Storage Basics:

A Guide to the World of Storage Technology

an internet.com Storage eBook

contents

Storage Basics: A Guide to the World of Storage Technology

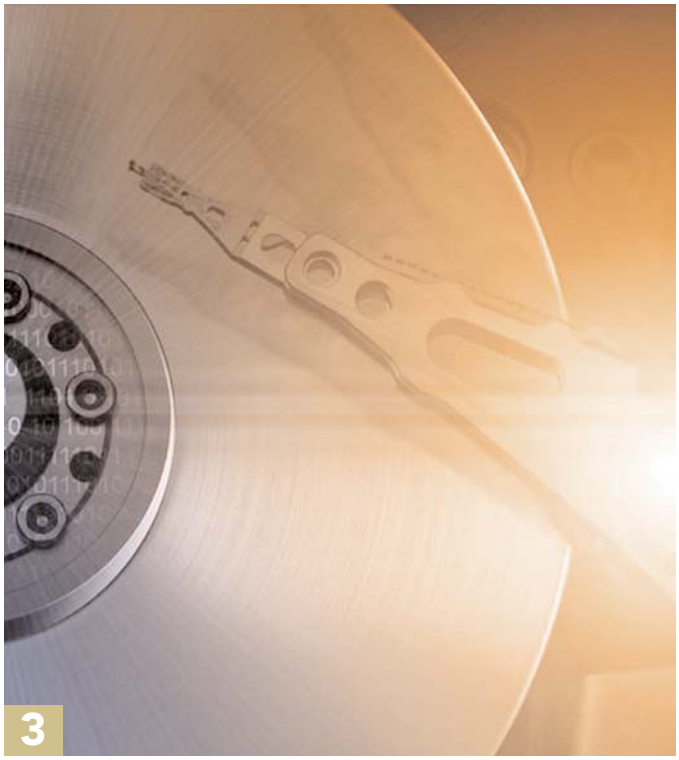

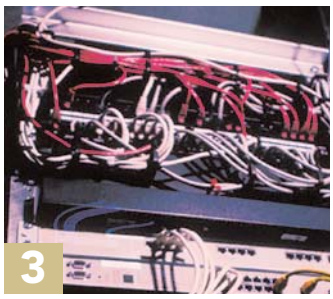

This content was adapted from EarthWeb's Enterprise Storage Forum Web site. Contributors: Dan Muse, Paul Shread, Drew Robb, Mike Harwood, and Henry Newman.

Solving Storage for Your SMB, An Internet.com Storage eBook.

Storage Basics:
A Guide to the World of Storage Technology

Introduction by Michael Pastore

Sales of storage products reached $3.9 billion in the fourth quarter of 2005, according to IDC, the best quarter for the storage market since IDC began tracking it in 2001. You can expect the numbers to keep rising.

Regardless of industry, size, or age, enterprises are awash in more data than ever before. Fewer business processes rely on paper, and the file cabinets that once filled offices for generations are now located on racks in the server room. Federal regulations require that data be stored, protected, and retrievable for a certain amount of time, and specific industry regulations add to the burden.

Jupiterimages

Storage is one of the most basic operations performed by computers, yet it continues to evolve. In the days of mainframes, data was stored physically separate from the actual processing unit, but was still only accessible through the processing units. As PC-based servers became more commonplace, storage devices went "inside the box" or in external boxes that were connected directly to the system. Each of these approaches was valid in its time, but as our need to store increasing volumes of data and our need to make it more accessible grew, other alternatives were needed.

Network storage is a generic term used to describe network-based data storage, but there are many technologies within it. Direct Attached Storage (DAS) is a storage device that is directly attached to a host system. The simplest example of DAS is the internal hard drive of a server computer, though storage devices housed in an external box come under this banner as well. DAS is still, by far, the most common method of storing data for computer systems.

Network Attached Storage, or NAS, is a data storage mechanism that uses special devices connected directly to the network media. These devices are assigned an IP address and can then be accessed by clients via a server that acts as a gateway to the data, or in some cases allows the device to be accessed directly by the clients without an intermediary.

A Storage Area Network (SAN) is a network of storage devices that are connected to each other and to a server, or cluster of servers, which act as an access point to the SAN. In some configurations a SAN is also connected to the network. SANs use special switches as a mechanism to connect the devices. These switches, which look a lot like a normal Ethernet networking switch, act as the connectivity point

Why is it important to learn the basics of storage technology? As mentioned earlier, how enterprises store data is becoming more than a best practice, it's becoming a legal matter as well, and the penalties for individuals and corporations can be severe.

Storage is also a growing area within IT, which means employment opportunities exist now, and should exist for some time. According to a one study, fewer than 25 percent of either Unix-/Linux- or Windows-based IT organizations had their own storage management team at the end of 2004. By the end of 2006, however, that number is expected to soar above 75 percent. ■

What Makes a Storage Server a Storage Server?

By Drew Robb

Ask people what a storage server is, and you can expect to hear a variety of answers. Some will say it is a regular server with added features, a few describe it as a stripped-down box dedicated to a specialized function, and still others believe the term refers only to a network attached storage (NAS) box.

Not Your Average Server

The typical server is configured to perform multiple functions. It operates as a file, print, application database, Web, or miscellaneous server. As such, it must have fast chips, more RAM, and plenty of internal disk space to cope with whatever end users decide to do with it.

Not so with a storage server. It is designed for a specific purpose, and thus configured differently. It may come with a little extra storage or a great deal.

"A general-purpose server typically has five or less disks inside," says Graham Lovell, senior director x64 servers at Sun Microsystems. "A storage server, on the other hand, has at least six, and more, usually 12 to 24 disks."

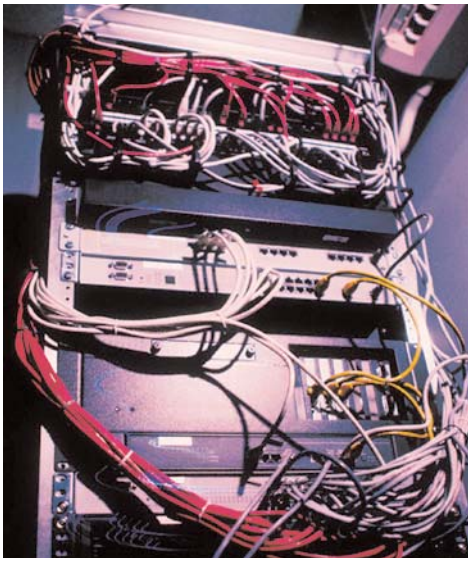

Jupiterimages

Storage servers are normally individual units. Sometimes they are built into a 4U rackmount. Alternatively, they can consist of two boxes - a storage unit and a server located nearby. Both boxes can then be placed side-by-side in a rack. The Sun StorEdge 3120 storage unit and SunFire X4100 server, for example, can be combined into a storage server and placed in a rack.

Apart from extra disks, what else is different about storage servers? In many cases, they come with a host of specialized services. This can include storage management software, extra hardware for higher resilience, a range of RAID configurations and extra network connections to enable more users to be desktops to be connected to it.

Just a NAS Box?

Interestingly, some vendors define storage servers purely in terms of NAS. A NAS appliance (also known as a NAS filer) generally has a slimmed-down OS and file system, and only processes I/O requests by the main file-sharing protocols. The big advantage of the NAS architecture is that it enables storage to be rapidly added by plugging the appliance into a network hub or switch.

"As far as HP is concerned, a storage server is NAS," says Jim Hankins, product marketing manager for HP's NAS division. "In essence, it is a dedicated file and print server."

HP has a number of its ProLiant models available as general-purpose servers or storage servers/NAS filer - each has the same basic hardware configuration. If licensed as a storage server, the user may not run general-purpose applications on that server. If the same ProLiant server is being used as a regular server, however, applications can be run on it.

In addition, HP's NAS-based storage servers have extra functionality built into the operating system - storage-specific management tools, "quota-ing" features, storage reporting capabilities, and a Web-based user interface that makes it easier to configure file and print.

So is NAS really just a storage server? The answer varies, depending on whom you ask. But it appears there is very little difference between them. NAS, it turns out, isn't really storage networking. Actual network-attached storage would be storage attached to a

storage-area network (SAN). NAS, on the other hand, is just a specialized server attached to a local-area network. All it does is make its files available to users and applications connected to that NAS box - much the same as a storage server.

"NAS is a marketing term," says Dan Tanner, an analyst at storage consulting firm ProgresSmart. "NAS is really nothing more than a file server, but specialized or adapted to the single purpose of serving files."

And what a marketing campaign it has been. From nowhere in the mid-1990s, Gartner projects the NAS market will exceed $2 billion by 2008, with an annual growth rate of 9 percent. And those numbers don't take into account a new NAS flavor called the NAS gateway. These gateways act as a file-serving portal into a SAN: There are disk arrays in a Fibre Channel SAN that have a storage server on the perimeter acting as a NAS gateway. This is a one way to marry up NAS and SAN assets.

"There are two flavors of storage servers," says Hankins, "NAS appliances that have the disk storage in the appliance, and NAS gateways."

What's Missing?

While some vendors use the same box as a plain vanilla server, others use a scaled-down version that is adequate for file serving. Steve Duplessie, senior analyst at Enterprise Strategy Group, defines a storage server as an optimized appliance designed to feed information, via a network, to a user or an application. As such, it is not typically compute heavy, but it has been designed from the ground up to provide specific I/O capabilities along with data protection capabilities.

A regular server has to be generic, it doesn't know what kind of load demands it will have - gaming is much different than running a database, for example. A storage server, such as a NAS box, is a contained appliance that does one thing really well, like file serving.

What does a "regular" server have that a storage server doesn't? According to Duplessie, it typically has more processing power, more RAM, and a more generic I/O structure and file system. As a result, most storage servers perform at 50 percent of the performance of a regular server for the same function, he says. This trend toward specialized computing elements is far from new. TCP/IP routing, for example, was a function

Storage Definitions

by Drew Robb

The world of storage can be forbidding to a novice. Even veteran IT personnel may be put off by the sheer volume of new terminology and alphabet soup that has evolved. Let's sample some basic terms:

Direct Attached Storage (DAS): The server stores data on disks that are in the same box. Redundant Array of Independent Disks (RAID) is used heavily in this approach.

Storage Area Network (SAN): A collection of computers and devices are connected over a high-speed network and are dedicated to the task of storing and protecting data. Instead of storing data locally, each server sends data across the network to a shared pool of storage.

Disk Array: A large array of disks in one box, it is often used as part of a SAN to store data for multiple servers. These servers typically connect to the disk array using Fibre Channel.

Fibre Channel (FC): Optical fiber cables transmit data at high speed in a SAN. Fibre Channel is the transport protocol used for this purpose.

Network-Attached Storage (NAS): NAS separates data from applications by storing data on filers attached to the LAN. Filers can share files across multiple applications, platforms, and operating systems.

Internet Small Computer Systems Interface (iSCSI): This standard enables storage and retrieval at high speed (1 GB/second or higher) over regular IP networks.

- **Drew Robb**, Enterprise Storage Forum

> 30 percent of 288 storage professionals surveyed said their companies' security policies did not include storage systems. -- Enterprise Strategy Group

that every operating system ran - until Cisco came out with a dedicated box that did it far better than hosting it on a general-purpose server.

"Any time you can optimize a function, it will be better [on a specialized box] than if executed on general-purpose gear," says Duplessie.

Dan Tanner, an analyst with the storage consulting firm ProgresSmart, agrees with Duplessie's view that a storage server is a specialized server or appliance.

"The server OS is cut down to address purely print server or file server functions, and often contains specially tuned or enhanced code," says Tanner. "Before NAS came along, though, Microsoft said you could use a regular server for file serving."

But using a vanilla server for file serving could lead to problems. Administering a general-purpose server is more complex. Further, someone might be tempted to use the server for multiple functions. Dedicated storage servers, therefore, have become the norm.

Not surprisingly, Microsoft introduced Windows Storage Server 2003 to distinguish it from general servers running the Windows 200x operating system. Windows Storage Server 2003 is a dedicated file and print server based on Windows Server 2003 and tailored to networked storage. It supports file serving and backup and replication of stored data. It can also be used to consolidate multiple file servers into a single box.

Storage Servers vs. Disk Arrays

Just as there is some confusion between ordinary servers and storage servers, there is also sometimes a misunderstanding between storage servers and disk arrays. Exactly where does one end and the other begin? A storage server can have as many as 24 disks - enough to qualify as an array. Disk arrays, however, can have hundreds of disks. So where do you draw the line?

Storage Server Differentiators

- **Lots of disks (12-24)**
- **A standalone unit**
- **Preinstalled software apps to manage the data or storage-specific peripherals**
- **Usually less powerful than its pre-installed counterparts**

"A storage server is usually standalone and not connected to other servers," says Lovell. "Multiple servers, however, typically connect to a disk array."

Disk arrays, too, often connect to a server that could be styled a storage server. The storage server is the intelligence that goes in front of the array. In this arrangement, the server can manage several tiers of storage. It can even arrange the replication of data from one tier to another.

"A storage server serves the storage, and the disk array is the storage," says Tanner. "Using a storage server lets you use multiple or different arrays."

Duplessie further separates the two terms.

"A storage server typically speaks to files and talks to people or applications over Ethernet," says Duplessie. "A disk array is a low-level block device that only speaks to an operating system." ■

Storage Strategies Made Simple

By Drew Robb

Storage is an immense and complex universe. Once you enter, your mind is soon swimming in strange, even alien concepts. Therefore, it is best to stick to what you know and keep it very simple - especially at the start.

One obvious way to avoid complexity is to use the services of a storage service provider. These are firms that lease storage from their own data centers and other services. Colorado Software Architects, for example, offers 1Disk.com. Sun, Arsenal Digital, and Iron Mountain are among the companies with similar services.

The advantage of a storage provider is that the vendor provides a variety of storage options for a fixed cost. This is a handy way to add storage capacity or meet regulatory compliance/archiving requirements without having to build new infrastructure.

Jupiterimages

Of course, simplicity can be taken to extremes (i.e., attempting to pass the entire storage burden to an external source or keeping everything stored on the same old servers using bigger and better disks). Such a strategy eventually runs into a wall; there is so much data stored on so many servers that it becomes impossible to manage.

Beyond DAS, then, where should the rookie storage guy go to ease his woes? Initially, at least, it might be smart to start with NAS and avoid SANs. At its core, a NAS filer is simply a specialized type of server that connects to the network. Storage is rapidly added by plugging the appliance into a network hub or switch. The likelihood is that the server administrator will run into very little that is new to him by buying a NAS box. Lower-end models that are relatively easy to use are available from Network Appliance, Snap Appliance (now owned by Adaptec), and HP.

The drawback of NAS is that filers and servers share the same LAN. As a result, network performance may eventually be affected. When that juncture is reached, it may be remedied by upgrading the LAN and adding higher-grade NAS equipment. A more long-term solution would be to roll out the first SAN.

Simple SANman Says

Undoubtedly, the land of the SAN can be forbidding. Continuing with our theme of simplicity, the transition to a SAN can be made smoother by beginning with rapidly maturing iSCSI technology. iSCSI allows the establishment of a SAN over an IP network. Thus, the IT department does not need to learn new protocols or add new skill sets to create a SAN. This also has the advantage of being much less-expensive than an FC SAN.

Super-Size It

iSCSI is especially appropriate for companies with IP backbones capable of handling gigabit traffic. While the technology is improving rapidly, it doesn't offer the same speed or capacities as a heavy-duty FC SAN. Similarly, SANs offer higher speeds and throughput than NAS systems. To do this, they offload data traffic to a separate network for storage devices.

On the negative side of the ledger, however, SANs may have difficulty supporting multiple operating systems and platforms. In addition, some users complain about being unable to integrate SAN solutions from different vendors.

Choose Wisely

The basic strategy for storage is to try to stick with the familiar. NAS and iSCSI are good starting points for competent IT departments already familiar with IP networking. FC SANs, on the other hand, should probably be avoided unless you have very large capacity and require the highest possible performance.

If so, it is best to recruit a dedicated storage team to wrestle this beast and bend it to your corporate will.

Although the cost and complexity are greater in the short term, the potential long-range payoff is greater than with NAS or iSCSI.

And for those that just don't want to involve themselves in yet another IT skill set, managed storage services now cover the entire spectrum. Sometimes it is just less-expensive, easier, or faster to call in the professionals and leave everything to them. ■

Storage Security Basics

By Drew Robb

Given the emphasis administrators and corporate managers place on IT security, it's hard to imagine an environment in which security implementations are not a primary concern. As such, many of today's network IT administrators carefully consider all aspects of security when deploying and managing their networks.

Despite all the well-documented threats and media attention, however, there is no shortage of networks that are still operating with minimal and poorly implemented security measures. This can be due to lack of knowledge about the real risks to data security, unaddressed vulnerabilities, and sometimes to a false sense of security due to reliance on inadequate security strategies.

Jupiterimages

Storage networking technology has enjoyed strong growth in recent years, but security concerns and threats facing networked data have grown equally fast. Today, there are many potential threats that are targeted at storage networks, including data modification, destruction and theft, DoS attacks, malware, hardware theft and unauthorized access, among others. In order for a SAN to be secure, each of these threats must be individually addressed. Fortunately, many of the security practices and protocols used to address traditional network vulnerabilities also help ensure the availability of storage networks by reducing common security threats.

At the ground floor of any security strategy are some basic security concepts, including authentication, authorization, encryption (confidentiality), integrity, accountability and access control. We'll start with access control.

Access Control

Access control is a cornerstone concept when designing a secure network environment. Access control is all about controlling who can and cannot access a network, a resource, a folder or file.

In order to effectively secure such resources, you must carefully consider and control the level of access granted to each network user and then deploy strategies to ensure that only required users actually have resource access. It is a fundamental concept, and the foundation for a strong and secure network environment.

There are several types of access control strategies, including mandatory access control (MAC), discretionary access control (DAC), and role-based access control (RBAC).

MAC represents the tightest form of access control. In this strategy, security policies prevent the creator of any information from controlling who can access or modify their data. Instead, administrators or managers maintain control over who can access and modify data, systems and resources. Mandatory access control systems are commonly used in highly secure network environments such as military installation or financial or medical institutions.

MAC secures information and resources by assigning sensitivity labels on objects and comparing this to the level of sensitivity a user is assigned. This label is a kind of confidentiality stamp; when a label is placed on a file it describes the level of security required to access that specific file and will only permit access by files, users and resources with a similar or lesser security label.

MAC assigns a security level to all information, and places security clearance to each network user to ensure that all users only have access to that data for which they have security clearance. For example, users may be assigned a security label such as top secret or

confidential, and data and resources are classified accordingly. MAC restricts access to objects based on a comparable sensitivity between the user-assigned levels and the object-assigned levels.

The administrator or the operating system policy does not force discretionary Access Control (DAC); instead, an object's owner controls access. In a DAC model, if a user creates a folder, that user decides who will have access to that folder.

DAC is associated with an access control list (ACL). The ACL maintains information on the rights a user has to a particular system object, such as a file, directory or network resource. Each object has a security attribute that identifies its access control list and the list has an entry for each system user with associated access privileges. The most common privileges include the ability to read a file (or all the files in a directory), to write to the file or files, and to execute the file (if it is an executable file or program).

> "Twenty percent of companies do not know or are not in a position to tell if their storage security has been breached."

Microsoft Windows 2000/2003/XP, Linux, UNIX and MAC OS X are among the operating systems that use access control lists, although the list is implemented differently by each operating system. In Windows NT/2000/2003, an ACL is associated with each system object. Each ACL has one or more access control entries (ACEs) consisting of the name of a user or group of users. The user can also be a role name, such as "secretary" or "research." For each of these users, groups, or roles, the access privileges are stated in a string of bits called an access mask. The system administrator or the object owner typically creates the access control list for an object.

In a role-based access control (RBAC) configuration, access decisions are determined by the roles that individual users have as part of an organization. In any organization network users are assigned specific roles such as marketers, salespeople, managers, secretaries and so on. Users with similar roles are grouped together, and access control is determined by the role those

Storage Security

by Paul Rubens

Back in the days when storage meant direct attached storage (DAS), storage security was included in overall IT security. But as storage architectures have developed with the introduction of high-speed, high-capacity Fibre Channel-based storage area networks (SANs) as well as more traditional Ethernet-based network attached storage (NAS) systems, storage security has become a discipline in itself. Neglect it at your peril.

The starting point for a systematic approach to storage security, according to Sal Capizzi, a senior analyst at Boston, Mass.-based Yankee Group, is to take stock of the various types of data being stored and classifying it according to how important it is and how costly it would be to the business if it were lost or stolen. Then for each classification, appropriate security policies should be set.

The next step, Capizzi says, is to enforce password and World Wide name identification (for Fibre Channel) and logical unit number (LUN) authorization to ensure that only authorized users, devices or applications can access data, and to implement LUN masking so that particular storage volumes can only be seen by authorized users, devices or applications.

Ensure that all actions, accesses and changes to data are logged to provide a clear audit trail of who did what to which data from where, and when. Without such logs it is very hard to tell if or how data has been compromised.

Finally, don't neglect the boring obvious stuff: Use anti-virus, and anti-spyware software and a suitable firewall, disable unused ports, change passwords frequently, and so on.

– Paul Rubens, Enterprise Storage Forum

users have on the network. Role-based access requires a thorough understanding of how a particular organization operates, the number of users and their exact function in that organization.

Access rights are grouped by role name, and the use of resources is restricted to individuals authorized to assume the associated role. For example, within a school system, the role of teacher can include access to certain data, including test banks, research material, memos and related material. School administrators may have access to employee records, financial data, planning projects and more.

When a user is associated with a role, the user should be assigned only those privileges necessary to do their job. This is a general security principal known as the "least privilege" concept and applies to all access control methods. In a role-based scenario, when someone is hired for an organization, their role is clearly defined: teacher, secretary, sales, marketing, manager, etc. A new account is created for the user and then placed in a group with those with the same role within the organization. Individual permissions do not need to be set; rather, the level of access control is inherited from the group in which they are placed. As an example, if a new teacher is hired for a school, the user account is placed in the Teachers Group. Once in the group, the new employee will inherit the same level of access as those already in the Teachers Group.

Role-based access control is actually a form of MAC, since access is dictated by an administrator and the criteria for object access in not in the hands of the owner.

Authentication, Authorization and Accountability

Poor user authentication and authorization are one of the most common weaknesses in networks, and storage area networks are no different.

Poor user authentication and authorization are important concepts in network security. Authentication refers to the process by which you verify that someone is who he or she claim they are. This traditionally involves a username and a password, but can include any other method of demonstrating identity, such as a smart card, biometrics, voice recognition, fingerprints, and so on. Authentication is a significant consideration for network and system security and an important part of maintaining secure access control. Authentication security is controlled through policies and protocols. In an IP LAN/WAN environment, CHAP, EAP and MS-CHAP are examples of authentication protocols. There are also authentication protocols unique to a SAN environment, including both a secret key design with DH-CHAP authentication and public authentication with FCAP (Fibre Channel Authentication Protocol).

Authorization refers to the process of determining if a user, once identified and authenticated, is allowed to have access to a particular resource. This is usually determined by finding out if that person is a part of a particular group that provides the correct permissions, rights or required level of security clearance to access a resource.

Accountability refers to the tracking mechanisms used to keep a record of events on a system. One tool often used for this purpose is known as auditing. Auditing is the process of monitoring occurrences and keeping a log of what has occurred on a system. It is largely up to the administrator what types of events should be tracked and which should not. By tracking events on a system, it is hoped that attempts to access the network or otherwise compromise data will be recorded and prevented.

Confidentiality and Integrity

In any security strategy, protocols are needed to prevent data from being read by intruders (confidentiality) and others to determine if data has been tampered with during transit (integrity).

To prevent data from being read, encryption is used. Encryption takes raw data and scrambles it in such a way that it is unreadable without the key. If the correct key is not available, the stolen data maintains its confidentiality. As an example, within IPSec, the Encapsulating Security Payload (ESP) protocol can encrypt data sent over Fibre Channel links. Regular Ethernet communications can also use IPSec encryption or other protocols such as the Secure Sockets Layer (SSL) protocol. All encryption protocols are designed to make intercepted data unreadable to ensure confidentiality.

Integrity refers to the checking of data to ensure that data has not been tampered with or modified in any way. As an example, during the IPSec key exchange process, initial negotiations use one of two integrity verification methods, the message digest 5 (MD5) or Secure Hash Algorithm (SHA), to ensure that data has not been tampered with during the process.

Storage Budgeting Tips

By Henry Newman

With the price per gigabyte of storage coming down rapidly, that line item is no longer the overriding consideration for most storage budgets. While that is some relief for storage users, in other ways it creates a new problem: how long should you wait for storage to get faster and cheaper before you buy?

Add to that the complexity of upgrading to new technologies - 2Gbps vs. 4Gbps Fibre Channel, for example, or SAS vs. SATA, SCSI or Fibre Channel - and you're confronted with an array of planning and budgeting issues when it comes time to upgrade or replace your storage architecture.

Budgeting for storage is not just about buying more density or the latest cool stuff; it is about determining your needs based on available technology, and making sure those requirements are met.

Jupiterimages

The important issues to consider when budgeting for storage are:

1. How will a new technology integrate into the current environment?
2. Will this technology meet user requirements for performance and reliability?
3. How does this new technology affect O&M (operation and maintenance) costs?

Integration

Integration of technology into the current environment is a large problem for several reasons. Let's take a real-world example from an actual site. They have servers from one vendor and storage from another. The storage vendor can provide a new storage infrastructure that will support 4Gb Fibre Channel RAID controllers, 4Gb Fibre Channel switches, and other storage components. That all sounds great, but can the the server side support the 4Gb architecture?

This is a big question that should be asked of every hardware vendor. A standard PCI bus running at full rate supports 536 MB/sec, but many PCI buses do not support this full rate, and even though the situation is better, the same is also true for a PCI-X bus running at approximately 1.1 GB/sec (twice the PCI rate). A two-port 2 Gb HBA can require up to 800 MB/sec (200 MB/sec for each port reading and 200 MB/sec for each port writing). Therefore, a standard PCI bus cannot support two-port HBAs running at 2 Gb, which would be the same as one port at 4 Gb.

From a failover point of view, having two ports with 2 Gb provides greater redundancy if an HBA port fails, which is more common than both ports failing. This assumes that you have an HBA failure and not a PCI bus failure. In the case of PCI-X, a two-port 4 Gb HBA far exceeds the PCI-X bus bandwidth, (1.1 GB/sec for PCI-X, and two ports of a 4 Gb HBA require 1.6 GB/sec for full rate), so performance is far closer to that of two ports of a 2 Gb HBA.

All of these performance numbers assume that the I/O being done is streaming I/O. If it isn't, then why even consider 4 Gb HBAs and infrastructure in the first place? Yes, you can get improved IOPS performance with 4 Gb HBAs from a larger command queue, but the performance improvement is not that great and is often very workload-dependent. Ranges I have seen are from 0%-20%, but your mileage may very. This improved performance is surely not a justification to run out and buy a 4 Gb infrastructure.

The bottom line is that any site considering 4 Gb technology must make sure that the servers can support this new performance level. More often than not, large

servers lag in bus technology, given the large lead time it takes to design the complex memory interconnects to the bus and the availability of new bus technology. You can buy PCI-Express bus technology from Dell on one, two and four CPU systems, but try to find that on large (greater than 16) multi-CPU servers today.

User Requirements

User requirements should be a major driver of technology upgrades. Many organizations do not have a good handle on what the user application profiles look like, what the growth requirements are, and worst of all, whether the system is configured and tuned for those application profiles. This lack of understanding of the environment can lead to poor decisions on what hardware and software is needed.

One system I recently reviewed did not have an emulation or characterization of their workload. This is especially important for large sites. Without this information, how could this large site test patches for performance degradation (yes, it happens all too often), test new technology to measure performance improvements, or test increases in workloads to see if the system can handle them?

User applications and requirements should be a large component in any decision to upgrade technology. If you do not know what users are doing with the system, how do you know what they need today, let alone plan for the future? This situation often turns into a fire drill when the system is overloaded, and management starts throwing money at the problem instead of executing a master plan for technology infrastructure upgrades.

O&M Considerations

Technology maintenance costs almost always follow the same pattern:

- The cost of O&M for new technology is high for early adopters.
- Over the next 6 to 18 months, the cost drops as the technology is more widely adopted.
- The cost continues to drop, and drops sharply when a technology replacement is released, until...
- The cost skyrockets as the vendor tries to phase out the technology. This value is far greater than the original cost of maintenance, and sometimes I have seen it go as high as five times greater, since the vendor no longer wants to support the technology because of its cost and wants you to upgrade.

This is the general lifecycle for O&M costs. It makes sense given vendor costs, and unless technology trends change, the pattern is likely to continue.

One other area that should be considered is the personnel cost to the organization of supporting old hardware and software. You're not likely to find a new hire who knows how to work on Fibre Channel arbitrated loop HBAs, RAIDs and switches, and finding training course for that hardware isn't an easy task either. Just recall the frantic search for mainframe COBOL programmers for Y2K - a clear example of personnel operations costs becoming unreasonable.

> "Fewer than 25 percent of either Unix-/Linux- or Windows-based IT organizations had their own storage management team at the end of 2004. By the end of 2006, however, that number is expected to soar above 75 percent."

Conclusions

The issues addressed here are the ones that drive the high cost of storage changes. Most sites know what their physical storage growth will be, or at least what the budget will allow them for physical storage growth. The major cost items are not adding a few trays of disks with 146 GB drives or swapping out 36 GB drives for 300 GB drives; the major cost drivers are the infrastructure. The real question is how do you determine what you need, how much it is going to cost, and how to fit it into your current environment.

One pitfall: sites think they can just jump into new technology without fully understanding the whole data path (the path from the application to the operating system to the HBA/NIC to the storage devices). Plugging 4 Gb

HBAs in current servers into a 2 Gb storage infrastructure does not generally improve performance unless you are aggregating the performance of multiple RAID controllers and multiple hosts. The science (some call this an art, but it is really based on scientific analysis and study of the data path) of determining what users need and when they will need it is the process of budgeting for storage.

You need a full understanding of:

- Your current environment, including the performance level that environment can support today and the performance level that environment can support given technology trends;
- User requirements for performance and growth, including the current workload and the trend line for growth (performance mapped to expected new technology); and
- Your current and future O&M costs. Don't wait until your maintenance contract ends to find out that the cost has sky rocketed - technology maintenance costs follow a pattern.

Budgeting for storage is considered by many to be a complex problem, but it's not very complex if the lines of communication between the affected groups are open and free-flowing. The key is to have the data - seeing the future does not require a crystal ball, just some understanding of what you have and what you use, mixed in with a bit of history. ■

About this information

This content was adapted from EarthWeb's Enterprise Storage Forum Web site. Contributors: Dan Muse, Paul Shread, Drew Robb, Mike Harwood, and Henry Newman.