

Brief Table of Contents

Preface xxix

Part One: The Business of Being an IT Manager 1

1 The Role of an IT Manager 3

2 Managing Your IT Team 19

3 Staffing Your IT Team 43

4 Project Management 95

5 Changing Companies 139

6 Budgeting 161

7 Managing Vendors 179

8 IT Compliance and Controls 201

Part Two: The Technology of Being an
IT Manager 229

**9 Getting Started with the Technical
Environment 231**

10 Operations 255

11 Physical Plant 283

12 Networking 311

13 Security 349

14	Software and Operating Systems	383
15	Enterprise Applications	411
16	Storage and Backup	433
17	User Support Services	451
18	Web Sites	469
19	User Equipment	499
20	Disaster Recovery	515

Storage and Backup

...all you need in life is a little place for your stuff...

—GEORGE CARLIN

CHAPTER TABLE OF CONTENTS

1. Managing the Data	434
2. Disk Storage Technology	437
3. Tape Storage and Backup	443
4. Information Lifecycle Management.....	448
5. Additional Resources	449

Even though the name of the entire field is *Information Technology*, the heart of that information is *data*. And the growth in the amount of data is proceeding at an exponential rate. It isn't uncommon for a presentation or word processing document to be several megabytes in size. Where once the transmission of a file that size would be considered prohibitive, it's now commonplace and routine.

In addition, individuals are more reluctant than ever to delete files, preferring to keep copies of everything just in case they may need it again. Certainly, compliance and regulatory requirements have encouraged this posture.

CHAPTER SIXTEEN

16.1

Managing the Data

If you're responsible for all that data, the first thing you need to know is where it all is. The easy issue is the servers: they're easier to identify and locate. More complex are the issues surrounding the workstations that users have — some in the office, and some in remote locations — and the data that is stored on those C: drives with no management and backup. Some user education and policies will help.

Data Retention

Working with your company's Legal department, you may want to establish some data-retention policies. Very often, the first item addressed regarding data retention is e-mail. Many companies are now setting specific periods of time for retaining e-mail, after which the items are automatically deleted. You may permit users to get around the auto-purge by moving items to another location, but this requires that the user take a specification.

One of the Reasons Data Retention Got a Bad Name

Everyone first understood the real value — and the danger — of computer data retention during the Federal Government's anti-trust suit against Microsoft in the late 1990s. David Boies, the government's lead attorney, "consistently used e-mail and other documents to impugn the credibility of Microsoft witnesses." Old e-mails suddenly became not only sources of historical interest but weapons to be used against their authors.

Source: <http://news.com.com/2100-1001-241335.html>

Related issues about the age of data include:

- ◆ How long should the files of a user that has left the company be kept?
- ◆ How long should backup tapes be kept?
- ◆ What happens to files in "shared" areas of the network (see the section Shared Data Storage on page 436)? If a file hasn't been touched in several years it could be pretty unlikely that anyone is aware the file is there, or would even think to look for it.

Many organizations are now using retention levels as an opportunity to minimize their risk of culpability in legal matters. If the data isn't there, then it

can't be found. However, this is a double-edged sword in that retention levels may work to eliminate files you need, as well as those you don't.

IT has to work with user departments in setting data-retention practices. It is important — and often difficult — to remember that the users are generally considered the “owners” of data, while IT is the “custodian.” While all sales data may look the same to you, all hardware looks the same to them.

User Education

It's important to educate users about good practices with their data:

- ◆ Users have to understand that unless data is stored on the network, it won't be backed up. Many organizations have implemented desktop configurations that prevent users from storing files locally.
- ◆ For users who travel, the use of off-line replications of network files can give them the best of both worlds — network storage and local access. See the section on Replication (Synchronization) on page 339 in Chapter 12, Networking.
- ◆ Users should be advised about removing large attachments from e-mail messages and saving them to a network drive. Multiple copies of large attachments in e-mail can waste enormous amounts of space.
- ◆ Some technically savvy users like to periodically back up their workstations, and they find that the network is a convenient place to do so. Since a backup like this would also include the operating system, and software applications, the vast majority of this would be unnecessary (since those items are easily restored from the original media). Work with the users to show them how to use a CD or DVD burner to back up their data, or explain to them that their network files are already backed up each night.

Users often don't understand the issues of storage. They know that disk drives with hundreds of gigabytes of space are fairly cheap at their local PC store. While that fact is true, cost isn't the only issue in a corporate environment, and the cost of high-end drives used in corporate environments is considerably more expensive than the disk drive in the family PC. Users usually don't understand that while disks are cheap, their IT department isn't in the position of making unlimited space available. IT departments know that when storage space grows, it impacts the backup procedures (number of tapes, drives, and time required), and in the event of a disaster, it impacts the amount of time needed for recovery.

It's common for IT departments to periodically scan servers for "junk," which could be:

- ◆ Files with extensions or prefixes that indicate they're temporary files, created by applications, that weren't properly deleted
- ◆ Copies of software installation images
- ◆ Games
- ◆ Music
- ◆ Personal photos and videos
- ◆ Backups of workstations

Shared Data Storage

For data stored in a user's private directory, or data stored in application databases, it's generally pretty easy to identify who owns that data. However, most networks also include space for shared data which might be for a team, a department, or a specific project. These shared areas can easily deteriorate into dumping grounds over time.

While it probably isn't possible (or desirable) to eliminate the use of shared directories, there are procedures to help manage them:

- ◆ Users shouldn't be allowed to create shared areas on their own; they should be required to submit a request to IT for that.
- ◆ The request should identify which specific users are allowed which type of access to that area.
- ◆ The request should identify who will serve as the overall manager of that space, so that when IT has questions or concerns about it or its files, they have a specific individual they can work with.

You might get pushback from some users (who have, for example, worked at other companies where users could create whatever shared spaces they wanted). But the days of limitless access and almost limitless disk space are over; while abuse of network disk space isn't the most pressing item on an IT Manager's agenda, it can be one that easily gets out of hand.

Quotas

Disk quotas are annoying, but they can be an effective tool for helping to manage disk space. The most common types of quotas put in place include the size of

- ◆ A User's mailbox
- ◆ A Network directory
- ◆ The total amount of space available to an individual user in all locations

Of course, as with any policy, there will be some exceptions for some users. However, quotas, along with occasional messages about exceeding them, as well as having users request more space from IT, can all be a great way of keeping users conscious of the fact that space is an issue.

16.2 Disk Storage Technology

Disk drives are the nuts and bolts of data storage, and they come in many shapes and sizes.

Direct Attached Storage

Direct Attached Storage (DAS) is the term used for storage that is part of a server. The primary difference between DAS and other storage configuration (like network attached storage and storage area network, discussed below) is that DAS storage is directly connected to (or part of) a server, while the others are essentially stand-alone units that attach to a network. For larger amounts of storage, DAS is considered to be an inefficient use of hardware since a considerable investment has to be made in servers, which may end up going under-utilized.

Network Attached Storage

Network Attached Storage (NAS) refers to storage hardware that connects directly to your Ethernet network. NAS combines traditional disk array technology with "intelligence." This intelligence comes from a processor and small operating system, embedded in the NAS unit. To a certain degree, since a NAS device does include a small operating system, it could technically be considered a server. However, the entire unit is optimized for storage and ease of use. NAS units can be installed very easily and quickly, right out of the box, without a lot of the complex configuration needs of a "traditional" file server.

NAS solutions are based on file-level access, unlike storage area networks, (below) which are based on block-level access. Files on the NAS are made available over the LAN using file-sharing protocols such as NFS or Microsoft's CIFS. Users on the network are able to access files from the NAS without having to go through a traditional server.

Storage Area Network

There is some debate within the industry as to what characterizes a Storage Area Network (SAN). At a minimum, most agree that SANs do block-level access, as opposed to file-level access of NAS solutions. Data traffic on SANs is very much akin to the traffic used with internal disk drives, like ATA and SCSI. In a SAN, the server issues a request for specific blocks of data from the disks. This method is known as “block storage.”

However, there are some that say SAN solutions are also characterized by a fibre channel network, as opposed to TCP/IP and Ethernet, which is discussed below in the section Storage Network Connectivity (iSCSI and Fibre Channel).

Just a Bunch of Disks

Just a Bunch of Disks (JBOD) is a term that has been developed to differentiate lower level storage solutions from higher level ones (like SAN and NAS). There are some in the industry that will say that JBOD is synonymous with DAS, and others say that JBOD refers to “unprotected” storage (e.g., no RAID or mirroring, see page 440 for the section on RAID).

Storage Network Connectivity (Fibre Channel and iSCSI)

Fibre Channel

Traditionally, SAN environments have used Fibre Channel (FC) for connectivity, which can provide gigabit speeds and has become the most common connection method for SANs. Despite its name, FC can be implemented on copper cabling as well as fiber-optic cables.

Since FC is considerably faster, it's often used as an alternative to SCSI for connecting servers with storage devices. FC requires an investment in specialized hardware, which is a deterrent for many users. Each server must have an FC host-bus adapter (HBA) to allow it to connect to an FC switch (which is

analogous to a traditional network switch). The combination of FC cabling, HBAs, and switches is commonly referred to as the Fiber Channel, or SAN, “fabric.” FC can provide speeds of 2 Gbps, and is considered the ultimate solution for enterprise and mission critical needs.

A variation of the Fibre Channel, known as Fibre Channel over IP (FCIP) is bridging the gap between traditional IP networking and FC. Intended for connecting geographically distant SANs, FCIP can only be used with FC technology, and associated special hardware; in comparison, iSCSI will run over existing Ethernet environments.

iSCSI

iSCSI (IP SCSI) overcomes many of the limitations of traditional Small Computer Systems Interface (SCSI) (such as distance limitations and number of devices) by using traditional Ethernet, which would then allow an almost limitless number of devices, and extends the geography dramatically — essentially to any device on the LAN or WAN. In addition, with traditional SCSI, the disk drives can only communicate with the device that has the SCSI controller. However, an iSCSI device can be shared by multiple servers. iSCSI has a significant advantage in that it can leverage your existing Ethernet infrastructure for your SAN solution. However, the trade-off for that convenience is speed. iSCSI provides about half the speed of FC technology.

Although iSCSI can use traditional Ethernet hardware, some installations are making use of iSCSI HBAs in their servers. These HBAs are specialized cards that can off-load TCP/IP and iSCSI processing from a server’s CPU, and increase overall performance.

Disk Drive Types

There are a number of different disk drive technologies available on the market:

- ◆ **ATA** (also known as IDE) is the type of drive most often used in desktop and laptops, and uses a 16-bit parallel interface. The latest iteration of the standard supports transfer rates of 133 MB/sec and is believed to be about the maximum available.
- ◆ **Serial ATA (SATA)** picks up where ATA ends. Its transfer rate is 150 MB/sec, and future iterations are expected to see transfer rates of 300 and 600 MB/sec. In addition, the SATA devices draw less power and are easier to install.

- ◆ **SCSI** has gone through a number of iterations and enhancements, with the latest being SCSI Ultra/320 with a transfer rate of 320 MB/sec. An SCSI environment consists of a controller and a cable, with all the disk drives connected to the same cable. Of course, it's common to have multiple controllers and multiple cables to ensure redundancy and eliminate single points of failure.

Factors that Impact Disk Drive Performance

While picking the right disk drive technology is one part of the process for optimizing performance, it isn't the only item to consider. The transfer rate that a disk can deliver depends on a variety of factors:

- ◆ **Rotation speed.** This is the speed the disk spins at and ranges from 5,400 to 15,000 rpm. The higher the speed the more often the data on the disk will be in the right position to be read by the drive heads, and the faster data can be transferred.
- ◆ **Average Access Time.** This is the average time it takes to position the heads so that data can be read. The faster the better.
- ◆ **Cache Size.** This is the size of the cache built into the disk drive hardware. Just like the cache on a chip or system board, the bigger the better. However, cache is expensive and reaches a point of diminishing returns beyond a certain size.
- ◆ **Internal Transfer Rate.** This is the speed that data can be transferred *within* the drive. This will be a bit faster than the actual transfer rate of data to other system components because of associated overhead.

In addition to the above issues, other factors like network performance, server performance, application efficiency, traffic loads, etc., can impact the overall result of how long it takes to access the data you need.

Redundant Array of Inexpensive/Independent Disks

Because protecting your data is such a vital concern, a number of technologies have been developed to ensure that the loss of a disk drive doesn't mean the loss of your data. Redundant Array of Inexpensive/Independent Disks (RAID) defines different levels of protection, each having various trade-offs of performance and cost. These trade-offs are shown in Table 16.1, below.

Table 16.1 RAID levels (Source: www.digidata.com)

Common RAID Levels (Note that RAID level numbers indicate different types of data layouts, not higher performance or availability)

RAID Level	Description	Data Reliability	Performance	Application Strength	Cost
RAID 0	Data striped across multiple disks	Not true RAID; any disk failure causes loss of data	Faster than single disk for reads because data striped across several disks can be read simultaneously; similar to single disk for small writes	General	Low; because there is no redundancy
RAID 1	All data copied onto two separate disks	Very high; can withstand selective multiple disk failures	Faster than single disk for reads because data mirrored on two disks can be read simultaneously; similar to single disk for small writes	General	High; requires twice as many disks for redundancy
RAID 10	Two copies of data striped across disks	Very high; can withstand selective multiple disk failures	Very high; for reads, access is very fast because data is both mirrored and striped (i.e., there are two disks from which to read any piece of data, and striping spreads data across more disks)	High data reliability and performance such as ERP and image processing	High; requires twice as many disks for redundancy
RAID 2	Data bit striped across disks with error correcting codes on additional disks	Very high; can withstand multiple disk failures	Slowest of all the RAID levels due to bit striping	Slow speed makes RAID 2 unattractive; it is not used commercially	

(Cont.)

Table 16.1 RAID levels (continued)

<i>RAID Level</i>	<i>Description</i>	<i>Data Reliability</i>	<i>Performance</i>	<i>Application Strength</i>	<i>Cost</i>
RAID 3	Data striped across disks on separate channels with dedicated parity disk	Much higher than single disk; can withstand single disk failure	Faster than a single disk, owing to parallel disk accesses	Video, prepress, medical imaging, and other large file applications	Low; requires only one disk per RAID group for redundancy
RAID 30	Data striped across disks on separate channels with dedicated parity disk and across disks on the same channel	Much higher than single disk; can withstand single disk failure	Faster than RAID 3	Video, prepress, medical imaging, and other large file applications	Moderate; designed for large arrays
RAID 4	Data block striped across disks on separate channels with dedicated parity disk	Much higher than single disk; can withstand single disk failure	Faster than a single disk, owing to parallel disk accesses	Video, prepress, medical imaging, and other large file applications	Low; requires only one disk per RAID group for redundancy
RAID 5	Data and parity striped across multiple disks	Much higher than single disk; can withstand single disk failure	High compared to single disk for reads but lower than single disk for writes	OLTP, e-mail, ERP, Web, CRM	Low; requires only one disk per RAID group for redundancy
RAID 50	Data and parity striped across multiple disks and across disks on the same channel	Much higher than single disk; can withstand single disk failure	Faster than RAID 5	Transaction processing with high read to write ratio	Moderate; designed for large arrays

16.3

Tape Storage and Backup

While disk is the technology for online storage, tape is the preferred method for backup solutions (see Table 16.2 below for pros and cons).

Table 16.2 Pros and cons of tape as a storage medium

PRO	It's very inexpensive. As such, multiple copies (such as having complete sets for different points in time, or copies in different locations) can be easily kept without incurring enormous expense.	CON	It's much slower than disk.
	It's easily transportable. This makes it convenient for storage at an off-site location and easy to retrieve, if needed, to a recovery site.		It's a "sequential" medium; that is, if the file you need is at the end of the tape, you have no choice but to read through the entire tape to get to it.
			It's a more fragile medium than disk; is more susceptible to deterioration over time.

Tape and Tape-Drive Technologies

Just as there are different kinds of disks, there are different kinds of tapes (see Table 16.3).

Variations on Backup

Although tapes are the standard for backup, there are a number of variations on this solution that have been adopted.

Disk-To-Disk-To-Tape

Because tape is a slow medium, and because backup does impact overall performance of the production environment, some environments have implemented

Table 16.3 Tape drives (Source: www.govconnection.com)*Tape Drive Comparison Chart*

Type	Vendors	Formats	Capacities (GB): Native/Compressed*	Transfer Rates (GB/h): Native/Compressed*
Travan	Certance	Travan 20, Travan 40	10/20, 20/40	3.5/7, 7/14
AIT	HP, Sony	AIT-1, AIT-2, AIT-3	35/91, 50/130, 100/260	14/37, 16/43, 33/86
DAT	Certance, HP, Sony	DDS-4, DAT72	20/40, 36/72	10/20, 12.6/25.2
DLT vs80/160	HP, Quantum	DLT vs80, DLT vs160	40/80, 80/160	11/22, 29/58
DLT 8000	Quantum, HP	DLT 8000	40/80	22/44
Super DLT	HP, Quantum	SDLT220, SDLT320, SDLT600	110/220, 160/320, 300/600	36/72, 58/116, 130/260
LTO/Ultrium	Certance, HP	LTO-1, LTO-2	100/200, 200/400	54/108, 108/216
Super AIT	Sony	S-AIT	500/1300	108/280

*All AIT/S-AIT speeds and capacities assume 2.6:1 compression. All other speeds and capacities assume 2:1 compression.

a backup solution where the disk storage is first copied (often referred to as taking a snapshot) to another disk environment, and that second disk environment is then backed up to tape. This offers a number of benefits:

- ◆ Because disk performance is much faster than tape, a disk-to-disk backup can happen very quickly. As such, the impact to the production environment (e.g., shutting down databases during backup) is minimized.
- ◆ Because the disk-to-disk backup happens fast, the performance hit to the production environment is also reduced.
- ◆ The backup of the snapshot image happens faster since the backup process isn't competing with any other processes or users for performance. The snapshot image exists only for the benefit of the backup environment.
- ◆ If a user needs a file restored and that file is still available from the snapshot image, it can be restored easier and faster than it could be restored from tape.

The cost of the device that holds the snapshot image is the only downside to this alternative. And, since it really doesn't have to be a high-end device (i.e., no redundancy is really needed since it's only an interim copy of the data), it may be quite affordable for small and mid-size environments.

Disk-to-Disk

Because there is a certain security risk with a portable medium like tape, some organizations have tried to eliminate the use of tape altogether by using disk as an alternative. However, there are some significant factors to consider with the elimination of tape:

- ◆ Because tapes are often used for “archival” purposes (i.e., allowing you to have complete copies of your environment from different points in time), the cost of doing the same with disk can be prohibitively expensive.
- ◆ Disk storage isn't as portable, so it could impact the flexibility of your disaster recovery plans.
- ◆ Because of the risk of having your live data and the backup copy in the same location, it's worth considering doing the disk-to-disk copy over a WAN connection to a remote site (perhaps your disaster recovery site). The cost of a line with sufficient bandwidth could be very expensive. And, while the need for replicating only the “deltas” to your backup site

(as discussed in the section Data Replication in Chapter 20) may make for relatively fast backups, restoring an entire server over a WAN connection may take an extensive amount of time.

Data Encryption

Again, because tape is a portable medium, there are certain security risks associated with it. Citigroup and Bank of America are just two of the companies that have had to deal with lost backup tapes that contained confidential personal data (see Chapter 8, IT Compliance and Controls and Chapter 13, Security). To limit the risks of backup tapes falling into the wrong hands, many organizations have started to encrypt their backup. Decru and Neoscale Systems are two companies that offer solutions for encrypting a backup.

A key up to 256-bits in length is commonly used to encrypt the data on the tape. When using encryption for your backup tapes, you must have copies of the key stored in multiple locations, most important at your disaster recovery site. Without the key, you won't be able to decrypt your backup tapes, and you will be unable to restore any data.

Dedicated Network

Another alternative for backup solutions is to set up a dedicated network for it. Because of the volume of data that travels across it, network performance could be adversely impacted as it competes with users and application traffic for network bandwidth and throughput.

By providing a dedicated network for backups, it allows the backup to perform faster, and eliminates any impact it may have had on the users and the production environment.

Backup Schedule

Backup traditionally comes in three flavors:

- ◆ **Full.** Captures everything in the environment, regardless of when the data was last modified or backed up.
- ◆ **Differential.** Captures those files that have changed since the last *full* backup.
- ◆ **Incremental.** Only captures those files that have changed since the last backup of *any* kind.

Traditionally, IT environments have done a full-complete backup of everything over the weekend, followed by differential backups on weeknights. However, with the incredible growth of data storage, backup has become almost a non-stop process. Operations managers are constantly juggling the backup scheduling of different systems in order to maximize the efficiency of all resources, and ensure that backup procedures don't interfere with other applications and activities.

Full

A full backup is a complete copy of everything in the environment, regardless of when files were last backed up, accessed, or modified. Full backups take the longest amount of time to complete and tie up the most resources such as tapes and tape-drives. A full backup is the handiest when a total restore is needed, since it has a complete image.

Differential

With differential backups, a complete restore of an environment would require the last full backup to be restored, followed by a restore from the last differential backup.

Incremental

Incremental backups execute the fastest, since they capture the smallest number of files. However, they make the restoration more complex. For a full restoration, the last full backup needs to be restored, along with each incremental backup, in order, that was taken since the full backup.

You should also check with your application and database vendors about the integrity of backups that are performed when files are open. You may need to shut down the application and/or database to get a reliable backup.

Backup Storage

Backup tapes are usually stored at an off-site location. The thinking is that if a disaster wipes out your primary facility, you don't want to lose your backup tapes at the same time. If you do, then you lose any ability to recover data.

Some companies may choose to store their tapes at another company location, most opt to use dedicated records storage facilities that have appropriate security and climate controls.

One of the challenges with off-site storage of the backup tapes is determining when to send them off site. For the highest level of protection in case of disaster, the tapes should go off site immediately after they are created. However, chances are that any requests you get for a restore will require the most recent backup tapes, so there is some convenience and cost efficiency involved in leaving them on-site for a short period of time.

The intermediate disk storage discussed in the disk-to-disk-to-tape backup solution (see section above) can also be used for restoring data when the tapes are sent off site. Another alternative is to make two sets of backup tapes, one kept on site for any restores that are needed, and the other sent off site to be used in case of disaster.

16.4 Information Lifecycle Management

Information Lifecycle Management (ILM) is one of the more recent topics to warrant a large amount of attention in the storage industry. In a nutshell, ILM reflects the fact that the value and use of data and information change over time. The cafeteria lunch menu isn't as important as the payroll data. And, whatever value the cafeteria menu has this month, it will be negligible next month; while the payroll data will be valuable for quite some time. Similarly, ILM recognizes that there are aged data (e.g. files that haven't been modified in some time), as well as static data (e.g. files that generally don't change, such as graphics and multimedia files). How data's use and value changes can vary greatly from industry to industry and from organization to organization.

With ILM, an organization has to define different use classes of data (e.g., critical, static, archival, transactional, etc.), and the needs of that data over time (high-speed access, real-time replication, etc.). With those two items, you can then define different classes of storage solutions (online disk, tape, DVD/CD, on-site/off-site, etc.).

Since the needs and use of data will change over time, an ILM solution (which is a combination of hardware and software) allows you easily manage the data and seamlessly migrate it to different storage solutions during the data's lifetime. With a comprehensive ILM solution, you're ensuring that you're not wasting a very expensive storage solution for aged data that is hardly ever accessed or neglecting to back up critical data that change radically over the course of a day. An ILM solution also ensures that you don't lose track of where the data is during its lifetime.

ILM has grown in popularity just as compliance issues (see Chapter 8, IT Compliance and Controls) have become a more serious concern in IT. To comply with various rules and regulations regarding the retention of data,

organizations have begun to adopt ILM solutions as way of helping to ensure compliance (i.e., the tracking of data). At the same time, companies are making cost-effective and cost-conscious decisions as to what technology to use to store data with different needs, so that the high-cost, high-performance, and high-availability storage solutions are used only for the data that needs it.

ILM versus Hierarchical Storage Management

The discussion of ILM may sound suspiciously like Hierarchical Storage Management (HSM). They both have some common elements in that data is moved to other storage forms at certain points.

The key difference between the two concepts is that HSM primarily relies on the measure of access frequency and/or age to determine when a file should be moved to another storage medium. HSM also is primarily one-directional, in that it generally moves data from primary to secondary, and from secondary to tertiary storage solutions. ILM, on the other hand, can move data based on a number of policies (in addition to age and access frequency) and to/from any storage solutions.

16.5 Additional Resources

Web Sites

- ◆ *www.ca.com* (backup software vendor)
- ◆ *www.decru.com* (storage security vendor)
- ◆ *www.emc.com* (storage solution vendor)
- ◆ *www.hitachi.com* (storage solution vendor)
- ◆ *www.hp.com* (storage solution vendor)
- ◆ *www.ibm.com* (storage solution vendor)
- ◆ *www.neoscale.com* (storage security vendor)
- ◆ *www.netapps.com* (storage solution vendor)
- ◆ *www.quantum.com* (storage solution vendor)
- ◆ *www.veritas.com* (backup software vendor)