

OPTIMIZE ENTERPRISE STORAGE

Data storage optimization helps companies deal with data growth while controlling costs, increasing productivity and gaining competitive advantage.

Executive Summary

Businesses of all sizes are producing more data than ever. Even in the frail economic environment of the past several years, data is experiencing an average 50 to 60 percent compound annual growth rate, according to the market research firm IDC.

By some estimates, data storage spending accounts for more than 20 percent of enterprise IT budgets. Furthermore, experts warn that there's a growing gap between the amount of data being created and the amount of available storage. And simply adding storage devices, without a plan in place, can be a slippery slope toward uncontrolled cost increases.

By taking a strategic approach to storage management, businesses can effectively manage data growth. Furthermore, state-of-the-art storage technologies, coupled with newer concepts like cloud computing and data governance, are geared to help firms cut costs, streamline management and position data stores as vital business assets.

The bottom line: By strategically utilizing technology to manage ever-expanding data stores and performance needs, IT professionals can more effectively respond to the needs of the enterprise. And this, in itself, can be the key to gaining and keeping a competitive edge.

Table of Contents

-
- 1 Executive Summary

 - 2 Data Stores Growing

 - 2 Data Governance

 - 3 Assessing the Need

 - 4 Evaluating Best Practices
 - Storage Area Networks
 - Dynamic Tiered Storage
 - Data Deduplication
 - Storage Virtualization
 - Thin Provisioning
 - Storage in the Cloud

 - 7 Improving TCO/Efficiencies

 - 8 CDW: A Storage Management Partner That Gets IT

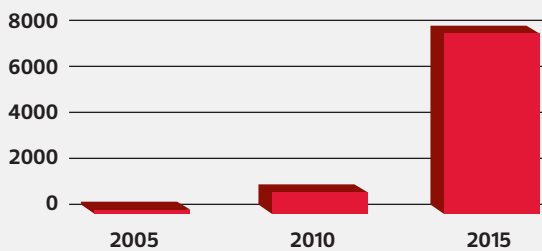
Data Stores Growing

For most organizations, 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. IT infrastructures are now storing video, e-mail, PowerPoint presentations and all kinds of unstructured data or information. And this is in addition to the structured data in the corporate database.

The research firm IDC's 2011 Digital Universe Study, *Extracting Value from Chaos*, sponsored by EMC, finds the world's data is doubling every two years with a massive 1.8 zettabytes to be created and replicated in 2011. And, according to IDC, over the next five years, the amount of digital information created annually will grow by a factor of nine.

A Decade of Digital Universe Growth: Storage in Exabytes



Source: IDC Digital Universe Study, *Extracting Value from Chaos*, sponsored by EMC, June 2011

According to the IDC survey, there are not enough people with the correct skills, experience and resources to manage the deluge of data and keep pace with growth. IDC also predicts that by 2015, nearly 20 percent of all digital information will be touched by cloud computing providers, compared to 2 percent now.

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 further strain on the storage infrastructure.

Another trend facilitating growing storage needs is virtualization. As Computerworld reported in September 2011, "The benefits of virtualization, particularly its ability to allow users to rapidly provision new workloads, are pushing demand for storage to new highs.

"In the pre-virtualization days, the process for requisitioning server space often required review by IT administrators." With virtualization, "a new application can be provisioned in minutes versus days, weeks or months." This capability is pushing up storage demands significantly.

The trend toward desktop virtualization is also affecting storage needs. Because this data does not "live" on desktop or notebook hard drives, no separate client security, backup

and update policies are needed. Still, bringing user data into the data center means a place must be created to store it.

The company looking to save on desktop operations through a virtual desktop initiative will likely have to build out more consolidated storage.

Data Governance

There's no question that data has value. It is a corporate asset that can be used for strategic business advantage. However, like all assets, it must be properly handled, processed and protected.

As the challenge of managing data increases, more businesses are implementing data governance. While the discipline is relatively new, it is seen as a way to effectively administer the information that serves as the lifeblood of the organization.

Data governance consists of a broad-ranging set of objectives, policies and processes designed to ensure that data assets are formally managed throughout the enterprise. For example, it defines data ownership, creation and updating, as well as arbitration measures should disagreements arise.

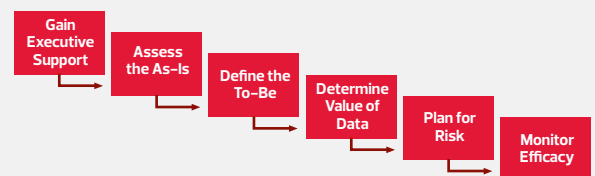
Other primary aspects of data governance include: decision-making authority, compliance monitoring, policies and standards, data inventories, lifecycle management, content management, records management, preservation, data classification, data quality, data risk management and data valuation.

Companies increasingly see data governance as a cornerstone to the success of master data management (MDM) projects. While data governance provides the procedures for proper data handling, MDM offers a technical foundation enabling the enterprise to link all its data into one "master" file or hub that provides a common point of reference.

Emerging Discipline

Data governance remains an emerging discipline. However, research shows that it is maturing faster than anyone could predict. This is especially true as businesses begin to realize the benefits derived from having good data.

The Data Governance Process



The data governance process is one that requires sustained energy of effort, maintaining commitments, communication and delivering results.

Source: *Data Governance – Managing Information as an Enterprise Asset*, NASCIO Governance Series

According to a 2010 study, conducted by The Information Difference, an analyst firm specializing in data quality, 71 percent of organizations surveyed have already implemented a data governance program or plan to within one year. This is 30 percent more than when the firm conducted a benchmark study on data governance in 2008.

"The two prevalent motivations for implementing data governance (for both those already in production and those planning implementations) were better quality and faster decision making," notes Dave Waddington, senior vice president and co-founder of The Information Difference. The ability to respond faster to business change also scored high in the study.

Of those companies with active data governance programs, 79 percent viewed those programs to be a success. In addition, the vast majority of those surveyed believe that it is imperative to have a data governance initiative in place prior to or in conjunction with implementation of an MDM program.

Data Governance and Technology

Enabling data governance requires a number of components. These include empowering people, setting up processes and getting significant help from technology. In fact, some consider leveraging technology to be one of the most critical parts of the data governance discipline.

By calling on a suite of technologies, IT decision-makers can achieve end-to-end data management on a highly effective cost basis. Effective data governance can be implemented using a number of technology tools including:

- **Data quality technology** – Used to examine and cleanse data, and to audit and measure quality
- **Data security tools** – Designed to protect data from intrusion and loss
- **Metadata management technology** – Used for all types of metadata – classification, description, guidance and control
- **Data profilers** – Intended to capture "real and true metadata" and used to understand data content and structure
- **Workflow and workgroup tools** – Offering process integrity, teamwork and collaborative data management
- **Wikis and portals** – Used to support two-way communication among data governance participants and stakeholders

Assessing the Need

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

A comprehensive data assessment will provide an idea of how data resources are being used and where there may be issues 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 by the organization.

For instance, an assessment would get into how much data a business currently has on the network. Other metrics include how much of that data is duplicated, where it's being created and where it's being stored.

An assessment plan should begin the process of evaluating and classifying the value of the data to the business. Consider 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.

Also consider how data is managed 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 shops can identify and eliminate data that holds no value.

IT shops also want to conduct an analysis of future needs and future technology solutions. For example, if the enterprise is planning to increase its utilization of server or client virtualization, this will need to be taken into consideration when evaluating future storage requirements. Cloud computing is another innovative technology solution that requires a level of centralized storage and storage management.



Evaluating Best Practices

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

- Maximize utilization of current storage media
- 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.

Today, there are highly effective ways to optimize storage. Technology solutions to explore include:

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.

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 the business to consolidate and manage storage requirements proactively, while creating the high availability of data required today.

Another major advantage of having a SAN is the ability to manage data based on its value to the organization. For example, purchasing high-performance drives for data that is not mission-critical is not necessary.

Also, when scaling up and purchasing newer storage devices, older devices can shift to handle secondary or tertiary data thereby extending the lifecycle of drives and reducing overall total cost of ownership (TCO). 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.

Storage Performance Considerations

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

- **Capacity Scaling:** This is the ability to easily upgrade storage without disrupting the system.
- **Performance Scaling:** This represents the necessity of maintaining acceptable service levels as storage capacity grows and the number of supported hosts increases.
- **Availability:** The concept of redundancy ensures failover or 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 system-wide service interruptions.
- **Manageability:** This includes the idea that systems should, as much as possible, automate scaling, manage capacity and ensure failover with minimal direct human intervention.

Source: SearchDataCenter.com

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 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 consider another way to optimize storage infrastructure. According to The Enterprise Strategy Group, IT analyst and business strategy firm, storage utilization rates are often in the 30-to-50-percent range, which means organizations are paying to house, power, cool and manage all of the capacity they do not use.

Studies show that a SAN can significantly increase storage utilization. The enhanced storage utilization found in a SAN is resultant from the ability to consolidate storage and create a shared storage resource designed for all server applications.

For many businesses, a comprehensive data assessment can help to figure out how to utilize SANs more efficiently: Throughout a business, there may be some departments that are exceeding their storage capacity on a regular basis, while others are underutilizing capacity.



This can create extreme inefficiencies and wind up costing the IT department significant amounts of money. The idea is to get storage under control in a centralized environment where IT professionals can flexibly manage resources.

Dynamic Tiered Storage

Also known as auto-tiering, tiered storage consists of categorizing and storing data based on its value to the business. 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 high-speed storage devices, such as hard disk drive arrays, while secondary data can be stored on slower devices, such as optical discs and magnetic tape drives. 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.

Common Tiered Storage Configuration

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 the enterprise that needs to effectively store high-performance data demanding high availability.

Tier 2 Storage: For seldom-used, noncritical databases (historical data, for instance), a Tier 2 storage model is a great option. Tier 2 data can generally be stored on less expensive media in a storage area network (SAN). Tier 2 storage is a good option for companies 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. This is because data can be stored on even less expensive media, such as recordable compact discs or magnetic tape. 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.

Source: GreenPages Technology Solutions: An IT consultancy

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.

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 enterprise – what's already there and also what is likely to be created – and to assign it a value based on a set of predetermined characteristics. These could range from how old the data is, 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 enterprise to establish rules and processes for the successful movement of data between tiers. With manual tiered storage solutions, administrators actually move the data. With automated tiered storage solutions, the data is moved between tiers based on a set of algorithms. While automated tiered storage solutions are a growing segment of the market, 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. By using tiered storage as a solution to manage data and reduce the number of drives to store data, companies can greatly reduce the space that storage requires in the data center.

According to a spring 2010 Storage magazine *Snapshot* survey, 50 percent of respondents said they currently use a form of tiered storage in their environments. In addition, 63 percent of those not currently using tiered storage strategies said they expect to in the future. The survey also found storage tiering becoming more automated, with 32 percent using an automation tool for tiered storage.

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.

The technique improves storage utilization. And it can be applied to network data transfers to reduce the number of bytes sent across a link. 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 level 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 deduplication, when a document is changed, only the changed data is stored again, making it more efficient than file-level data deduplication in terms of the amount of data stored. The tradeoff is that it takes more processing power to complete the deduplication 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 it allows for quickly replicating 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 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 storage assessment.

Some vendors claim they can reduce storage requirements by as much as 80 percent with deduplication across physical and virtual backups. However, realistically most companies can expect to reduce their storage requirements by about 20 to 40 percent with data deduplication.

Storage Virtualization

Storage virtualization is a pooling of physical storage resources from multiple storage devices within the 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 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.

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

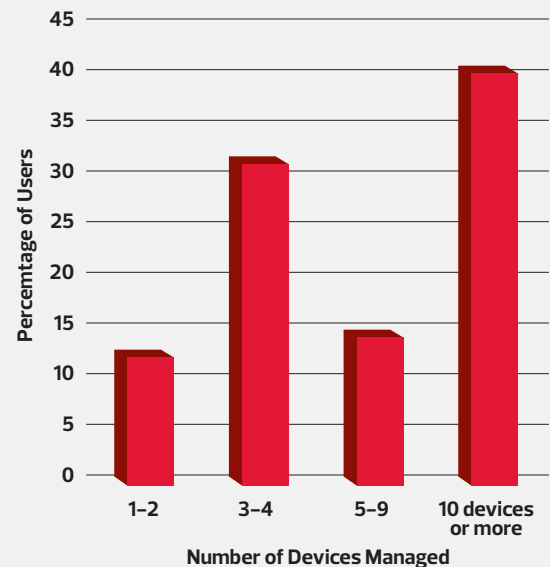
Storage virtualization also enables manufacturer independence, since it supports multiple devices and protocols. This means organizations don't have to be tied in to one vendor as new generations of technology emerge.

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 also complements server virtualization. In fact, the full benefits of server virtualization can only be realized when teamed with the virtual storage layer, which works in tandem with the virtual server layer. Both technologies help increase flexibility within the IT infrastructure. They also help IT shops be more responsive to business needs.

This is mostly because a virtualized infrastructure is easier to change as alterations become virtual, not physical. With storage virtualization, management of the storage infrastructure is more automated and less disruptive to the organization.

IT Managers Managing More Storage Devices



Nearly half of IT managers report managing 10 devices or more for primary storage, backup and disaster recovery.

Source: Survey conducted by Nimble Storage, 2010

It is much easier to scale up as data storage requirements grow. And it enables the organization to save money by adding devices that deliver appropriate levels of performance for the task.

Thin Provisioning

This is a way to essentially provide storage space for data and not actually create that space within the physical infrastructure of the storage system until it's needed. The theory here is that most applications and the people who run

them typically overstate what they actually require. Or, they build into their storage requirement a sense of what their growth of the data is going to be over time.

Thin provisioning is an effective way to raise capacity utilization within the storage system. This will help to alleviate the need to buy as much storage as would be needed with a physical infrastructure.

Storage in the Cloud

Storage in the cloud, or storage as a service (SaaS), is on its way to becoming an increasingly appealing option for IT shops as well as infrastructures. While many firms still rely on hard drives and external storage devices, others are discovering the benefits of the cloud.

Some of the more salient advantages of cloud computing include:

Accessibility – The cloud utilizes the Internet to drive the sharing of data. Therefore, information can be accessed regardless of location.

Redundancy – Cloud service providers typically offer backup services, thereby adding to data recovery due to outage or disaster.

Security – Instead of storing information on a physical device locally, the storage service provider offers both physical and virtual protection.

Scalability – Cloud data storage services allow firms to purchase more or less storage space, depending upon the business environment.

Economy – Storage as a service can result in the purchase of less hardware than a firm would need with internal storage.

Business Continuity – Because data is stored in an offsite location, by nature it is protected.

Cloud Disadvantages

Cloud storage service providers are in business to serve a number of clients. Therefore, one legitimate concern is that one business's data is stored on the same hardware as that of another firm.

Should one firm's data become compromised, there is a chance that other clients' data becomes compromised as well. To overcome this potential problem, some service providers use discreet devices for each business customer.

Cloud service providers also choose and implement their own security measures. Consequently, companies may feel insecure not knowing which other clients have the same security measures in place. One solution is to use different encryption keys for each individual corporation using storage services. Still, this is dependent on the service provider.

Essentially the storage system knows that when data comes into the system, it will simply allocate space for it. There's no need to have the space sitting there unused ahead of time.

Improving TCO/Efficiencies

Going from disparate storage devices to a central storage environment will help cut costs and improve efficiencies. When a business reduces its storage needs, it has a ripple effect that can help to contribute to a reduced TCO.

One of the advantages, for example, is delayed hardware purchases. Consider that if a firm is reducing storage needs by 20 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, the unused 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, when the enterprise comes to unified storage – or if it is already being done but needs to be more efficient and cost effective – it can accomplish a great deal in terms of improving business continuity.

Unified storage can improve disaster recovery plans, backup, redundancy, compliance and security. With unified storage, organizations are backing up less data, which means backing up over a WAN is a realistic option.

In addition, companies can store data longer on primary storage and disk targets before moving it offsite. This provides faster data recovery times.

CDW: A Storage Management Partner That Gets IT

Explosive data growth is an inevitable reality of today's business operations. Companies are not only storing more data, they are keeping their existing data longer.

The solution: A comprehensive effort to optimize infrastructure. This is one where storage needs can be appropriately sized per business conditions, redundant data can be eliminated and all available storage can be treated as a single pool for greatest utilization.

Data storage management solutions can benefit a business in a number of ways including:

- Reducing administrative overhead and freeing up IT resources
- Providing a central data storage management system
- Reducing data storage maintenance needs and expenses
- Achieving greater scalability for future growth and flexibility

CDW can assist with storage optimization in a number of ways. Along with offering products and services, we start with a server and storage assessment. From there we'll build a solution that's scalable, highly available and unified.

Your CDW account manager and solution architects are ready to assist with every phase of choosing and leveraging the right storage 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 manufacturer evaluations, recommendations, future environment design and proof of concept
- Procurement, configuration and deployment of the final solution
- Ongoing product lifecycle support

To learn more, contact your CDW account manager at 800.800.4239 or visit CDW.com/datastorage



Organizations of every type rely on the timely retrieval of information to facilitate transactions and decision making. While a typical organization is experiencing double-digit data growth, IT budgets, staffing and traditional storage capabilities are not keeping pace. As a result, IT organizations are under constant pressure to employ more efficient storage strategies and increase the amount of data their staff can manage without additional headcount. Customers are looking to storage vendors for innovations to solve these challenges in the same way that server virtualization has enabled them to experience greater efficiency by pooling server resources and dynamically provisioning compute power according to business needs.

CDW.com/emc



For most organizations, storage represents a significant portion of the IT budget. By improving the efficiency of your storage environment, you can help boost overall IT efficiency. Just imagine the possibilities if you could reclaim time, money and resources that are currently underutilized by simply looking at storage efficiency optimization through a new lens.

With NetApp®, you can improve storage utilization and optimization, store more data with fewer systems, use less space and power, and lower your operational costs – all while maintaining or even boosting IT performance.

CDW.com/netapp



HP 3PAR T400 and T800; Get agile, get efficient, get 3PAR

Fast, efficient single-system array, delivering maximum consolidation and performance headroom for enterprises and service providers.

HP 3PAR Utility Storage was designed from the ground up to exceed the economic and operational requirements of today's most demanding IT environments. It is a storage platform that delivers the performance, scalability and availability required of Tier 1 Storage along with unique technology benefits not available with legacy platforms.

CDW.com/hp



Back Up, Recover, Discover and Retain Information Faster and More Efficiently with Symantec.

Unforeseeable events should not bring your business to a halt. Symantec delivers leading backup, recovery and archiving technology to help you recover systems, data and e-mail – in minutes, not hours. Symantec allows you to archive e-mails and reduce storage costs by minimizing the base of information that must be protected. Also, safeguard information wherever it is stored – on desktops, notebooks, servers or e-mail.

CDW.com/symantec



The information is provided for informational purposes. It is believed to be accurate but could contain errors. CDW does not intend to make any warranties, express or implied, about the products, services, or information that is discussed. CDW®, CDW-G® and The Right Technology. Right Away® are registered trademarks of CDW LLC. PEOPLE WHO GET IT™ is a trademark of CDW LLC.

All other trademarks and registered trademarks are the sole property of their respective owners.

Together we strive for perfection. ISO 9001:2000 certified

88473 – 111125 – ©2011 CDW LLC

