Choosing the best architecture for data protection in your SAN

Introduction

The business need for reliable data protection continues to increase. While several options exist for IT organizations to deploy disk-based data protection solutions, tape libraries with removable cartridges still provide the most complete, comprehensive and effective means to protect enterprise wide data.

Today, the majority of enterprise backup and recovery solutions are deployed in a Storage Area Network (SAN). A SAN-based solution provides the most consolidated, cost-effective and scalable choice for most enterprise customers.

SAN backup and recovery solutions have introduced many new business requirements. Organizations have leveraged the SAN to consolidate backup and recovery operations requiring 24x7 operations. As a result, backup and recovery operations are taking place around-the-clock, making reliability and availability of these systems absolutely crucial.

Reliability and performance of high-end tape drives have improved with each generation. Today's tape drives offer both traditional parallel SCSI as well as native fibre channel (FC) connectivity. Many organizations are looking to adopt native FC tape drives going forward with the expectation that this will provide them a more efficient, higher performance and more reliable SAN backup solution.

However, factors affecting the reliability of backup solutions in a SAN extend far beyond a single tape drive. The important question for enterprise IT organizations is not which tape drive interface to deploy, but how to choose the best architecture to meet current and future needs with respect to reliability, functionality and performance – all at a reasonable cost.

This paper introduces the HP StorageWorks Extended Tape Library Architecture as a framework to address key reliability and interoperability concerns. Advanced features and functions made possible by the HP Extended Tape Library Architecture are also discussed. These advanced features satisfy enterprise-wide business needs in the areas of security, high availability, performance, and 24x7 manageability and control.

SAN Solution Reliability

SAN tape library configurations provide consolidated backup and recovery solutions for large enterprise environments consisting of many components including tape libraries, fibre channel switches and host bus adapters, servers, disk storage, as well as backup and application software. SAN backup solutions introduce more complexity into the IT environment than simply attaching tape libraries directly to servers. As a result, SAN configurations often suffer from solution level reliability issues that impair the ability of IT administrators to adequately protecting their business-critical data.

Several factors contribute to the lack of reliability in SAN-based backup environments:

1. Large enterprise storage area networks can experience network conflicts and errors which can cause backup and recovery operations to fail.
2. Due to their complexity, a large administrative effort can be required to not only deploy these solutions, but maintain them as well.
3. When problems do occur, it can be very challenging to collect all the relevant information and isolate issues in a complex enterprise SAN. Troubleshooting issues often requires a significant level of effort and expertise.
4. Frequent changes in an enterprise SAN due to the addition of storage devices or hosts introduce configuration conflicts and incompatibilities, and create instability in the SAN solution.

The following is a discussion of some of these issues in more detail.

Network conflicts and collisions from multiple hosts

In large enterprise environments, backup servers share the SAN with application servers, file and print servers, management stations, and end-user hosts. Tens if not hundreds of these hosts each have the ability to send I/O requests to exposed tape drives and other storage devices in the SAN. This unregulated SAN traffic – be it from operating system calls, diagnostic software, SAN management software, or bus scans – can disrupt, delay, or abort critical backup and restore jobs.

Operating systems have been designed to periodically scan for storage devices, in order to determine the status of known devices or to automatically discover new ones. Bus scans can be automatically executed as part of the operating system services, or manually initiated by a user executing commands at an OS prompt. In addition, many servers today have built-in diagnostic or management applications that periodically issue I/O commands to storage devices to accomplish their intended purpose. When these I/O operations occur during a backup, poor tape drive performance or failed backup jobs can result.

When tape drives are busy streaming data, they are typically operating in an optimized mode in order to maximize data throughput and performance. When unrelated I/O requests come in from a third-party host, the tape drive may need to drop out of streaming mode in order to respond to these outside requests. In large SANs, disruptive requests are increasingly likely, resulting in poor performance of the backup operation.

Additionally, the tape drive may delay its response to unexpected I/O requests during streaming. If multiple requests occur, several of these requests may time out, causing the initiating host to begin an error recovery procedure where even more I/O requests are issued to the tape drive. Recovery procedures vary depending on the OS, application or agent, but in the worst case, such procedures can result in the host sending a target reset to the tape device, or the host overwhelming the tape device with I/O requests that results in a write operation failing or timing out.

As the SAN grows, these problems become more and more likely due to the number of hosts in the environment and the probability of a collision or error occurring.

Network events and errors

Low-level SAN events and errors are an additional source of backup failures and performance degradation.

Change Notification Events

Change notification events occur in a FC fabric network whenever a host or device is added, removed, reset or power-cycled. Event notifications are sent by the FC fabric to all nodes which registered for them. Depending on the device or HBA, these events can generate different results. Some HBAs will respond to change notifications by re-logging into every device it knows about again. Others may abort one or multiple I/O operations in progress in order to assure that their target devices are still valid. Both of these cases can unintentionally disrupt backup operations between a host and a tape drive.

Extra traffic generated by these events can be seen by tape storage systems on the SAN in the form of event notifications and/or port or process logins. These additional FC events can result in the same undesirable consequences as network conflicts and collisions.

Link and Path Failures

Link failures occur at a lower level in the FC protocol and can be due to any of several factors, including loss of synchronization between nodes, problems or marginal behavior with the physical connection, and resource or buffering issues with the FC protocol chips in the storage devices or hosts. Link failures can interrupt I/O requests and cause hosts or devices to either abort operations or initiate error recovery routines.

Path failures typically occur when a connection is broken or unplugged from hosts or devices connected to the SAN. When a path fails during a tape operation, it almost always results in the failure of the associated backup or restore job.

Enterprise disk arrays typically provide the ability to recover from these failures through coordinated logic between host software and the disk array controller, yielding the ability to retry any failed I/O requests using the same or an alternate path. However, tape devices require additional logic to provide a similar capability because of the "stateful" nature of tape. Failed I/O requests can not be simply retried as in the case for disk because the position of the tape in the tape drive is important. A more sophisticated recovery method is needed, requiring more intelligent capability in the tape device.

Configuration challenges

Configuration of data protection solutions in enterprise SAN environments can be a significant effort, due to the size and complexity of these environments.

Each of the components in the SAN needs to be installed and configured properly. One of the most notable includes the host bus adapter in each server connected to the SAN. Example configuration settings include the port settings (e.g. N_port vs. NL_port), maximum I/O transfer length, and settings that determine how the device driver will respond to timed-out I/O requests. Some drivers respond to a timed-out I/O request by sending a target reset to the device. If the device is simply busy and being used by another host, such a reset can result in backup or restore failures. This behavior typically needs to be disabled for to ensure proper HBA operation for backup in a SAN environment.

Enterprise data protection solutions typically employ one or more large tape libraries with many tape drives in a SAN environment with many hosts. The configuration process requires determining the mapping from each tape drive position in the library, through each connection or path in the SAN, to device filenames or handles presented by the operating system, and then finally to the backup application software. For hosts configured with multiple HBAs connected to the SAN, the host may see multiple images of the same tape drives on the network, each representing a different path. These multiple paths need to be reconciled so that truly unique tape drives are used and optimal paths chosen.

Until recently, many of the automated configuration tools provided by backup software applications did not work for SAN configurations. The result was that configuration and verification of a complete environment could be a cumbersome, tedious process representing a large labor-intensive effort. For example, a tape library configured with 16 drives and connected to 20 hosts on the SAN requires 320 mappings which must be determined and configured for all hosts to have access to all tape drives. Today, a few backup applications are now providing tools to automate some of the configuration process in a SAN. These tools may still have limitations and may not work properly for more complex configurations including mixed-drive libraries, hosts with multiple paths to the tape drives, or hosts with multiple HBAs.

After initial configuration, maintaining these mappings is can also be difficult. As the SAN upgrades or expands to include new hosts, switches or storage devices, mappings from backup applications to tape drives can change and become broken.

Troubleshooting

Large complex SANs can be very difficult to troubleshoot when problems do occur.

A significant first step involves verifying that the configuration is still correct. A host that has a tape drive incorrectly configured can interfere with backup operations from other hosts. Due to the number of mappings in large enterprise environments, finding one incorrect mapping amongst these can be a time-consuming and error-prone task.

When a backup operation fails, isolating the root cause often requires correlating detailed diagnostic data from multiple hardware and software components including the tape drives, FC switches, operating system device drivers and backup application software. The troubleshooting process requires a significant level of expertise to analyze and determine root cause of a problem.

Even with appropriate knowledge, troubleshooting problems with a SAN solution can be a time-consuming task. The larger the SAN environment, the more frequently issues can occur, and the more effort it takes to isolate them.

What's been done to address reliability concerns?

Organizations have employed several options to help manage the solution-level reliability issues discussed above. Examples include:

- implementation of zoning in the SAN
- use of SCSI reservations for shared access to tape drives
- segmentation of the configuration into smaller simplified SANs

Each of these options is discussed in more detail below.

Zoning

Zoning logically segments a SAN into groups of hosts and devices which can access each other. Through zoning, a user can mitigate issues seen by tape drives on a larger SAN by limiting the number of hosts which have access to them, thus reducing the number of I/O conflicts and collisions. While this method effectively addresses some of the reliability concerns, it is limited in several areas.

During configuration, zoning represents an additional step to an already complex process. In addition to the determination of mapping from backup software through the SAN to the tape drive, the proper zoning between these devices on a SAN must be identified and configured.

When a problem does occur, verification of the zoning configuration is often required. Zoning configurations can become very complicated, especially where overlapping zones are employed. Understanding which hosts do or do not have access to a specific device can be challenging. This further complicates the already challenging task of isolating and troubleshooting a problem.

Zoning configurations in the SAN are very fragile. As the SAN is upgraded or changed, zoning configurations must be updated. New switches, hosts, or replaced devices can affect the zoning configuration. Updating the zoning configuration can also require that mappings configured in the backup application must be verified or redone.

Reserving tape drives

The SCSI protocol provides a mechanism to manage shared access to a device through the use of RESERVE and RELEASE commands.

Using these commands, a host may reserve a tape drive for use, blocking access from other hosts. While reserved, the tape device refuses tape access requests from other host devices until the host with the reservation issues a release command. This provides some benefit in preventing third-party hosts from interfering with a backup operation. As with zoning, however, drawbacks exist.

One drawback is that reservations do not block all requests from third-party hosts: inquiry, status and some diagnostic commands are not blocked. Responding to blocked and unblocked requests can still interfere with performance of the tape drive as well as lead to error recovery conditions as mentioned earlier.

Another drawback involves determining which hosts or applications *should* get a reservation. In some cases reservations are implemented in the tape device driver providing reservations to all applications. Management or diagnostic applications can issue and get reservations which could interfere with the intended backup operations. When a backup application does not get a reservation, for whatever reason, a backup failure can result.

When backup jobs fail, a tape device can be left in a reserved state. These "orphaned" tape drives are inaccessible by other hosts or applications until released – an operation that can require manual intervention by a system administrator. Concerns relating to orphaned devices and "lost" reservations make the system appear to be less reliable, and require more administrative effort to maintain availability.

Tape drives understand little about the environment or configuration they "live" in, or the applications and policies that dictate drive behavior. Individual tape drives do not have the information or knowledge required to properly manage access on their own.

Simplified SANs

One way to avoid the pitfalls presented by large SAN configurations is to keep the SAN itself smaller. Many administrators have opted to configure their data protection solutions using several small SANs to connect their tape libraries.

This is done by breaking the set of tape drives within a library into smaller pools and connecting these pools of drives to a subset of the potential hosts in an isolated small SAN environment. This approach keeps the configurations simple, with less chance for network errors or conflicts to interfere with backup operations.

While successfully minimizing reliability issues, this type of configuration greatly limits the potential benefit of leveraging a shared pool of tape drives. In small, isolated SANs, performance bottlenecks are introduced and backup capacity is not fully utilized. Managing separate SAN "islands" also becomes difficult, as few tools can span the entire environment.

While several of options exist for mitigating or minimizing solution-level reliability issues in enterprise SANs, each option comes with significant drawbacks and limitations. Furthermore, none of these methods completely addresses all the problems discussed earlier, leaving significant reliability and performance issues unaddressed.

Evolution of network storage

What, then, is the best approach for ensuring reliability, robustness, and ease-of-use in a complex SAN? To better answer this questions, consider how disk storage has evolved to address some of these same issues.

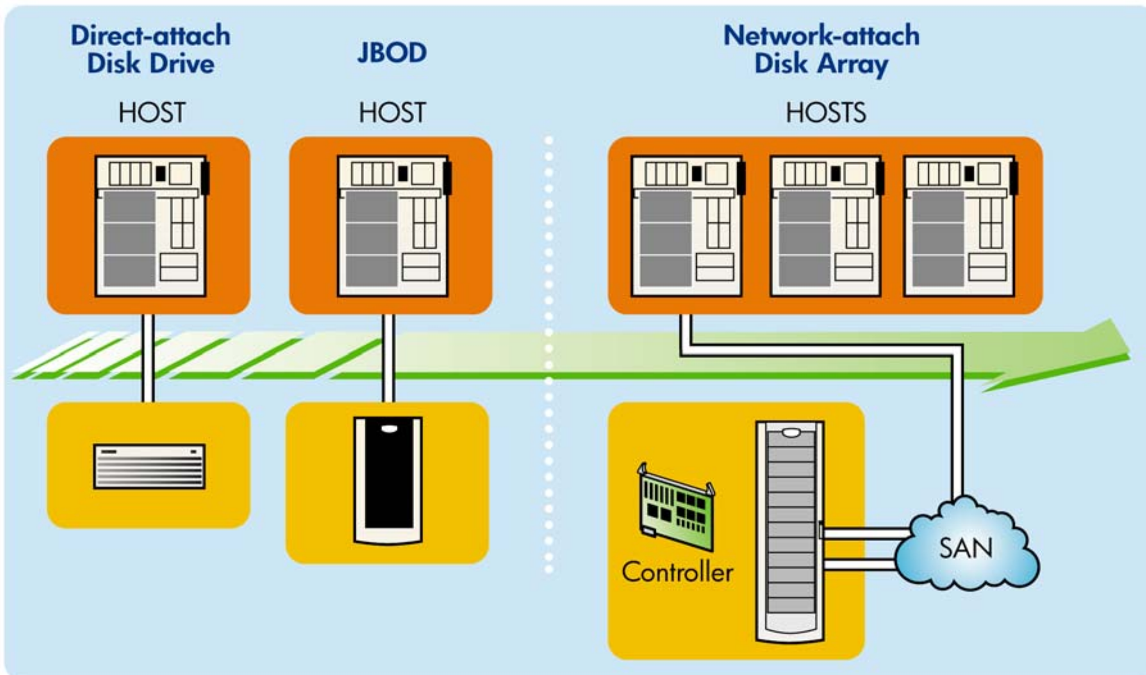

Disk storage began as dedicated direct-attached devices to hosts. As applications and networks grew, the need for more consolidated and scaleable storage also increased. Disk storage expanded to include JBOD, or 'just-a-bunch-of-disks' configurations, to meet the need for consolidation and scalability.

Over time, external disk arrays with integrated controllers evolved to include attributes such as fault tolerance and high-availability, as well as becoming significantly more network capable. Today, disk array systems exhibit a number of attributes making them the optimal choice for providing shared disk storage in a SAN. These disk arrays provide consolidation, scalability, share-ability, security, reliability and redundancy, primarily through an array controller that sits between the JBOD and storage network. Few, if any, administrators would consider utilizing a JBOD in a SAN environment since JBODs lack many of these attributes.

Today's tape libraries typically include slightly more network capability than a JBOD, but significantly *less* network capability than a SAN-attached disk array. Tape libraries have lacked the security, network reliability, share-ability, and manageability required for enterprise SAN environments.

To meet the increasing demands and requirements of their business, end-users require more functionality and value from their investment in data protection systems and solutions. The HP StorageWorks Extended Tape Library Architecture intelligently addresses these demands and requirements.

A controller-based architecture for tape libraries

The HP StorageWorks Extended Tape Library Architecture introduces next generation architecture for tape libraries in enterprise SAN environments. The architecture is aimed at meeting high reliability and interoperability requirements by incorporating intelligence and advanced capabilities into the tape storage subsystem. In addition, the architecture enables comprehensive manageability and monitoring of the tape storage subsystem.

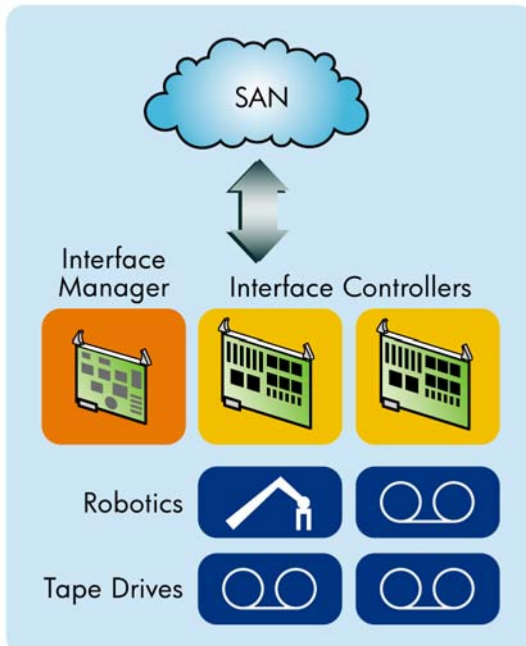

The architecture adds a layer of intelligence between the tape drives and the SAN by incorporating a set of Interface Controllers similar to the array controllers in a disk array. These controllers connect externally to the SAN, providing the network interface for the tape library, and internally to either fibre channel or SCSI tape drives. As a result, the controllers can manage shared network access and provide much of the reliability and security required.

The brains of the architecture lie in the Interface Manager, which provides a central point of knowledge of the tape library and the SAN environment. The Interface Manager extends the intelligent management capability of the architecture and enables the storage subsystem to become truly self- and network-aware. Rich, remote, and integrated device management is also made possible by the Interface Manager.

The HP StorageWorks Extended Tape Library Architecture is designed to meet current and future enterprise needs by focusing specifically on the following areas:

- Reliability
- Interoperability
- Advanced Functionality
- Manageability

Each area is discussed in detail below, illustrating how the architecture addresses these needs.

Reliability

The HP Extended Tape Library Architecture improves SAN solution reliability through device-level access control, caching, event reporting, and simplified, automated configuration.

Device-level access control allows administrators to specify which hosts are allowed to access tape drives, and enables a controller to block I/O requests from third-party "rogue hosts" which may interfere with an active backup operation.

Through caching, a controller is able to respond on behalf of a tape drive while the drive is busy during a backup or restore operation. Caching can improve drive performance in addition to addressing conflicting I/O requests from multiple hosts.

Intelligent handling of network events further prevents disruption of backup or restore operations. For example, a target reset command sent by a third-party host can disrupt active backup or restore operations. With more knowledge of the storage network environment and the tape library, a controller is able to modify or limit the affect of potentially disruptive SAN events.

An embedded controller also provides increased network reliability through SAN-smart event reporting. Controllers enable monitoring, tracking and detection of errors affecting SAN backup and restore operations, including fabric events, aborted I/Os, and reservation conflicts. The controller also actively identifies configuration issues and conflicts in the SAN. Examples include recognizing when two hosts collide through a reservation conflict or attempt to write to the same tape drive at the same time.

The Interface Manager provides simplified, automated and robust configuration of the tape library in a SAN environment by intelligently and consistently configuring the system based on knowledge of the tape library and the SAN. The Interface Manager also adapts to changes within the tape library and SAN to automatically correct and maintain consistency.

The Interface Manager contains on-board Flash memory to provide a persistent history of the tape library and storage network health. This persistent history can be automatically analyzed to verify and detect configuration issues.

Interoperability

The HP Extended architecture represents an evolutionary step that highly leverages existing fibre channel routers for tape libraries. Years of engineering and testing provide a mature fibre channel interface that boasts the broadest heterogeneous SAN interoperability in the industry.

This is in contrast to native FC tape drives, which have relatively new FC interfaces and will be evolving their interoperability over time.

With the HP Extended architecture, interoperability is independent of the tape drive itself, allowing the use of LTO1, LTO2, SDLT220, or SDLT320 tape drives behind the Interface Controllers. The HP Extended architecture enables users to upgrade their tape libraries with new generations of tape drives, while maintaining stability and interoperability with their SAN environments. Users can also retain their older drives as SAN interface technologies change and new features or capabilities become available.

Examples of interface changes and upgrades include moving from loop to fabric, moving from1 Gb/s to 2 Gb/s fibre channel, moving from 2Gb/s to 4 Gb/s or 10 Gb/s in the future, and moving from fibre channel to iSCSI SANs. A controller-based architecture enables these types of changes while leveraging a user's existing investment in tape libraries and drives.

Advanced functionality provided by the HP Extended Tape Library Architecture allows users to seamlessly expand or extend their SANs without negatively impacting current solutions.

Advanced Functionality

Enterprise IT organizations make significant investments in the area of data protection. As business demands continue to increase, organizations need to realize more value from their data protection solutions.

The controller-based HP Extended Tape Library Architecture is able to deliver increased value through advanced features and functionality, now and in the future. Some of these advanced features and functions are described below.

Security

Device-level access control is a key element for providing reliability in an enterprise SAN, preventing non-backup hosts and applications from disrupting critical backup and restore jobs. Access is typically controlled by specifying a set of authorized host world-wide names (WWNs) for each tape drive or library. An embedded controller then verifies each host login or I/O request against the authorized WWNs for the target tape drive. Configuration of access controls for a large tape library and SAN environment can be tedious and error-prone, and proper management tools are critical. By providing a single point of configuration and control for the entire library subsystem, the Interface Manager enables easier, more predictable management.

Access controls also let administrators protect data in a tape library from unauthorized access in the SAN. This becomes more of a business requirement as SANs grow to be large and may extend outside of the data center. Physical access becomes more difficult to control. Unauthorized hosts can obtain a physical connection and have access to sensitive data stored in unprotected tape libraries.

The Interface Manager and embedded Interface Controllers also provide a platform for delivering future data security capabilities, including encryption. Data encryption provides an even higher level of security for protecting access to data stored in a tape library. Encryption can require significant memory and processing resources, or dedicated hardware in order to meet performance requirements. Interface controllers can provide the additional resources or hardware needed for this functionality, in a way that stand-alone tape drive can not.

Another security-enhancing mechanism for enterprise SANs is library partitioning. Partitioning allows users to divide large tape libraries into two or more logical storage systems. Each library partition can be given separate access controls or could even be connected to separate SANs such that media or tape drives cannot be accessed between them.

Administrators can segregate different departments or business sensitive data pools to provide better protection and security. In other cases, a service provider may need to separate access from different end-user customers being hosted in their environment. Again, the HP Extended Tape Library Architecture makes this type of functionality possible, without requiring additional hardware purchasing, installation, or configuration.

Availability

Availability is another important requirement in today's enterprise SAN environments. Several availability features commonly provided by high-end disk arrays – including path failover, mirroring, and device sparing -- can be delivered through a controller-based architecture for tape libraries.

Path failover, as the term implies, provides the ability to fail-over to an alternate data path between a host and the tape library in a SAN. Because of the sequential-access nature of tape, failed or timed-out operations can not simply be retried, without knowing the current tape position. Additional logic and protocol-level interactions are required in order to provide a true fail-over or failure-recovery mechanism for tape. A controller-based architecture is able to provide the additional intelligence required to deliver this capability.

A controller-based architecture also enables tape mirroring, which can split a single backup stream from the host into multiple streams sent to the tape drives. If a tape drive or media failure occurs, a mirrored backup can still complete successfully using one of the remaining tape drives without disrupting the backup operation.

Sparing is the ability to automatically replace a failed or offline tape drive with a new drive, or "hot spare." If the tape drive to be replaced sits behind an intelligent controller, the controllers can map a spare drive into the same physical and logical location as the old drive. This change is then transparent to the backup application and does not require any reconfiguration of the environment.

Performance

High performance tape drives are doubling their data transfer (or "streaming") rate with each generation. While these tape drives are capable of adjusting their streaming rate to match the host, there is a minimum performance threshold a host must be able to meet in order for the tape drive to maintain a steaming mode of operation.

A controller-based architecture can provide more intelligent data gathering and reporting of performance metrics in a backup solution, alerting users when tape streaming is not maintained. Metrics such as host-to-drive throughput, compressed data rates, port utilization, and the frequency of start-stop operations by the tape drive are all relevant to assessing backup performance. Additionally, these metrics need to be understood within the context of a particular backup or restore operation, host-to-tape drive pairing, and specific network/data path. The ability to report these metrics in a useful context requires a significant level of analysis and intelligence that can be best provided through a controller-based architecture.

Third-party copy or serverless backup – the ability move data directly from disk to tape -- has been a long-awaited promise of SANs since their emergence. However, obtaining the necessary integration between storage hardware and backup applications – a requirement to enable full serverless backup solutions – remains a challenge.

Solutions available today tend to target very specific environments and configurations. Each targeted solution offers differing levels of integration with disk arrays, business applications, and tape libraries, resulting in varied levels of capability and value. As these solutions continue to improve and expand, serverless backup solutions will become more valuable and prevalent in enterprise environments.

The ability to move data directly from disk to tape requires dedicated processing power and memory resources. Integrated controllers are well suited to provide these additional resources – something stand-alone or commoditized tape drives can not offer. Because such controllers are integral to the tape library subsystem, they automatically scale as backup capacity increases.

Disk-based backup solutions have recently emerged in the market. One potential benefit of these solutions is the ability to address the performance weaknesses of tape. Tape positioning issues do not exist for disk systems, making partial or single-file restore operations much faster. Disk-based backup systems also do not require a minimum threshold for streaming write operations offering more reliability and potentially better performance at lower data rates.

A controller-based architecture allows for the integration of disk into the tape library system providing users with the best performance of both storage media. Integrating disk tightly into the architecture instead of as a separate component on the SAN provides users with a more turn-key, reliable and manageable solution. A combined disk and tape system enables a more seamless integration with backup application software and operations, providing users more flexibility and policy-based management of their data protection solution.

Manageability

By providing a single point of configuration and control, the Interface Manager enables complete remote manageability for the whole tape library subsystem. HP StorageWorks Command View for tape libraries, a key software component of the HP Extended Tape Library Architecture, communicates out-of-band with the Interface Manager to provide rich, web-based library manageability from any location. Through Command View software, users can:

- Select from a list of multiple supported libraries to manage
- View asset information for the library system
- View health summary for all components of the library system: drives, robotics, Interface Controllers, and Interface Manager
- View library inventory: slots, media, and drives
- Search for media by barcode label
- View details on drives, robotics, Interface Controllers, and Interface Manager, including WWNs, speeds, and connection types
- Configure drives to be used with the Direct Backup Engine
- Configure hosts/HBA-to-drive mapping and masking settings with Secure Manager
- Automatically discover hosts and HBAs on the SAN or add your own for access control and security purposes
- Upgrade firmware on the Interface Controllers and Interface Manager cards
- Configure IP network parameters for the management station and Interface Manager cards in each library
- Configure SNMP trap/event destinations to receive alerts asynchronously
- View persistent event logs for the library system

Many of these features and functions are also available through a telnet or serial interface on the Interface Manager itself.

The HP Extended Tape Library Architecture enables not only remote manageability, but also integration with leading management applications such as HP OpenView Storage Area Manager, HP Insight Manager for Proliant, HP OpenView Network Node Manager, and others. Future management technologies, such as SMI-S, also become possible through the HP Extended architecture.

Because the Interface Manager and Command View software communicate out-of-band, over the LAN, additional SAN traffic is not generated and library management does not interrupt backup and restore operations.

What about native fibre channel tape drives?

Several high performance tape drives offer native-FC interfaces instead of parallel SCSI. Many IT organizations may consider deploying these tape drives in their enterprise SAN environments. However, native-FC tape drives connected directly to the SAN do not provide many of the benefits offered by a controller-based architecture.

It is important to note that a controller-based architecture *can* utilize native-FC tape drives to provide additional value and functionality.

The HP Extended Tape Library Architecture is able to incorporate native-FC tape drives by providing an intelligence layer of Interface Controllers between the drives and the enterprise SAN.

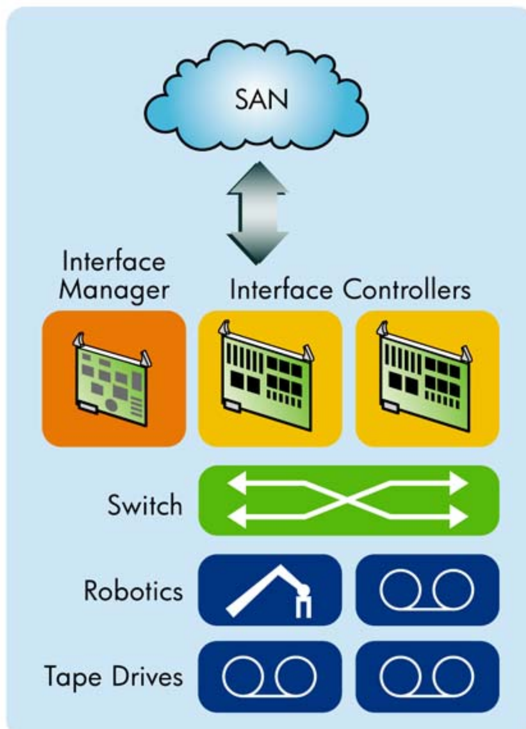

This architecture provides all the same benefits of reliability, interoperability, manageability, and advanced functionality described in previous sections. Native-FC tape libraries *without* controllers do not offer these same benefits.

The HP Extended architecture also provides the option of incorporating switching technology between the interface controllers and the tape drives to enable additional functionality. With a switched-back-end architecture, advanced functions such as mirroring, partitioning, and failover are able to span the entire tape library without limitations due to controller boundaries. This architecture also provides the ability to fail-over *between* controllers for a higher degree of availability.

A switched-back-end architecture also enables the development of advanced data protection techniques including the ability to integrate disk with tape into the data protection subsystem as mentioned in the previous section.

Doesn't a controller introduce a single point of failure?

Some users may be concerned that adding a controller will introduce a single point of failure into the system, and therefore lower the overall reliability.

The reality is that the most important factors affecting reliability are not hardware failures. On the contrary, electronics and controllers are typically one of the lowest failure-rate components in a tape library. Backup and restore failures most frequently occur due to configuration and network errors in the SAN. These reliability issues are discussed in more detail earlier in this paper.

Because the HP Extended architecture is designed specifically to address SAN interoperability and reliability problems, it provides considerably better overall solution reliability, not worse. In addition, the architecture provides capabilities to isolate and resolve problems more quickly in complex environments, resulting in increased uptime and availability.

What about performance?

Performance of backup operations depends primarily on the throughput or bandwidth from the source of the data, through a SAN infrastructure path, all the way to the tape drive. The number of components in the SAN path alone, is seldom is a factor in performance issues.

A good analogy is water flowing through a pipeline. The amount of water or rate of water coming out the end is similar to the performance of a backup operation. An additional section of pipe has little effect on the rate of water flowing through it. However, a small diameter section at some point in the pipeline can introduce a bottleneck, restricting the flow rate of water through it.

HP's fibre channel Interface Controllers are specified to match the generation of high-end tape drives they are connected to. These controllers are specifically designed to prevent the introduction of bottlenecks to achieve optimal performance during backup operations.

Controllers are able to optimize performance through the immediate reporting of status on write operations. For example, a controller can acknowledge the completion of a write operation as soon the last byte of data is received, allowing the host to send the next write command while the data is being written to the tape drive. Logic between the host and controller (and tape drive) assures that all data is written to tape before the last I/O request for a backup job completes successfully. This same mechanism is used by the majority of tape drives in the industry today, and a similar capability is an important part of the cache management schemes used by disk array controllers to improve performance.

As discussed previously, controllers can improved the overall performance of a tape drive by blocking disruptive I/O requests that would drop the tape drive out of streaming mode. In addition, a controller-based architecture can offer advanced functionality which can improve performance in a complex SAN environment.

Isn't a controller-based architecture expensive?

HP's Extended Tape Library Architecture is designed to be cost-effective. Controller hardware represents a small fraction of the cost for an enterprise tape library solution.

Furthermore, the architecture needs to be compared with the alternative solution connecting a non-controller-based native-FC tape library to a SAN. Typically, for a large tape library a user will want to aggregate the number of fibre channel tape drives per link when connecting into the SAN. This is due to the practicality and cost of maintaining full connectivity from all tape drives in a library to all hosts in a large enterprise SAN. Aggregation is provided by controllers. A non-controller based tape library requires additional fibre channel switches in order to provide this aggregation. This comparison can frequently yield a neutral cost difference between these alternatives.

Through self-management and advanced diagnostic capabilities, the architecture significantly reduces the administrative effort to deploy and maintain the solution. This can be a significant savings in terms of effort and administration.

Finally, a controller-based architecture provides many advanced features and capabilities that result in significantly higher value and return on investment for the data protection solution.

Summary

This paper has highlighted several important criteria for selecting a data protection solution for an enterprise SAN environment. Of these criteria, the most critical involve how well the architecture delivers solution-level reliability and interoperability in large complex network configurations. Driven by the end-user's demanding business needs, today's data protection solutions must provide greater value and a better return on investment than ever before.

A controller-based architecture for tape libraries is designed to best meet these needs in enterprise SANs. The HP StorageWorks Extended Tape Library Architecture maximizes return on IT investment by incorporating greater levels of intelligence into the tape library in order to mitigate network conflicts and security issues. The architecture also serves as a platform to provide advanced functionality and deliver greater overall value to enterprise tape library users.

The HP StorageWorks Extended Tape Library Architecture defines the next-generation of tape libraries in enterprise SAN environments, taking tape libraries to the next level by making them fully network-capable. As in modern disk arrays, the architecture incorporates a layer of intelligence between the tape library and the SAN to offer increased levels of data protection and enable IT organizations to transform their SANs into Adaptive Enterprises that manage storage as a utility.

For more information

www.hp.com/go/automated
www.hp.com/go/storage
Optimizing Storage for an Adaptive Enterprise, HP, 2003
ENSAextended: Storage for an Adaptive Enterprise, HP, 2003
HP ENSAextended technical overview, HP, 2003

5982-1838EN, 9/2003