



Storage Management

Choosing the right solutions will help enable your organization to handle ever-expanding data requirements while controlling costs, delivering performance and responding to the needs of the business.

Table of Contents

- 1 Executive Summary
- 2 Assessing Your Needs
- 2 Evaluating Best Practices
 - Storage Area Networks
 - Storage Tiering
 - Data Deduplication
 - Storage Virtualization
 - Improving TCO/Efficiencies
 - Business Continuity
- 6 Taking That First Step
- 7 CDW: Storage Management Partner

Executive Summary

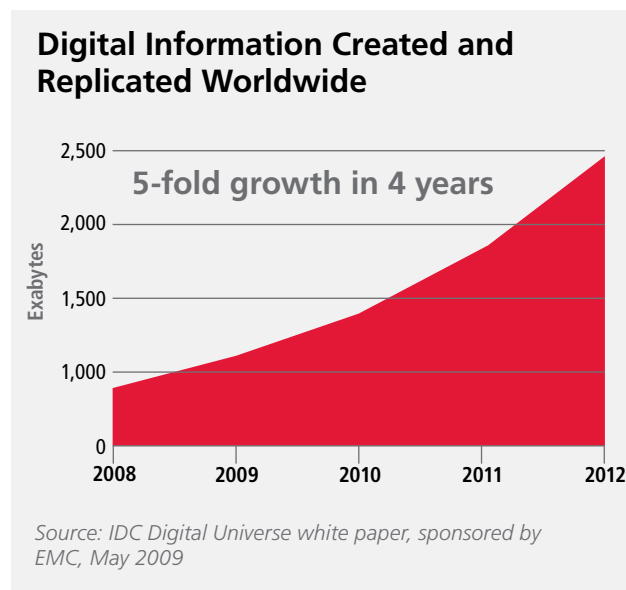
Managing data storage is no longer simply a matter of adding new drives when your organization needs more capacity. Businesses of all sizes are producing more data than ever before and, with stringent compliance and e-discovery requirements, they are also keeping more data for longer periods of time. Many businesses are already spending up to 20 percent of their IT budgets on data storage and experts warn that there's a growing gap between the amount of data being created and the amount of available storage. Simply adding storage devices without a plan in place can put you on a slippery slope towards uncontrolled cost increases.

However, by taking a comprehensive, strategic approach to storage management, Fortune 1000 companies can not only manage data growth, they can also control costs, improve business processes and support innovative technology initiatives such as virtualization and cloud computing. State-of-the-art technologies in data storage management — such as storage-area networks, data deduplication, storage virtualization and storage tiering — are all designed to help organizations treat data storage as a vital business asset.

The bottom line: By utilizing technology strategically to manage data storage volume and performance, IT will be able to be more responsive to the needs of the business. And isn't that what every IT professional is being asked to deliver this days?

Assessing Your Needs

If you are a Fortune 1000 business, there really isn't any way to avoid dealing with growing data storage requirements. The Internet and digital media are contributing mightily to this swelling demand and businesses will never be the same. Files are getting bigger and IT infrastructures are now storing video and e-mails and PowerPoints and all kinds of unstructured data in addition to the structured data in the corporate database. The research firm IDC recently released a report stating that the amount of digital information across the globe grew by 62 percent between 2008 and 2009. Over the next 10 years, the amount of digital information created annually will grow by a factor of 44, IDC reports.



In addition, changes in compliance requirements and tough e-discovery laws mean that businesses are not only creating more data, they are also required to save it for longer periods of time. This puts a further strain on the storage infrastructure. No one in IT wants to be the one to tell the CEO that certain e-mails are not available if the business is facing civil litigation.

The reality is that building a strategic approach to data storage is no longer a luxury: It is a requirement. Fortunately, the leading vendors in the data storage industry have been able to come up with innovative ways in which businesses can cost-effectively and efficiently manage their growing data storage needs through advances in network storage, storage tiering, data deduplication and storage virtualization.

To figure out which mix of solutions will be best for your organization, you should start with an analysis of what you have in place and how your needs may be changing. It is always a good idea to bring in an experienced team to help you do an assessment of what you have, where you may have opportunities to save money and what kind of solutions will help you to meet your future demands.

In the world of data storage, where there are so many different variables to examine, you should work with a team that is not tied in to any one particular solution, but can help your organization plot out a course based on the individual needs and goals of your company. Your company's ideal solution may include a mix of products from different vendors, so there's no need to limit your options when a company such as CDW can offer a wide range of best-of-breed solutions in addition to extensive data storage expertise and experience.

Evaluating Best Practices

Typically organizations will see a variety of symptoms that would indicate an inefficient storage environment. These could include exceeding backup windows, missing recovery SLAs, not meeting retention guidelines for internal and external archive requirement, or low utilization rates due to the addition of dedicated storage for a new application.


A comprehensive data assessment will give you an idea of how your data resources are being used and where you may have problems with disparate data stores, large amounts of duplicate data, static data that is not being managed through its lifecycle, or devices that are not up to the task of delivering either the volume or performance required for your organization. For instance, an assessment would get into how much data you actually have on the network, how much of that data is duplicated, where it's being created and where it's being stored.

An assessment plan should also begin the process of evaluating and classifying the value of the data to the business: Which data is primary and must be accessed in real-time for mission-critical applications and how much is secondary and perhaps can be stored on a lesser-performance type of drive. How do you want to manage the data through its lifecycle, as it shifts from primary to secondary and eventually to archived data? By establishing policies to determine data protection, retention and access policies, IT organizations can identify and eliminate data that holds no value or, worse, could be unnecessary and potentially damaging to the company in litigation.

You will also want to conduct an analysis of future needs and future technology solutions. For example, if your organization is going to increase its utilization of server virtualization or client virtualization, you will need to take that into consideration when you evaluate future storage requirements. Virtualization and cloud computing are innovative technology solutions that require some level of centralized storage and centralized storage management. A recent EMC, sponsored IDC study reported the number of files, images, records and other digital information containers will grow by a factor of 67, each needing to be managed, secured and protected. Despite this growth, the number of IT professionals globally will grow only by a factor of 1.4.

Top enterprise storage challenges

In general, what are your organization's greatest challenges with respect to its storage environment? (Percent of respondents, N=504, multiple responses accepted)



Source: ESG Research Brief, Enterprise Storage Priorities Emphasize Information and Infrastructure Efficiency, January 2009

A number of new technologies have evolved that address the growing storage needs of today's enterprises. Although the technologies vary widely, ultimately storage management boils down to a few basic strategies:

- Maximize utilization of the storage media at your disposal
- Reduce the size of the data that needs to be stored
- Store each piece of data in the most efficient and cost-effective way possible

Historically, each application — be it a server, a shared network drive or an end-user system — had its own storage associated with it. This forced IT departments to manage each piece of storage directly and separately. Plus, most of this storage went unused for great lengths of time.

Among the key issues and technology solutions your organization will want to explore are:

Storage Area Networks

For many businesses, the first step in consolidating data storage and getting storage under central management is moving to a storage area network. A SAN offers remote storage capacity to servers by utilizing an architecture that mimics storage attached to the operating system. It takes disparate silos of disk arrays that had been attached and dedicated to a single application and links them into a high-speed network.

Four Performance Considerations

Most storage managers grapple with four major issues when it comes to storage performance:

- **CAPACITY SCALING:** The ability to easily upgrade storage without disrupting the system.
- **PERFORMANCE SCALING:** The necessity of maintaining acceptable service levels as storage capacity grows and the number of supported hosts increases.
- **AVAILABILITY:** The concept of redundancy that ensures failover; the ability to immediately switch to a redundant system when a failure occurs. Failover ensures availability of data despite events that would otherwise make information inaccessible, possibly resulting in systemwide service interruptions.
- **MANAGEABILITY:** The idea that systems should, as much as possible, automate scaling, manage capacity, and ensure failover with minimal direct human intervention.

Source: SearchDataCenter.com

The obvious advantages of utilizing SANs are that they simplify storage administration, centralize management, make it simple to scale up and add new devices, and put the pieces in place for organizations to manage data through the entire course of its lifecycle. A SAN enables organizations to consolidate and manage their storage requirements proactively while creating the high availability of data required by most businesses today.

Another major advantage of having a SAN is the ability to manage data based on its value to the organization: You don't necessarily have to keep purchasing high-performance drives for data that is not mission-critical. Also, as you scale up and purchase newer devices, you can shift older devices to handle secondary or tertiary data so you wind up extending the lifecycle of your drives and reducing your overall total cost of ownership. Utilizing and repurposing existing equipment can be a major advantage of upgrading to a SAN.

SANS are available in a variety of configurations:

- **FIBRE CHANNEL:** FC is a gigabyte technology that migrated from the supercomputer field to become the standard protocol for storage area networks in the enterprise storage environment. FC is relatively expensive. It can use fiber-optic cables, as its name implies, as well as twisted-pair copper wires.
- **INTERNET SMALL COMPUTER SYSTEM INTERFACE:** iSCSI is a networking storage standard based on an Internet protocol. As such, iSCSI can run on existing network cabling, providing cost savings not available with expensive Fibre Channel. IP compatibility also allows storage of data to and retrieval from remote locations.
- **FIBRE CHANNEL OVER ETHERNET:** FCoE is a relatively new standard that adapts high-performance Fibre technology to more ubiquitous and cost-effective Ethernet lines.
- **INFINIBAND:** This technology, commonly used in supercomputer applications, is another SAN option that seeks to compete with FCoE and iSCSI. InfiniBand provides higher availability to stored data, but at a price.

Utilization is an important consideration when evaluating whether to deploy a SAN or whether to figure out a better way to optimize your SAN infrastructure. According to The Enterprise Strategy Group, utilization rates are often in the 30 percent to 50 percent range, which means organizations are paying to house, power, cool and manage all of the capacity they do not use.

For many businesses, a comprehensive data assessment will help to figure out how to utilize SANs more efficiently: Throughout a business, you may find that some departments are exceeding their capacity on a regular basis, while others are underutilizing capacity. This can create extreme inefficiencies and wind up costing the IT department a lot of money. The idea is to get this under control in a centralized environment where IT can manage resources flexibly.

Storage Tiering

Storage tiering is basically the process of categorizing and storing data based on its value to the organization. Its main purpose is to reduce storage costs by moving data quickly and efficiently to its most appropriate device or tier of storage. Mission-critical primary data, such as sales orders, is stored on the highest-performance devices, while secondary data can be stored on older devices. The other major advantage of tiered storage is that performance improves throughout the storage infrastructure, but particularly at the top tier where performance matters most.

Once data is created, it begins its lifecycle. Its value changes over time and as its value changes, storage tiering enables the organization to store it appropriately. Keeping data on the same tier throughout its lifecycle can be a significant waste of precious storage resources and could be a major drain on throughput and I/O.

Common tiered storage configuration provided by GreenPages Technology Solutions, an IT consultancy include:

TIER 1 STORAGE: For operations-critical 24x7 databases, file servers, e-mail applications and data warehouses, a redundant, cache-based tiered storage model (called Tier 1 storage) is the best option. The Tier 1 storage model offers quick response times and fast data transfer rates. As such, Tier 1 storage is a great solution for organizations that need to effectively store high-performance data that demands high availability.

TIER 2 STORAGE: For seldom-used, noncritical databases (historical data, for instance), a Tier 2 storage model is a great option, because Tier 2 data can generally be stored on less expensive media in a SAN. Tier 2 storage is a good option for organizations that have a large amount of data that does not require 24x7 availability or extensive backup. Tier 2 storage can also help reduce hardware costs and management overhead.

TIER 3 STORAGE: For rarely accessed data, a Tier 3 model offers further economies of scale, because data can be stored on even less expensive media, such as recordable compact discs. Tier 3 storage is a convenient and simple way for IT administrators to protect large amounts of noncritical data from fire, theft and computer malfunctions.

Storage tiering can be manual or automated. Under each scenario, it is important to go through a process of analyzing all of the data in the organization — what's already there and also what is likely to be created — and to assign it value based on a set of characteristics determined by the organization. These could range from how old is the data, to what applications it supports, to how often it has been accessed during a given period of time.

Either way, the assessment will help the organization to establish rules and processes for the successful movement of data between tiers. With manual tiered storage solutions, administrators actually move the data, and with tiered storage solutions, the data is moved between tiers based on a set of algorithms. Automated tiered storage solutions are a growing part segment of the market, but they are more expensive to deploy.

With the significant growth in data storage requirements and the pressure to manage costs, many IT decision-makers are turning to tiered storage solutions for at least a segment of the storage infrastructures. According to a survey last year in *Storage* magazine, nearly 60 percent of IT professionals are using some sort of tiered storage solution. Of those, approximately 36 percent move data manually from tier to tier. More than 40 percent said classifying data is the biggest pain point related to their tiered storage solution, so it is important to recognize that if you move in this direction, there is a lot of prep work to be done in advance. The advantage of doing this prep work is that it only really has to be done once and then you have in place a system that can be much more flexible and scalable as your storage demands expand, which, based on all of the research, is probably a foregone conclusion.

Data Deduplication

Data sprawl can be a significant drain on storage resources if it's not managed carefully. Think about a PowerPoint presentation or a video that might be sent to dozens of people throughout a department. Is it necessary for each individual to be saving a copy? And do they all have to be backed up? With data deduplication, one copy of each file can be saved, significantly reducing costs and overhead.

Data deduplication uses algorithms to search for duplicate data on systems and replaces that duplicate data with a "pointer" back to the one master copy stored within the network. There are two common levels at which data deduplication works — file and block level. At the file level, data deduplication searches for redundant files on the network and only saves one copy of that file. When a change is made to a document in the network, another copy of the file is saved. With block-level data duplication, when a document is changed, only the changed data is stored again, making it more efficient than block-level data deduplication in terms of the amount of data stored. The tradeoff is that it takes more

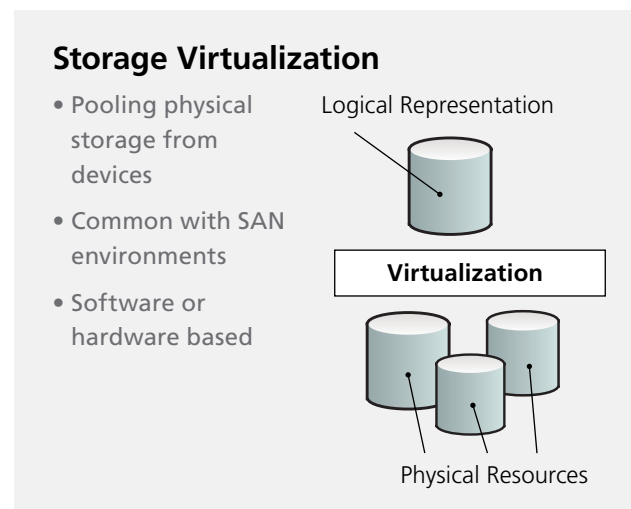
processing power to complete the duplication process at the block level.

There are two main strategy layouts for deduplication — inline and post-process deduplication. With inline deduplication, the data is deduped while it is being backed up. The advantage of this is that you can quickly replicate the backed up data to the secondary site. The drawback is that it is a power-intensive process. With post-process deduplication, the data is deduped after the data is backed up to the disk target, but before the backup job ends. The advantage is that there is no impact on the backup and recovery throughput. The disadvantage is that additional storage is required to hold the pre-deduped backup data.

Vendors offer a variety of ways in which to deduplicate data: It can be done inline or post-process and it can be purchased as software or as an appliance. It can be also done at the server level or throughout client devices as well. Choosing the best solution — if any — should be the result of an expert assessment plan. Some vendors claim they can reduce storage requirements by as much as 80 percent with deduplication across physical and virtual backups, but realistically most companies can expect to reduce their storage requirements by about 20 percent to 40 percent with data deduplication.

Storage Virtualization

Storage virtualization is a pooling of physical storage resources from multiple storage devices within your network to appear as a single storage device. Storage virtualization will typically support a variety of storage devices at one time, including Fibre Channel, SATA and SAS drives and will optimize their usage based on the tasks they are performing.



Storage virtualization is typically used within a storage area network and can be software based or could come as part of your network hardware. Among the key features: central management and the ability to dynamically move data from one physical device to another without disruption to the business.

Storage virtualization also enables manufacturer independence, since it supports multiple devices and protocols. This means organizations don't have to be tied in to one vendor as new generations of technology emerge. It also can help the business achieve significant cost savings as well as a rapid return on investment by enabling consolidation, reducing energy consumption and reducing IT overhead through simplified management. Storage virtualization will also enable your organization to increase its use of server virtualization throughout the enterprise.

Benefits of Storage Virtualization

Storage virtualization helps the IT manager or storage administrator perform the tasks of backup, archiving and recovery more easily.

Other value includes:

- Single point of administration
- Nondisruptive data migration
- Information lifecycle management
- Improved allocation efficiencies
- Heterogeneous replication

Storage virtualization, in fact, complements server virtualization, and both technologies can help increase flexibility within the IT infrastructure and help the IT organization be more responsive to business needs. This is mostly because a virtualized infrastructure is easier to change: Changes become virtual, not physical. With storage virtualization, management of the storage infrastructure is more automated and less disruptive to the organization. It is much easier to scale up as your data storage requirements grow, and it enables your organization to save money by adding devices that deliver appropriate levels of performance for the task.

As you assess your storage requirements, you want to look at the level of server virtualization currently within your organization and also plan on future demands.

According to the analyst firm Gartner, approximately 50 percent of server workloads will be run on virtualized servers within the next couple of years, up from less than 20 percent today.

Clearly, virtualization is going to be more and more important to Fortune 1000 businesses, and as you go further in that direction, the idea of having a SAN and managing storage centrally will be critical.

Improving TCO/Efficiencies

Going from a disparate to a central storage environment will help you cut costs and improve efficiencies. When an organization reduces its storage needs, it has a ripple effect that can help to contribute to a reduced total cost of ownership. One of the advantages, for example, is delayed hardware purchases: If you are reducing storage needs by 20 percent to 40 percent with deduplication, that could impact purchase cycles. This can also lead to savings in energy consumption, software licensing fees, maintenance and support contracts.

There can be savings in IT resources involved in deploying new devices, launching new applications and maintaining infrastructure at remote locations. These are critical concerns to consider when analyzing the advantages of moving to a more centralized storage infrastructure. Centralized storage, for example, can significantly reduce maintenance costs. IT no longer has to send personnel to remote sites to install and troubleshoot new storage devices.

In addition, devices can be hot swappable and adding systems on the fly is much simpler, faster and less costly. Centralized storage also saves money by improving efficiencies: If one department is using only 20 percent of its capacity, that capacity can be easily shifted to another department that needs it.

Business Continuity

For many businesses, there is a tipping point when it is time to turn to centralized storage management. Sometimes a business will have multiple locations and will be looking to avoid having different storage solutions in each spot. Sometimes organizations will be exceeding their backup and will begin looking at their overall storage infrastructure. Others will recognize that their storage requirements are growing significantly and it doesn't make sense to just add more devices.

However your organization comes to centralized storage — or if you are already doing it but need to make it more efficient and cost effective — you can accomplish a lot for your business in terms of improving business continuity. Centralized storage can improve your disaster recovery plans, your backup, redundancy, compliance and security. With centralized storage, organizations are backing up less data, which means backing up over a wide area network is a realistic option. In addition, organizations can store data longer on primary storage and disk targets before moving it offsite, which provides faster data recovery times.

Taking That First Step

It's important to understand the technology solutions that are available to you and it's equally important to understand the roles that data storage and storage management can play in your company's overall business strategy. For IT decision-makers at Fortune 1000 businesses, one of the goals decreed

by corporate management is for the company to “move at the speed of business.” By taking a strategic approach to storage management, IT organizations can help to deliver on that promise while managing costs and improving overall IT efficiencies. The first step is to make that commitment to manage data as a strategic resource and then to bring in the right partner to turn that commitment into a state-of-the-art actionable plan. Don’t wait until your IT infrastructure is overrun with data and your bosses and employees are in a panic. Now’s the time to take that first step.

CDW: Storage Management Partner

With expertise in every facet of storage management solutions, CDW can assist with the selection and implementation of a solution that fits your business needs.

We spend the time necessary to understand your business goals to help you identify which storage management solution provides the best fit for you. In addition, CDW provides a single, reliable resource for systems and storage technologies — from planning and implementation to support services. We’re certified to design, sell and install the storage management solutions of numerous vendors across the United States.

Getting Started with Storage Management

Your CDW Account Manager and certified storage specialists are ready to assist you with every phase of choosing and leveraging the right data storage management solution for your IT environment. Our approach includes:

- An initial discovery session to understand your goals, requirements and budget
- An assessment review of your existing environment and definition of project requirements
- Detailed vendor evaluations, recommendations, future environment design and proof of concept
- Procurement, configuration and deployment of the final solution

When you partner with CDW, you partner with the best in the business.



Second only to your people, your information is your organization’s most important asset. EMC provides the technologies and tools that can help you release the power of your information. We can help you design, build, and manage flexible, scalable, and secure information infrastructures. And with these infrastructures, you’ll be able to intelligently and efficiently store, protect and manage your information so that it can be made accessible, searchable, shareable and, ultimately, actionable.

You can also use EMC information infrastructure as the foundation for implementing your information lifecycle management strategies, securing your critical information assets, leveraging your content for competitive advantage, automating your data center operations, reducing power and cooling costs, and more.

CDW.com/emc

HITACHI Inspire the Next

Hitachi Data Systems helps organizations transform raw data into valuable information by making it more accessible and simpler to manage. Our vision is that IT must be virtualized, automated, cloud-ready and sustainable. Hitachi Data Systems provides best-in-class information technologies, services and solutions that deliver compelling customer ROI, unmatched return on assets (ROA) and demonstrable business impact. As the leader in storage virtualization, only Hitachi Data Systems offers a common, virtualized platform for all data and information. Our solutions cut costs, reduce operational complexity and improve IT agility. Data drives our world — and information is the new currency.

CDW.com/hitachidatasystems



Looking for a better way to manage your demanding and all too often unpredictable IT environment? Then look to HP.

HP StorageWorks Resource Management solutions bring stability to your business environment. Whether securing resources, measuring quality of service or managing through change, HP management solutions deliver the control you need to provide the right information to the right people at the right time.

HP’s approach helps simplify operations by enabling you to manage from the view that best meets your business needs. You can manage from a “single pane of glass” for a clear view of specific arrays, all storage, applications or broad IT view. The choice is yours. In addition, event notification features give you more time to plan responses. That way — whether it’s expanding capacity or resolving issues before they become problems — you can stay in control of your environment.

CDW.com/hp

About CDW

CDW is a leading provider of technology solutions for business, government, healthcare and education. Ranked No. 41 on Forbes' list of America's Largest Private Companies, CDW features dedicated account managers who help customers choose the right technology products and services to meet their needs. The company's technology specialists offer expertise in designing customized solutions, while its technology engineers and solution architects can assist customers with the implementation and long-term management of those solutions. Areas of focus include notebooks, desktops, printers, servers and storage, unified communications, security, wireless, power and cooling, networking, software licensing and mobility solutions.

CDW was founded in 1984 and as of March 31, 2010, employed approximately 6,150 coworkers. In the 12 months trailing March 31, 2010, the company generated sales of \$7.6 billion. Intently focused on responding to customers' technology needs with a sense of urgency, CDW helps customers achieve their goals by providing the right technology products and services they need — when they need them.

