to another; the copying is typically by the storage unit controller. Normally, replication occurs when data is written to the original volume; the controller writes the same data to the original volume and the replication target at the same time. This replication is usually synchronous, meaning that the I/O operation isn't considered complete until the data has been written to all destination volumes. Hardware replication is most often performed between storage devices attached to a single storage controller, making it less suited to replicating data over long distances. Most hardware replication is built out of SAN-type storage or proprietary NAS filers.
- **Software replication** integrates with the Windows operating system to copy data by capturing copies of file changes as they pass to the file system. The copied changes

are queued and send to the target server, while the original file operation continues its normal activity. These protected volumes may be on the same server, separate servers on a LAN or connected via storage-area network (SAN), or across a wide-area network. As long as the infrastructure is adequate, there is no restriction on the distance between source and target. The result is cost effective data protection. The example for this is NSI® Software's Double-Take®.

Causes of Data Loss

To best understand how to protect a particular set of data, it's also important to consider what the data is being protected *from*. Evaluating the usefulness of replication for particular conditions requires us to examine four separate scenarios in which replication might lead to better business continuity:

- **Loss of a single resource.** In this scenario, a single important resource fails or is interrupted. For example, losing the web front end that customers use for product ordering would cripple any business that depends on orders from the web; likewise, many organizations would be seriously affected by the loss of one of their primary mailbox servers. For these cases, some companies will investigate fault-tolerant architectures, but less will invest in fault-tolerance technology for file and print servers—even though the failure of a single file server may simultaneously prevent several departments' employees from accessing their stored data. Most planning for this case revolves around providing improved availability and failover for the production resource.
- **Loss of an entire facility.** In this scenario, entire facilities, and all of its resources, are unavailable. This can happen as the result of natural disasters, extended power outages, failure of the facility's environmental conditioning systems, persistent loss of communications, or terrorist action. For many organizations, the normal response to loss of a facility is to initiate a disaster recovery plan and resume operations at another physical site.
- **Loss of user data files.** This unfortunately common scenario involves the accidental or intentional loss of important data files; the most common mitigation is to restore the lost data from a backup, but this normally involves going back to the previous RPO—often with data loss.
- **Planned outages for maintenance or migration.** The goal of planned maintenance or migrations is almost always to restore or repair service in a way that's transparent to the end users.

Replication Methods

A thorough understanding of replication technology is useful for choosing the best way to protect critical data. This understanding begins with examining the logical layers of server I/O. Figure 1 shows the four levels, as they relate to server storage operations. With these layers in mind, we can now begin to appreciate the differences between replication philosophies.

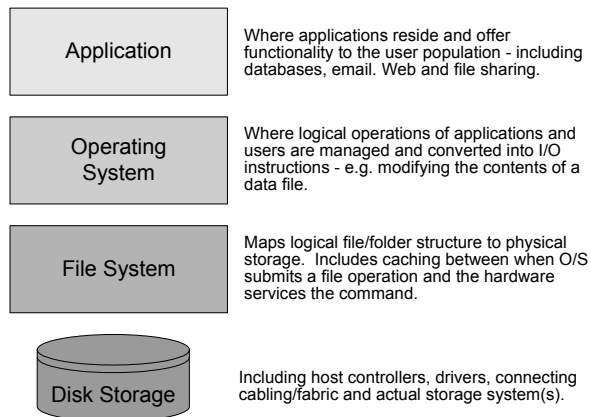


Figure 1: the logical layers of a server (for point of reference)

Whole-file Replication

The simplest method of replicating data is to copy the files, either manually or automatically. Methods include Windows Explorer drag-and-drop copying, scheduled XCOPY jobs, and automatic file copy tools. Whatever the method, whole-file replication copies only closed files (files that are not currently in use), and it lacks structured reporting, management, or auditing. Because of these restrictions, it's mostly useful as an ad-hoc method of distributing relatively static files (e.g. login scripts). The need for on-demand copies is still there—particularly for documents that need to be widely generated but have only one creation point. To provide a degree of automation and auditing, Windows server operating systems include support for the File Replication Service (FRS), and third-party vendors offer a variety of tools that distribute files automatically.

Limitations

Whole-file replication has two significant problems, both related to bandwidth usage. First is the problem of replicating file changes. If a user changes only a small fraction of a file, the file itself is still changed, and the modification date/time stamp reflects this. During the next replication cycle, the entire file will be transmitted, even though only a small portion of the file may have changed. That's why most tape backup arrangements perform both full backups (to accurately capture all data) and incremental backups (to capture changes without making unnecessary copies of unchanged files). Unfortunately, even when only part of any particular file is changed, tape backups must secure the entire file. To protect a file with any finer granularity, something other than whole-file technology is required.

Also, file-level replication tools generally don't provide any way to throttle the amount of bandwidth used by the copies; during the replication process, file copies may consume all the bandwidth between source and target—and this bandwidth includes the wasted copying of the unchanged portions of the data. Despite these limitations, in some environments, this approach can be effective. There are two primary requirements: the files must not be shared

between users (so they can be replicated without conflict), and the file size must be relatively small.

Application Replication

Application-centric replication takes advantage of special knowledge about an application's inner workings (including how often its data changes and which data items are most important) to tune replication for the best performance and utility. The application explicitly sends portions of its data to another instance of the application. For example, Microsoft's SQL® Server database has the ability to copy its transaction logs to another server at periodic intervals. This "log shipping" process preserves the log files, which are critical to recovering the database after a failure. Figure 2 shows an example.

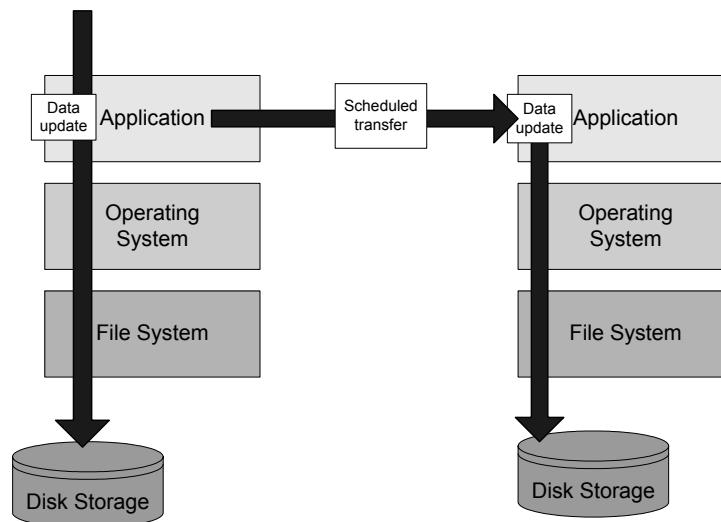


Figure 2: application replication

The application's architecture and capabilities have a great influence on replication; some applications can replicate only the data that has changed since the last job, while others must routinely compare the sets of data to see what's changed. The granularity of the data may be a field, record or complete table. Application replication is almost always a scheduled process, not a real-time copy.

Application-centric replication has the advantage that both instances are usable at the same time. Depending on the frequency of the scheduled replication job, this offers benefits like report generation and perhaps some level of redundancy/load-balancing.

Limitations

The biggest drawback to application-centric replication is that it's tied to a single application; applications that don't support replication must have their data copied by other means. In addition, application replication is a scheduled process, so the age of the data is based on how frequently the replication job occurs. Running it too seldom might cause an intolerable loss of data. Because replication uses CPU and memory resources on both the source and target servers, running it too often will degrade overall application performance for users. By contrast, the next two replication models can both be used to replicate data for many applications that have no built-in replication support. In fact, software replication can consolidate the protection of data from multiple application servers by using the highly scalable Windows® Storage Server 2003 as a target.

Hardware Replication

Unlike the other three replication models, in which the data continues to be available to the outside world in some fashion, hardware replication focuses on protecting the data so that it will always be available to the original server. This offers no protection against failures that damage the server, its operating system or applications, or other hardware components, so hardware replication is typically used in conjunction with clustering or other high-availability technologies. Figure 3 shows a typical hardware replication configuration.

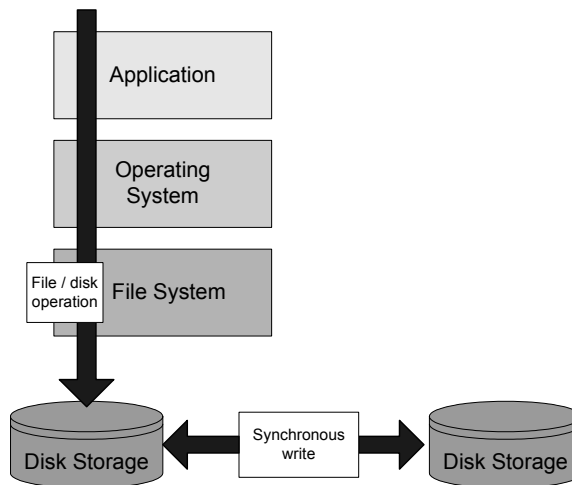


Figure 3: Hardware mirroring

Most hardware replication solutions use two, typically identical, and proprietary storage units. The units are joined by a Fibre Channel or other interconnects. Replication is handled by proprietary software (usually from the same vendor) that runs on the storage controller. In most cases, the storage, interconnect, and software all come from a single vendor and are sold, implemented, and maintained as a unit.

Functionally, hardware replication exists entirely in the lowest level of the server layers (see Figure 1). As disk write requests are passed from the server to the attached storage unit, the replication system takes over. For most hardware/synchronous solutions, the disk instruction does not immediately go to the primary storage unit. Instead, the request is queued while the write operation is performed on the secondary storage unit. Once the secondary unit has confirmed receipt of the instruction, the queued instruction is executed on the primary storage unit. This ensures that I/O is only committed to the primary unit *after* it's been replicated.

Limitations

Performing mirroring in the manner described above ensures that all transactions are the same between both copies of data, because the two storage devices are in lockstep. However, the drawback is that if the devices are separated so that they cannot use local interconnects, one of two things will occur. If bandwidth is not adequate, both the source and target systems will fall behind. Purchasing high-bandwidth connections almost always raises the TCO by requiring expensive long distance lines between storage solutions. Hardware mirroring solutions require more expensive (and duplicate) hardware, and the requirement to keep both devices in lockstep can be a performance limitation. The most common place for these solutions is where the value of data lost (between copies) greatly exceeds the cost of the solution (as is the case with real-time stock trading).

Software Replication

Software replication can be thought of as hardware replication without the hardware. Double-Take®, by NSI® Software, replication software installs drivers that filter and capture I/O requests from the OS as they pass to the filesystem; before the request is given to the hardware. At this point, the transaction can be sent to the remote replication target over any network interconnect (usually simple server-based IP communication).

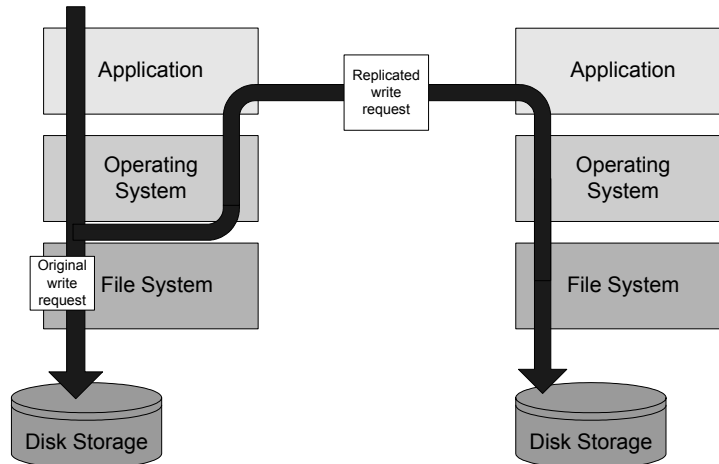


Figure 4: Software replication

The advantage to this approach is that it is application-independent; applications do not have to be modified or re-installed to use it; and in most cases, the application will never be aware that its data is being replicated. Coupling software replication with Windows Storage Server 2003 devices allows source systems to easily be replicated to a local or remote storage appliance, providing quick recovery in the event of a failure.

NSI's software products replicate individual I/O operations at the byte level, so that if a file change encompasses 12 bytes, then 12 bytes are queued for replication (not the 64KB block within the storage array and not the entire file). Software, or "Host-Based", replication tools protect files and folders on a volume; which means that 20GB of data on a 300GB volume does not require a 300GB volume on the target – merely 20GB of available storage. Double-Take allows both throttling and queuing of replication traffic – so that it can be deployed over existing WAN infrastructure links. This queuing technology enables solutions where rapidly changing data must traverse slow WAN links.

Clustering

Most discussions of data availability include mention of clustering. However, clustering is not particularly useful for storage protection. Clustered systems provide a method by which two or more server nodes have a logical relationship; work may be shared between the nodes or moved from node to node as failures occur. Windows clusters share physical storage devices, with the logical volumes on those devices being owned by one cluster node at a time. Figure 5 illustrates a traditional cluster.




Figure 5: Traditional clustering

In clustering's purest form-- two nodes and one shared disk-- this configuration results in a single point of failure. Adding more redundant storage is one way to work around this problem; however, the problem remains: clustering provides *service* availability, not any intrinsic replication or protection for *data* availability. A common implementation is to cluster the servers *and* replicate the storage (using some form of cluster replication software like NSI's GeoCluster™ product). Figure 6 shows a typical implementation.

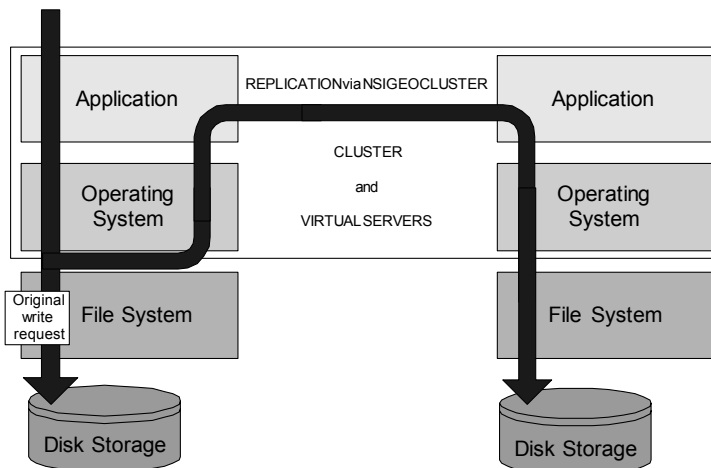


Figure 6: Clustering with NSI GeoCluster

Choosing a Replication Technology

For most environments, whole-file and application technologies fit only a small portion of organizational requirements for data replication. That raises the question of whether hardware or software replication is a suitable answer for applications for which whole-file and application replication don't work well. This is especially true in light of the fact that for the small portion (where whole-file and application does suffice), hardware and software might work as well. This would allow for a single protection model within the enterprise, regardless of data type – which is compelling to many corporations. The question then becomes whether hardware or software is a better implementation. There are a number of factors involved in this choice.

Cost

First and foremost, hardware mirroring systems require duplicates of what is already relatively expensive hardware, so the solution tends to be expensive. According to the Gartner Group, only about 0.4% of deployed servers actually require the expense and level of fault-tolerance provided by synchronous hardware mirroring.¹ Statistics from Enterprise Storage Forum indicate that approximately 15% of servers have data that is perceived as valuable enough to merit special protection, meaning that protecting all servers with hardware replication is unlikely to be cost-effective. Double-Take is inexpensive compared to hardware replication systems, and the industry-standard hardware used in Windows Storage Server 2003 makes the combination a very cost-efficient way to protect one's data.

Replication Granularity

The amount of data that must be replicated after a change is also important – regardless of the total size of the file. If an application produces a 150-byte write request, then software replication solutions replicate 150 bytes (plus some minor command overhead). Hardware replication systems use the disk block as their basis for replication, since that's what they have access to. Most storage systems opt for larger block sizes (64-128KB) to provide more efficient striping performance; for a 150-byte change, instead of software replication's 150 bytes, hardware block-level mirroring would need to transmit 64,000 bytes.

For write operations whose contents span more than one block, multiple blocks must be replicated. This means that hardware replication demands significantly higher bandwidth interconnects, which will be more expensive than the LAN or WAN-speed links that software replication can use, and therefore raise the TCO of the protection solution.

Because Double-Take has access to the actual file change instructions, it is able capture and replicate data at a granularity that is unequalled – regardless of the application. This contributes to its efficient network usage. The hardware and software architecture of Windows Storage Server 2003 devices optimizes those file instructions for high performance under high network loads, making them a natural target for ongoing replication.

Latency and Load

In order to ensure that hardware solutions are synchronous, the flow of data from the production server to the production storage must be detoured. Write operations on the production machine are queued on the production storage system; this queuing allows transactions to be sent to the redundant array, acknowledged, and applied synchronously to the production disk. While it is true that both arrays will have identical data, keeping both arrays in sync requires an expensive amount of I/O bandwidth. The alternative is to maintain both servers in lockstep, with the result that large numbers of application I/O requests remain queued on the source.

Asynchronous replication doesn't cause this problem; the production server's write requests are applied to the production disk at normal speeds. A copy of those changes is then sent to target platform at best available network speed. In most cases, the target will be seconds or less behind the source; while the source disk remains current at all times.

With all replication technologies, there is some minor amount of latency. The question is whether it's better to have that latency between network-connected servers or between the application and its own storage resources. When the purchase and maintenance costs of hardware solutions are considered, most organizations will find that software-based replication is a better fit for their needs and budget.

¹ Gartner, IT Trends for 2002.

Scenarios for Data Protection

The approach of replicating data in real time offers a potential escape from the cost-versus-recoverability dilemma. The phrase “business continuity” covers a broad spectrum of technologies, processes, and planning approaches; evaluating the usefulness of replication for particular conditions requires us to examine four separate scenarios in which replication might lead to better business continuity: high availability, disaster recovery, backup and restore, and migration.

Providing High Availability

Perhaps the most commonly envisioned approach to continuous business operations is that of failover, in which users are transferred from one computer to another in the event of a failure. Failover-based approaches assume that the users still have desktops, power and connectivity – so the outage is that of a failed server resource. The goal for high availability (HA) solutions is to keep the users productive by quickly restoring access to the failed resource. With this goal in mind, let's examine the technologies described above.

- Whole-file replication can provide access to an alternate copy of important data. This doesn't usually happen automatically, unless the sites are using a technology like the Windows Distributed File System (DFS) that abstracts the user from the actual location of their data. The primary problem with file-based replication is that the data is as old as the last scheduled replication push.
- Application-centric replication behaves similarly, but the user client software would most likely have to be manually redirected to the alternate application server.

For this reason, when most IT professionals think about high availability or failover, whole-file and application-based replication solutions are not typically satisfactory; they make failover visible to clients. Other technologies are better suited to high availability designs:

- Clustering is designed exclusively for high availability and handles it well. Unfortunately, as described earlier, a typical cluster still has a single point of failure: its shared storage subsystem. By definition, a highly redundant system should not have single points of failure; hence the need for hardware mirroring or software replication as a supplement.
- Hardware mirroring involves making exact copies of the data; the storage controller already abstracts the servers from the storage. As long as the server is functional (or can be rebuilt or repaired), it can simply access the redundant array transparently, without concern about which replica it is actually using. Of course, if the server cannot quickly be restored, hardware mirroring doesn't help, hence its usual deployment in conjunction with clustering.
- Software replication fills an important gap in the HA world. While one's most critical servers might already be clustered or protected with hardware replication, the remaining vast majority -- which are important, but for which HA isn't perceived as cost-effective—can be protected by replicating their data. In many corporate environments today, file servers tend to be unprotected, even though software replication provides an easy and reliable way to copy many servers' data to a single replication target.

Double-Take and Windows Storage Server 2003 provide an effective means of replicating data from file and application servers. As shown in Figure 7, this solution can be combined with hardware replication and clustering to provide cost-effective protection of a larger percentage of the most critical data.

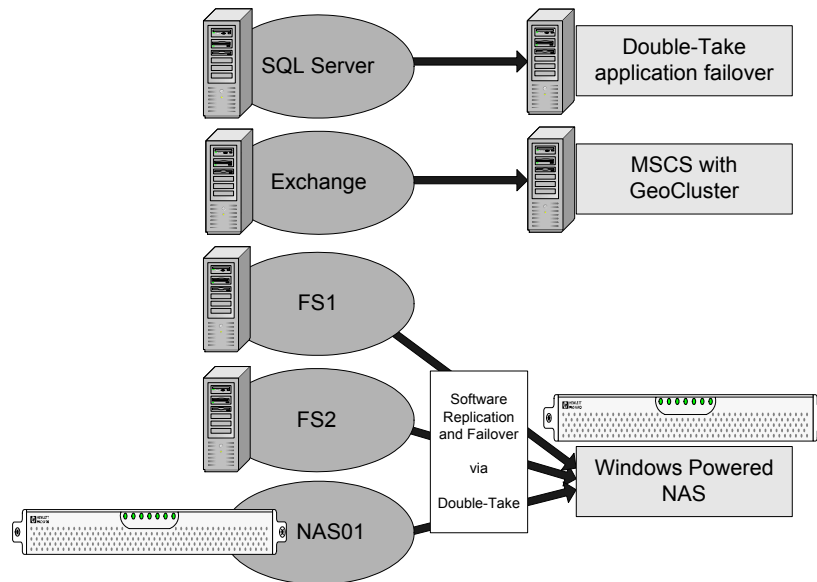


Figure 7: Replication for high availability

Providing Effective Disaster Recovery

Many people use the term "business continuity" and "disaster recovery" interchangeably. In this paper, business continuity refers to the entire realm of efforts, while disaster recovery (DR) is focused exclusively on protecting the business by protecting the data at an alternate location.

Depending on the crisis that drives the recovery, DR may take several different forms. In the most complex scenario, the complete failure, destruction, or interruption of access to a computer room might necessitate moving the company's operations and personnel to an alternate set of servers at another location. Simpler recoveries might involve restoring operations after damage to the primary copy of the data. For the purpose of this discussion, we will focus on the survivability of the data alone (and hope that the clients reading this already have a plan for how to recover, now that they have confidence that their data survives). Examining the five technologies outlined earlier, we see that:

- Whole-file replication can move the data to an alternate location; however, due to bandwidth considerations, whole file replication across a WAN has all of the same performance detriments that a tape backup across the WAN would have, so it is rarely a viable solution for ongoing recovery planning.
- Application replication is subject to the same bandwidth concerns as whole-file replication; additionally, not all applications support it.
- Clusters use shared storage, so there is no benefit to separating the nodes while only one location has the data store. Without distance, it isn't "disaster recovery".

For these reasons, most organizations provide disaster recovery by making tape backups and storing the tapes in secure offsite facilities. However, hardware and software replication offer some compelling advantages. Hardware mirroring is capable of protecting its storage across extended distances; various storage manufacturers sell hardware extends storage protection across hundreds or thousands of kilometers. Unfortunately, as described earlier, for both copies of the data to be synchronous, there must be near-zero latency between sites. Any latency over the distance will cause both copies of the data to be equally aged. However, ensuring very low latency over long distances requires paying for large amounts of available bandwidth, raising the ongoing cost. For most companies, hardware mirroring is not cost effective beyond the boundaries of large cities.

Double-Take's greatest strength in this area is that it can operate efficiently by only replicating the data that's changed. Combined with bandwidth throttling and queuing, this allows software replication to work well over long distances, even with Internet-quality WAN links. Furthermore, Double-Take can easily mirror several servers to a single target; the source servers can be individual servers or Windows clusters. Concentrating replication to a single target server plays to the strong scalability and robust performance of Windows Storage Server 2003 solutions. Figure 8 shows a sample implementation.

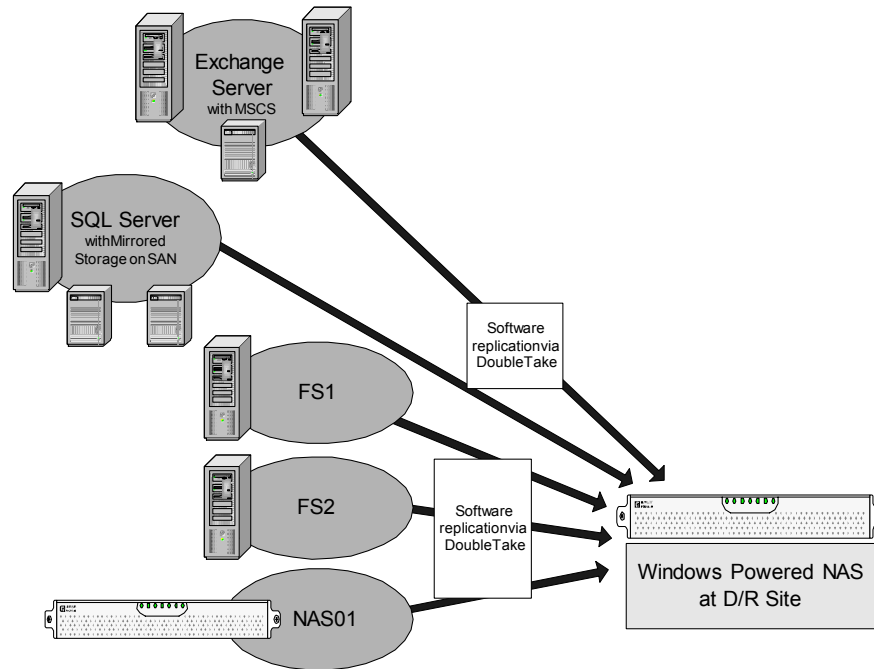


Figure 8: Software replication makes disaster recovery more efficient

Enhancing Backup and Restore

For a surprising number of companies, tape backup continues to be their only preparation for business continuity. The challenge with this approach is the ever-increasing restore times driven by the growth in data volume and change rates. Consider a typical scenario involving offsite storage; assume that full backups are done every weekend, with nightly incremental backups. Off-site storage is used for continuity protection. A failure that occurs at 4pm Tuesday must be recovered with the previous weekend's full backup and the Monday night incremental—but if that tape has already gone offsite, it must be retrieved, which can add hours (if not days) to the recovery time. Even if the tape can be retrieved with only a four-hour lead time, that still means that users won't have access to the Monday version of their data until sometime on Wednesday (and Tuesday's data is completely lost). For many companies, this is not practical. Let's examine the storage technologies discussed earlier in regard to this problem.

- Clustering offers no way to enhance backup or restore processes. In fact, there's an additional complication: backup subsystems are often kept on external hosts, and the cluster is backed up remotely. This adds extra time and complexity to the backup process.
- Whole-file replication does provide a second copy for backup purposes (within the latency parameters discussed earlier). However, most replication applets do not properly handle the situation where the target data is being backed up; if the backup

software locks files as it backs them up, replication may fail until the files are unlocked again.

- Similarly, most application replication tools do not deal well with the target data set being locked for backup.

As with disaster recovery, hardware and software replication offers more flexible approaches. Most hardware replication solutions offer various backup enhancements, including freezing one set of data while the other is given over to the backup (which may be host- or storage-attached) and making “snapshot” or point-in-time copies of the data. The only potential caveat is the re-synchronization time required for the frozen data set once it’s thawed and updates are allowed to happen.

Software replication via Double-Take can offer similar benefits with a different twist (see Figure 9). Unlike hardware solutions, where one logical copy of the data exists in two arrays, the two data sets in software replication are only loosely coupled. This means that while the production data is locked and in use, the redundant copies are natively in a closed state (except of course when each file is actually being updated). During the remainder of the time, tape backup software has easy access to the replicated copies, and they can be backed up without placing any additional network or CPU load on the production server; and without the need for expensive backup agents. Changes will continue to be sent from the source to the target and applied after the tape backup is complete. This, in combination with the Windows Storage Server 2003 ability to freeze iterations of a data set, provides for a great enhancement over traditional backup, while still allowing you to leverage your existing investments in tape hardware and software.

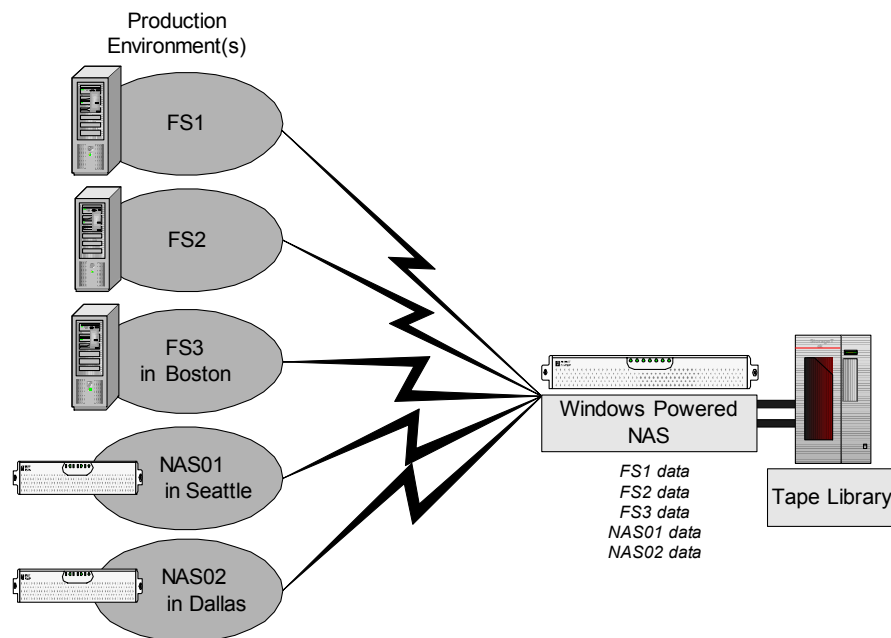


Figure 9: Software replication makes backup and restore more efficient

As a file-handling platform, Windows Storage Server 2003 devices are especially well-suited as the replication target and tape backup media server. Because Windows Storage Server 2003 devices scale very well with additional storage, they are particularly useful for providing centralized backup of branch or remote offices.

Migration Projects

One important business continuity aspect is often overlooked: not all outages are unplanned. As an example, server migrations are typically planned to occur over weekends, holidays, or other periods of reduced user demand. However, during those times, the server is still unavailable to the users; by definition, this means that the business is not continuous.

Data migrations usually involve extended outages—and thus consume weekends or holidays—for two reasons: the files must typically be left dormant long enough to move them without users updating them in mid-move, and a "point of no return" must be defined so that if the migration is unsuccessful, it can be rolled back so that production resources are available for the next business day.

- Whole-file copies (in the guise of ad-hoc copy/move operations) is the most common way to accomplish migrations today. This approach suffers from both the dormancy and point-of-no-return requirements described above.
- Application replication cannot normally be used for migrations because it's specific to particular applications.
- Likewise, clustering generally doesn't offer any benefits for migration projects.
- Hardware mirroring solutions aren't cost-effective when used solely for migrations. More commonly, the requirement is to migrate from a local disk to a managed hardware storage solution, in which case hardware replication solutions can't be used because the hardware is not yet in place. In those cases, software replication is typically used to get the data to the hardware array; thereafter, it could be protected by any hardware or software.
- Due to the relatively low cost of software replication solutions, they are gaining ground in the migration market. Software replication allows migrating from one version of Windows to another; it can also be used to migrate from old hardware to new or from conventional servers to consolidated Windows Storage Server 2003 devices.

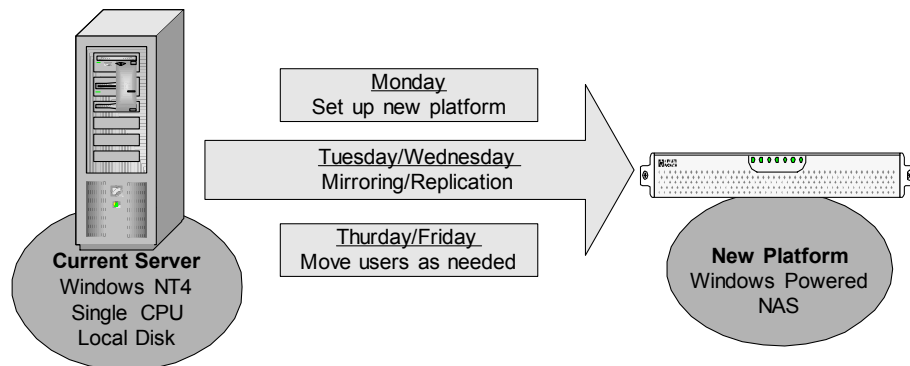


Figure 10: Replication and migration from stand-alone servers to Windows Storage Server 2003 devices
Replicating and migrating from stand-alone servers to Windows Storage Server 2003 appliances (as shown in Figure 10) offer some attractive benefits:

- Software mirroring (establishing the baseline) and replication (which copies changes to the baseline) can start during the workweek, so that the data on the existing production server is sent to the new platform. By using Double-Take's scheduling and throttling features, the mirror can be tuned to have minimum impact on the production server.
- As soon as the initial mirror completes, test users can be pointed at the new resource. If all goes well, the remaining users can be redirected sooner than originally planned; if

problems occur, the production server is still online, available and current. In this scenario, there is no “point of no return”.

- Instead of disabling user access at 6 PM on Friday and working all weekend, the new target can be brought online on Monday morning, migrated during the week, and the I/T team’s weekend is saved.

Because replication allows the data to be moved while preserving user access, many migration and consolidation projects no longer require weekend efforts.

The Best Solution

The ideal solution for protected storage would combine low acquisition and maintenance cost with fine-grained replication that could be scheduled or throttled to avoid placing excess load on production systems or networks. The combination of Windows Storage Server 2003 devices and NSI's Double-Take provides these benefits:

- **Low cost.** Instead of requiring expensive, proprietary hardware replication solutions, Double-Take can replicate data on servers running Windows NT, Windows 2000, or Windows 2003 to Windows Storage Server 2003 devices, which offer better scalability and lower costs than traditional storage systems. Customers like Continental Airlines and the Texas Children's Cancer Center are realizing lower acquisition and maintenance costs for their storage after moving to Windows Storage Server 2003 solutions.
- **Flexible replication.** Double-Take only replicates the bytes actually changed by each write, not the entire block or the whole file. When compared with block-mode replication solutions, this approach offers lower load on the production servers, faster update, and the ability to send replication updates across wide-area networks. Because Double-Take allows replication updates to be scheduled and throttled, administrators can tailor replication resource usage for their specific environments.
- **High scalability.** Microsoft's Windows Storage Server 2003 partners have designed devices that offer high degrees of storage scalability; organizations can choose the right amount of storage to start and easily grow to larger devices as their requirements change.

In addition, the wide range of vendors supporting Windows Storage Server 2003 makes it easy to choose the right hardware feature set. Because Windows Storage Server 2003 devices operate as members of a Windows domain, they seamlessly integrate with existing infrastructure, so they're easy to set up and manage. Double-Take is able to protect all of them.

Summary

Storage protection strategies fall into four general areas: high availability, enhanced backup and restore, disaster recovery, and migration. Each of these areas is important. Most organizations focus on high availability and disaster recovery only for the systems they perceive as most critical, based on the belief that protecting file and print servers costs too much. Likewise, enhancements for backup/restore and migrations are often dismissed for cost reasons.

The combination of Double-Take's powerful, flexible replication software and the Windows Storage Server 2003 platform enables cost-effective, easy-to-manage protection for file, print, and application. When taken together, Windows Storage Server 2003 devices and Double-Take provide superior scalability, flexibility, and TCO compared to hardware- or application-based replication systems.

For More Information

- Microsoft Windows Storage Server 2003 home page: <http://www.microsoft.com/storage/>
- NSI Software's Double-Take home page: <http://www.nsisoftware.com/nas/>



The information contained in this document represents the current view of Microsoft Corporation on the issues discussed as of the date of publication. Because Microsoft must respond to changing market conditions, it should not be interpreted to be a commitment on the part of Microsoft, and Microsoft cannot guarantee the accuracy of any information presented after the date of publication.

This White Paper is for informational purposes only. MICROSOFT MAKES NO WARRANTIES, EXPRESS, IMPLIED OR STATUTORY, AS TO THE INFORMATION IN THIS DOCUMENT.

Complying with all applicable copyright laws is the responsibility of the user. Without limiting the rights under copyright, no part of this document may be reproduced, stored in or introduced into a retrieval system, or transmitted in any form or by any means (electronic, mechanical, photocopying, recording, or otherwise), or for any purpose, without the express written permission of Microsoft Corporation.

Microsoft may have patents, patent applications, trademarks, copyrights, or other intellectual property rights covering subject matter in this document. Except as expressly provided in any written license agreement from Microsoft, the furnishing of this document does not give you any license to these patents, trademarks, copyrights, or other intellectual property.

Unless otherwise noted, the example companies, organizations, products, domain names, e-mail addresses, logos, people, places and events depicted herein are fictitious, and no association with any real company, organization, product, domain name, email address, logo, person, place or event is intended or should be inferred.

© 2003 Microsoft Corporation and NSI Software. All rights reserved.

Microsoft, Windows Powered, Windows, Exchange, and SQL Server, are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries.

Double-Take and NSI are registered trademarks of Network Specialists, Inc., GeoCluster is a trademark of Network Specialists, Inc. and all are used with permission of the trademark owner.

The names of actual companies and products mentioned herein may be the trademarks of their respective owners.