Protecting Data with HP ProLiant Storage Servers and HP OpenView Storage Mirroring

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.

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 is 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.

Because of this limitation, tape backup alone is considered insufficient when developing a business continuity strategy. Modern, high-available systems will include several layers of hardware, service, and data redundancy to ensure that information is always flowing:

- Duplicate components—Many servers contain duplicate hardware components that can be swapped out when failed, without taking the server down.
- Duplicate services—Software services such as file serving, print serving, mail serving, and web serving are all examples of services that can be interrupted if a server goes down. Clusters of servers provide duplication of services that allow continuity of service should one server fail.
- Duplicate data—Data duplication comes in two primary forms:
  - RAID technology—RAID is used to duplicate data across disk drives within an array to protect against one or more drive failures.
  - Replication technology—Replication (mirroring) is employed to generate real-time copies of specific data sets across two separate storage systems. This goes beyond RAID technologies by ensuring data availability even when an entire storage system goes down.

This paper examines in detail the various data replication technologies available today and compares their relative strengths and weaknesses. The discussion also includes a survey of several common business continuity scenarios, and introduces potential solutions for each.

Causes of data loss

To best understand how to protect a particular set of data, it is also important to consider what the data is being protected *from*. Evaluating the usefulness of replication for particular conditions requires examination of 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 fewer 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 their 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 is 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, you can now begin to appreciate the differences between replication philosophies.

Figure 1. The logical layers of storage I/O in a server

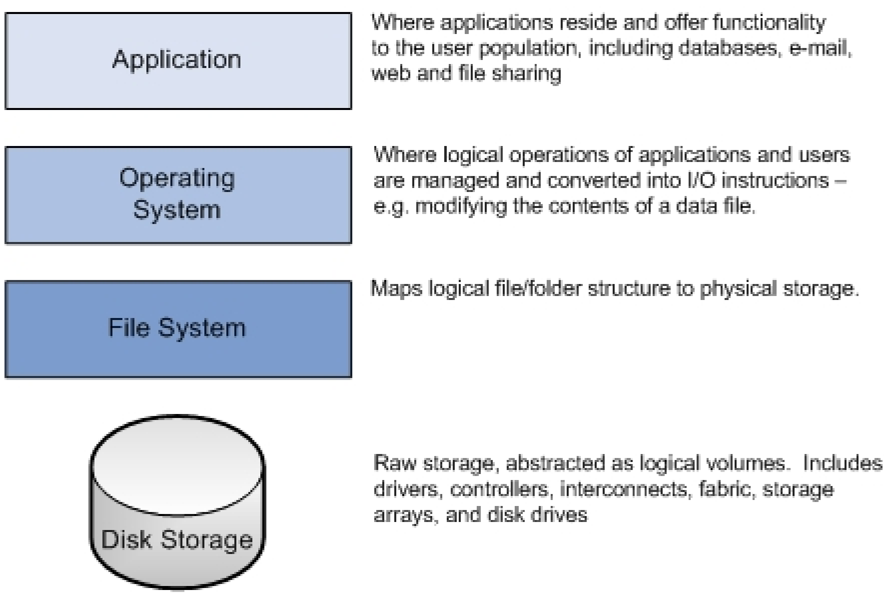

Whole-file replication

Whole-file replication copies files in their entirety. The simplest method of replicating data is to copy the files, either manually or automatically. Methods include Microsoft® 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 is mostly useful as an ad-hoc method of distributing relatively static files (for example, login scripts). The need for on-demand copies is still there, particularly for documents that must 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 is 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. In addition, file-level replication tools generally do not 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 replication generates copies of a specific application's data. 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

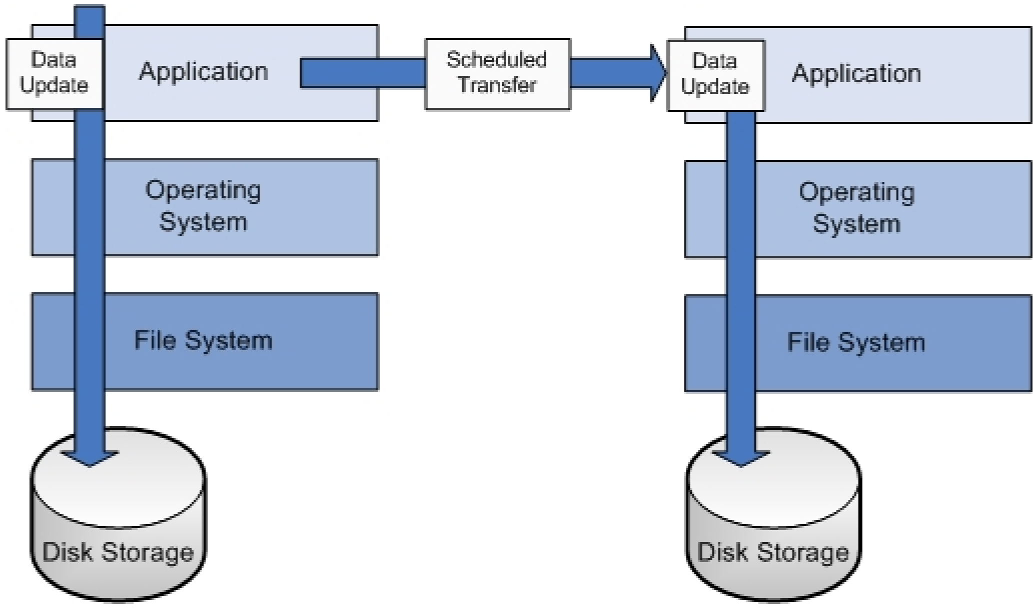

The application's architecture and capabilities have a great influence on replication; some applications can replicate only the data that has changed since the last job, while others must routinely compare the sets of data to see what has changed. The granularity of the data may be a field, record, or 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 is tied to a single application; applications that do not 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. This model is almost exclusively implemented for database-type applications.

Hardware replication

Hardware replication involves the copying of data between storage array devices at the block level. Storage arrays provide storage in abstract units called logical volumes, or LUNs. Each LUN may be replicated on separate arrays for protection and availability. Normally, replication occurs when data is written to the original volume; the array 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 is not considered complete until the data has been written to all destination volumes. Hardware replication is most often performed between storage devices over Fibre Channel networks (SAN), making it less suited to replicating data over long distances.

Unlike the other 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 replication

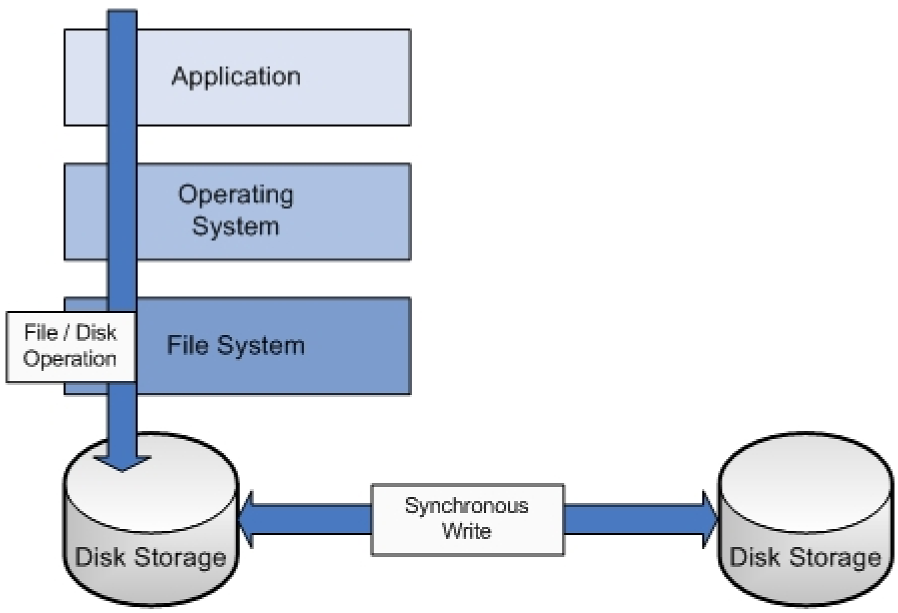

Most hardware replication solutions use two, typically identical, and proprietary storage units. Replication is handled by proprietary software (usually from the same vendor) that runs on the storage controller. In most cases, the storage, interconnect, and software all come from a single vendor and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

Functionally, hardware replication exists entirely in the lowest level of the server layers (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. When 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 has been replicated.

Limitations

Performing mirroring in the manner previously described 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 integrates with the operating system to copy data by capturing copies of file changes as they pass to the file system. The copied changes are queued and sent 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 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 of software replication is HP OpenView Storage Mirroring.

Software replication can be thought of as hardware replication without the hardware. HP OpenView Storage Mirroring 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

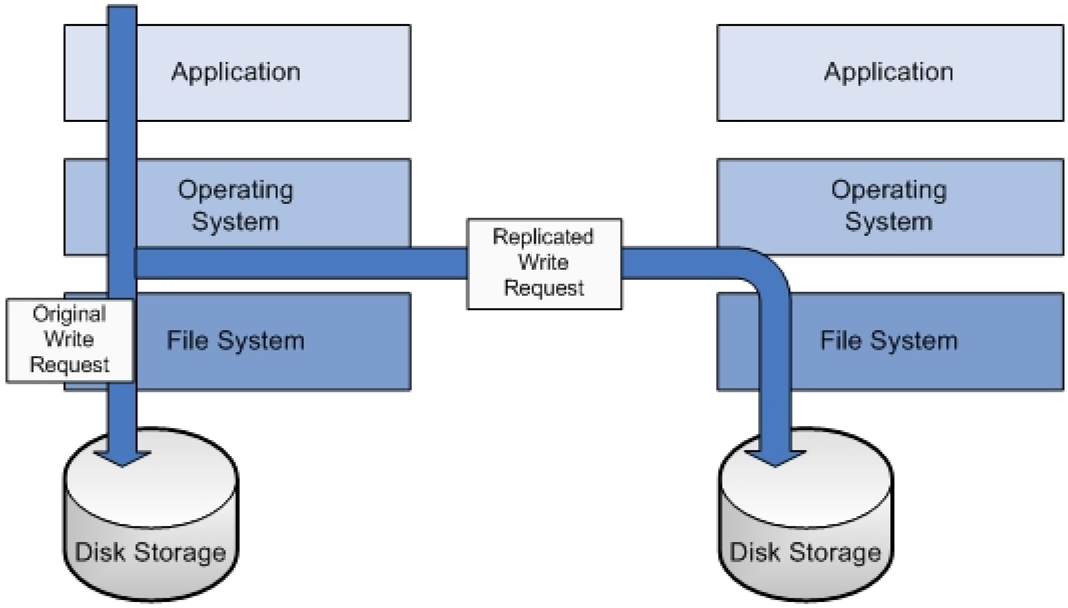

The advantage to this approach is that it is application-independent; applications do not have to be modified or re-installed to use it; and in most cases, the application will never be aware that its data is being replicated. Coupling software replication with HP ProLiant Storage Server allows source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

HP OpenView Storage Mirroring replicates 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 64-KB 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 20 GB of data on a 300-GB volume does not require a 300-GB volume on the target—merely 20 GB of available storage. Storage Mirroring 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.

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 do not 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. 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. HP OpenView Storage Mirroring is inexpensive compared to hardware replication systems, and the industry-standard hardware and software used in HP ProLiant Storage Server 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 is what they have access to. Most storage systems opt for larger block sizes (64–128 KB) 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 HP OpenView Storage Mirroring has access to the actual file change instructions, it is able capture and replicate data at a granularity that is unequaled—regardless of the application. This contributes to its efficient network usage. The hardware and software architecture of HP ProLiant Storage Server optimizes those file instructions for high performance under high network loads, making them a natural target for ongoing replication.

Latency and load

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 does not have 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 is 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.

Scenarios for data replication

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 examination of 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 limited to that of a failed server resource. The goal for high-availability solutions is to keep the users productive by quickly restoring access to the failed resource.

High-availability solutions typically come in one of three forms:

- Clustering—A high-availability cluster is designed exclusively for high availability of a set of services by duplicating the services on multiple platforms (servers). Unfortunately, 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 or software replication as a supplement.
- Hardware mirroring—The replication technologies available in storage arrays provide for the redundancy of the data. These solutions are very fast, very reliable, and very expensive. Unfortunately, the solution does not provide for the redundancy of the services that serve the data to clients. Therefore, hardware mirroring is usually deployed in conjunction with clustering to provide full redundancy.
- Software mirroring—The advantage of software mirroring is that it is designed to provide redundant services and data in one package. It is important to note that software mirrors do not duplicate all the features and failover capabilities of a cluster, and they do not maintain complete consistency between mirrors as in hardware mirrors. However, because of the cost, clusters and hardware mirrors are usually reserved for only the most critical corporate data. All other data, which often makes up the vast majority of corporate information, is typically not highly available. This is where software mirroring may be deployed as an economical high-availability solution. Corporate servers that would otherwise go unprotected can benefit from replicating the data from many servers to a single replication target. If any server goes down, the replication target may temporarily take over and serve the duplicate data in place of the failed server.

HP OpenView Storage Mirroring and HP ProLiant Storage Server provide a cost-effective and efficient means of replicating data on file servers. As shown in Figure 5, this solution can be deployed to provide a layer of protection and availability for a large percentage of important corporate information.

Figure 5. Replication for high availability

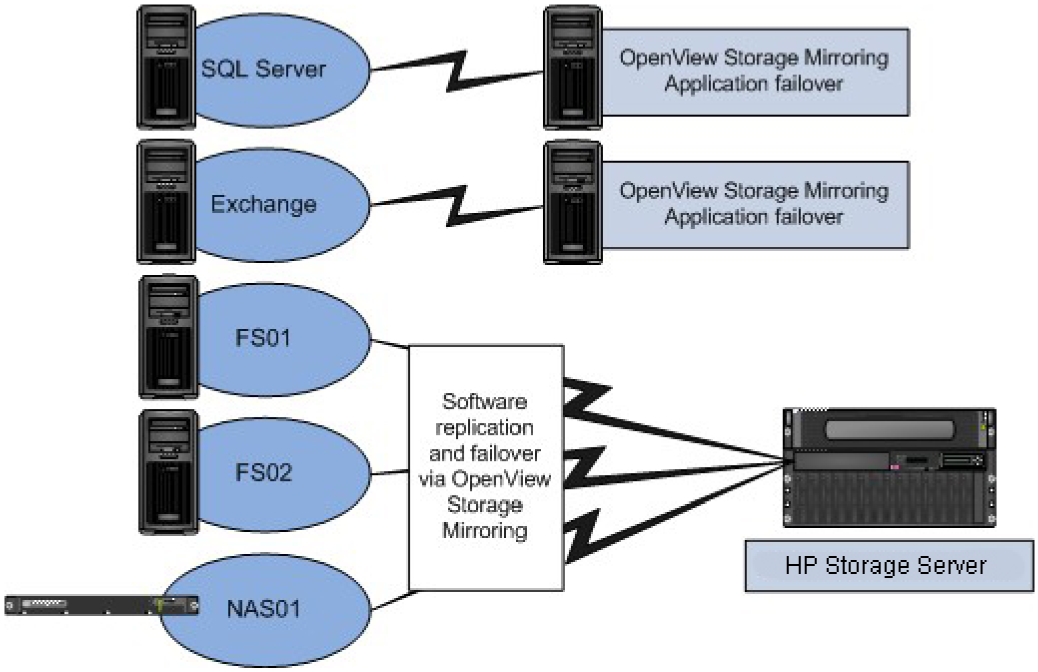

Providing effective disaster recovery

Natural and man-made events that interrupt the availability of information must be accounted for in any business. Planning on how to restore normal operations after an interruption is critical to business success. A key element in the planning is to determine not only the value of the information, but the value of the "availability" of the information as well. A database of scientific research may contain very valuable information, but it does not necessarily have to be immediately available after a disaster. On the other hand, a customer ordering and billing system can afford little downtime in most businesses.

In the simplest terms, disaster recovery involves the duplication of data and the storage of that data in a separate location. The intent is to ensure that the copy of the data will survive whatever disaster that the company anticipates. Beyond this, disaster recovery implies the process of restoring the data online after a disaster has occurred. The company must also evaluate the value of having the data "online" and available for use. This decision will dictate the technology used to implement a disaster recovery plan.

Most organizations provide disaster recovery by making tape backups and storing the tapes in secure offsite facilities. While this method is cost-effective and provides a fair measure assurance that a copy will survive a disaster, the solution lacks the ability to quickly recover. Following a disaster, if all that is left is a stack of tapes, it will be many days, and maybe weeks, before that tape can be re-loaded onto new equipment and utilized. On the other hand, if the disaster recovery strategy includes the use of replication between sites, then data can be back online in a matter of minutes. This is because the replicated site contains not only a duplicate copy of the data; it also contains a duplicate copy of the hardware and software needed to serve the data.

One of the challenges of using replication for disaster recovery is the amount of bandwidth required to maintain mirrors over WAN links. Some high-end solutions use dedicated lines that are prohibitively expensive beyond the limits of a metropolitan area. A cost-effective solution would utilize public lines and utilize the bandwidth efficiently. The strength of HP OpenView Storage Mirroring in this area is that it can operate efficiently by only replicating the data that has changed. Combined with bandwidth throttling and queuing, this allows software replication to work well over long distances, even with Internet-quality WAN links. Furthermore, Storage Mirroring can easily replicate 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 HP ProLiant Storage Server solutions. Figure 6 shows a sample implementation.

Figure 6. Software replication makes disaster recovery more efficient

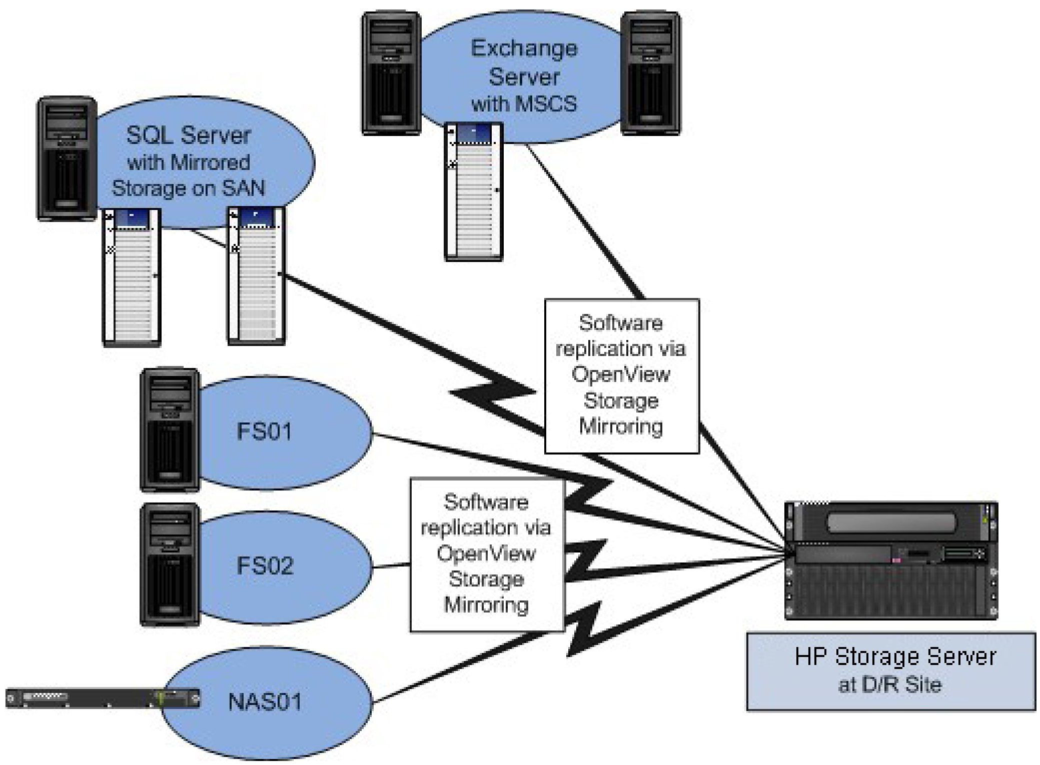

Enhancing backup and restore

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 4:00 PM 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 will not 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.

The combination of HP ProLiant Storage Server and HP OpenView Storage Mirroring can alleviate these issues by staging the data for backup (Figure 7). Instead of backing up directly to tape, the data on existing file servers can be replicated to a central NAS appliance using Storage Mirroring. The mirror on the NAS appliance can be used for immediate restoration if the primary server ever fails. Thus, recovery can be accomplished in minutes, rather than hours or days from tape.

As an added benefit, since HP ProLiant Storage Server is based on Microsoft Windows Storage Server 2003, the Volume Shadow Copy Service (VSS) built into Windows Storage Server may be employed to improve data availability further as follows:

- Generate Shadow Copies periodically on the NAS appliance to provide a level of versioning for the data that is changing. If a file is inadvertently changed or deleted on the primary server, a previous version is easily recovered from the NAS appliance.
- Generate Shadow Copies for backup to tape. A Shadow Copy of a data volume is a point-in-time snapshot of the data. Since this view into the volume is static, it makes an ideal source to create a backup to tape.

Figure 7. Software replication makes backup and restore more efficient

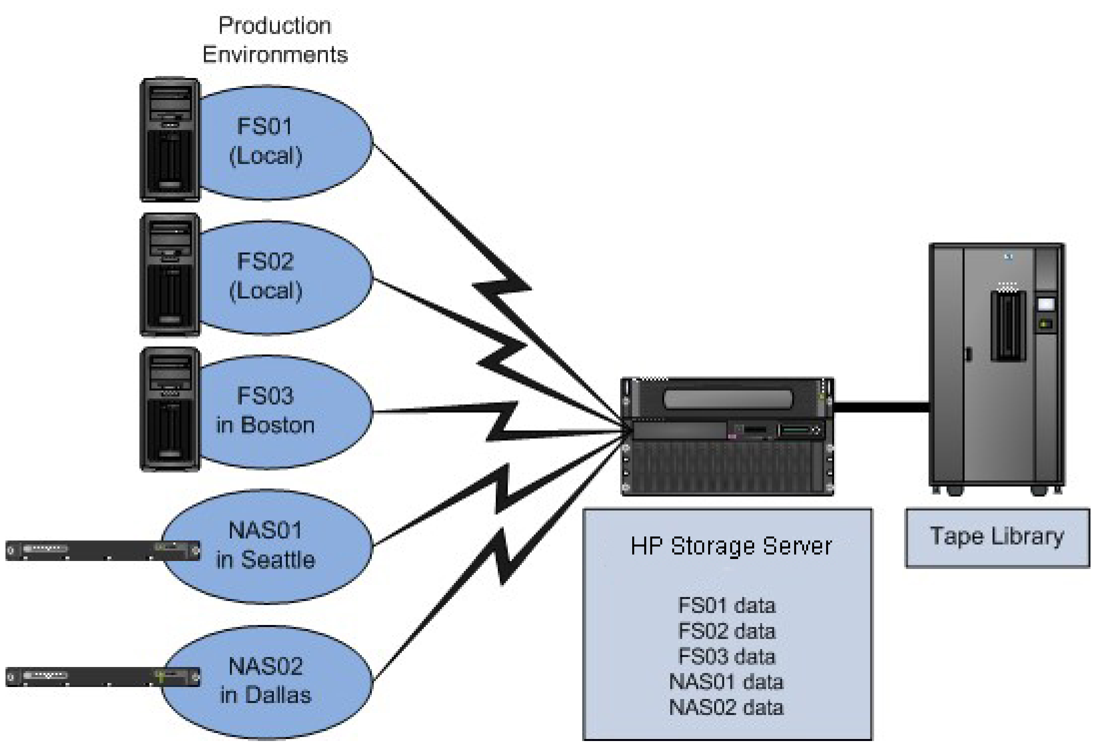

The combination of HP OpenView Storage Mirroring and ProLiant Storage Server for staging backups has many advantages:

- Multiple file servers can be serviced by a single NAS appliance for replication and backup.
- The existing tape backup solution may be leveraged to back up the NAS appliance.
- Backup windows are no longer an issue, since data is backed up from the NAS appliance and not the primary servers.
- Data availability is improved, as lost files can be quickly recovered from the NAS appliance.
- Using VSS, previous versions of files can also be quickly recovered from the NAS appliance.
- Static (point-in-time) copies of the data using VSS can be generated for backup purposes.

Data availability is further improved by deploying HP ProLiant Storage Server appliances as the primary file servers. The VSS built into HP ProLiant Storage Server may be used to generate multiple versions of any given file. The feature allows the end user to retrieve previous versions of a file directly without intervention by the administrator.

This solution is especially well suited for providing centralized backup of data stored in branch or remote offices. Instead of investing on tape backup strategies for each location, a centralized location may be employed to replicate the data from several locations. This allows the tape backup process to be managed from a single location.

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.

Software replication is an attractive approach to data migration for several reasons. First and foremost, the data can be migrated during normal business hours because data can be replicated while clients are online. Second, the migration generates an exact copy of the original, complete with proper file permissions and shares. Software replication using HP OpenView Storage Mirroring allows migrating from one version of Windows to another, so it can be used to migrate from old hardware or old versions of Windows to consolidated HP ProLiant Storage Server solutions.

Figure 8. Migration from legacy file servers to HP ProLiant Storage Server

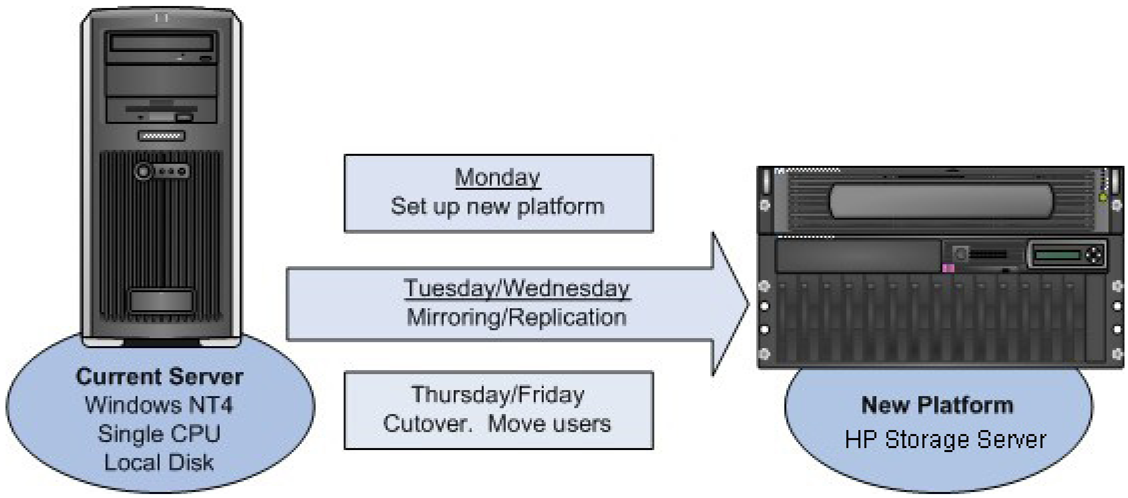

Replicating and migrating from legacy file servers to HP ProLiant Storage Server appliances (as shown in Figure 8) offer some attractive benefits:

- 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 the scheduling and throttling features of HP OpenView Storage Mirroring, 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 6:00 PM on Friday and working all weekend, the new target can be brought online on Monday morning, migrated during the week, and the IT team's weekend is saved.

Because replication allows the data to be moved while the data remains online and accessible to the end users, 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 HP ProLiant Storage Server and HP OpenView Storage Mirroring provides these benefits:

- **Low cost.** Instead of requiring expensive, proprietary hardware replication solutions, Storage Mirroring can replicate data on servers running Windows NT®, Windows 2000, or Windows 2003 to HP ProLiant Storage Server appliances, which offer better scalability and lower costs than traditional storage systems.
- **Flexible replication.** HP OpenView Storage Mirroring 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 Storage Mirroring allows replication updates to be scheduled and throttled, administrators can tailor replication resource usage for their specific environments.
- **High scalability.** HP ProLiant Storage Server appliances are designed to offer high degrees of storage scalability; organizations can choose the right amount of storage to start and easily grow to their storage as their requirements change.
- **Standards-based Solutions.** HP ProLiant Storage Server is based on industry-standard components, such as Microsoft Windows Storage Server 2003, HP ProLiant servers, and HP StorageWorks disk array systems. Because Windows Storage Server 2003 operates as a member of a Windows domain, HP ProLiant Storage Server seamlessly integrates with an existing Microsoft Windows infrastructure.

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 HP OpenView Storage Mirroring replication software and the HP ProLiant Storage Server platform enables cost-effective, easy-to-manage protection for file, print, and application. When taken together, HP ProLiant Storage Server and HP OpenView Storage Mirroring provide superior scalability, flexibility, and TCO compared to hardware- or application-based replication systems.

For more information

HP ProLiant Storage Server: www.hp.com/go/storageservers

HP OpenView Storage Mirroring: http://h18006.www1.hp.com/products/storage/software/sm/index.html

Microsoft Windows Storage Server 2003: http://www.microsoft.com/storage/

5982-8055EN, Rev. 1, 08/2005

invent