A STRATEGY PAPER FROM

CENTER FOR DIGITAL EDUCATION

Managing the Storage Storm in Education

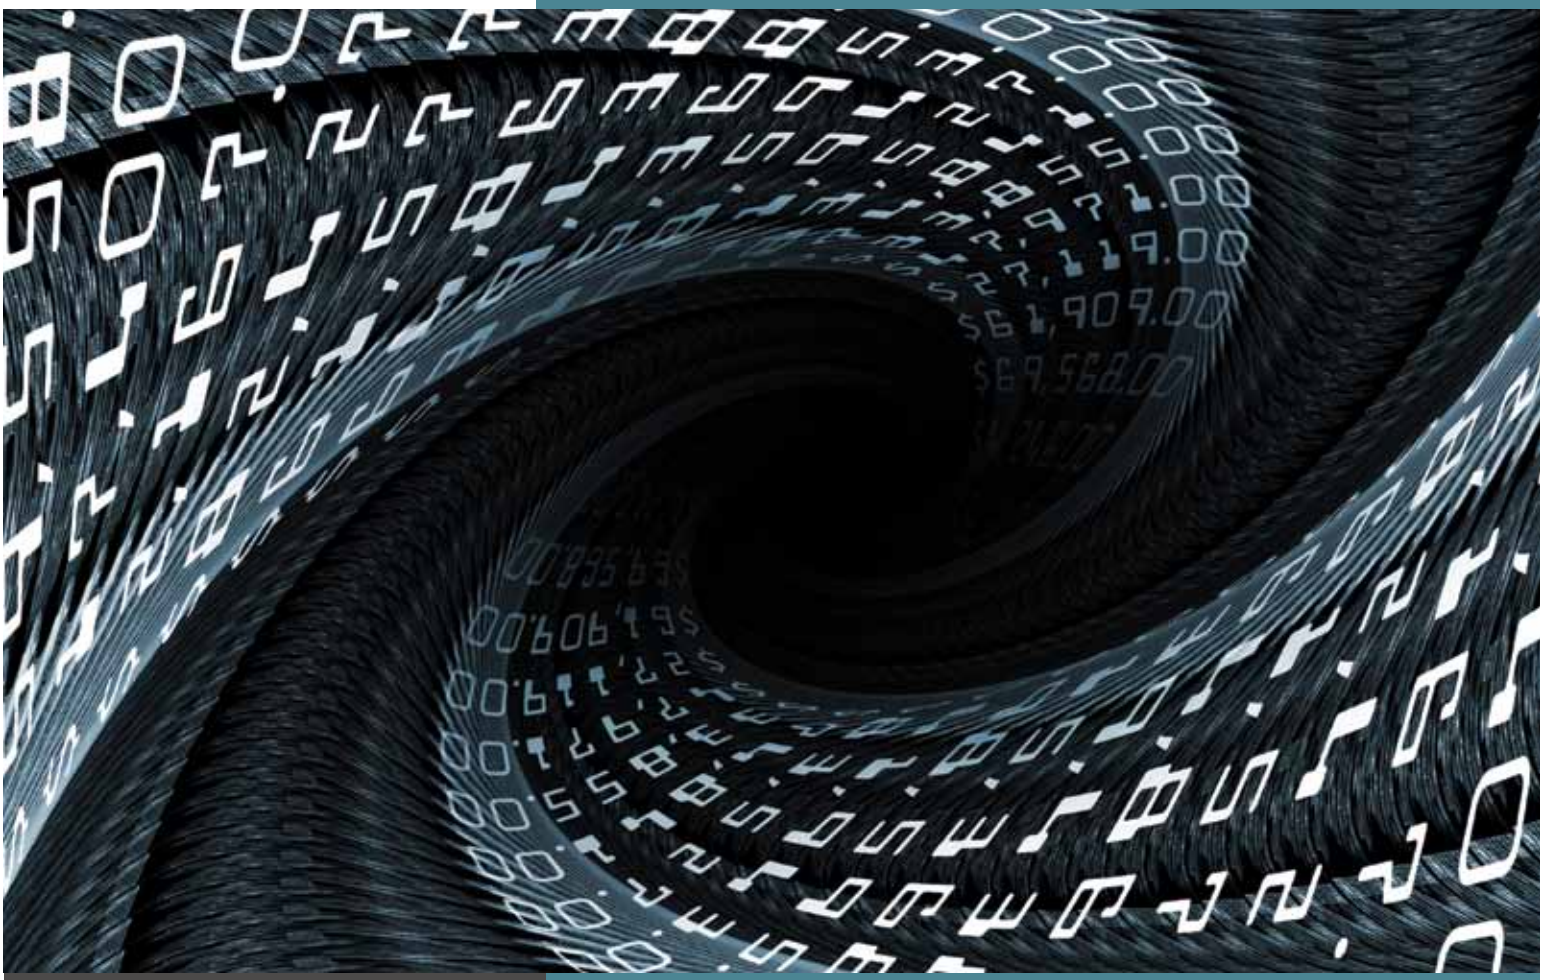

Managing the Storage Storm in Education

Executive Summary

As K-12 school districts, community colleges, and four-year colleges and universities cope with the increasing volume and complexity of digital content driven by new technologies, it is becoming critical that storage systems support the use of these latest technologies. While storage solutions may not be the primary focus of a campus in an age of hot new technologies, the importance of properly planning and managing storage will only increase in the years ahead as data growth rates continue to climb. All of this is happening as tight budgets are forcing IT managers to meet new demands without additional financial and staff resources.

Fortunately, balancing these challenging requirements is becoming easier, thanks to modern storage technologies and best practices that give IT administrators greater flexibility in managing data and using resources more effectively. Educational institutions need to develop a strategic end-to-end view of storage to ensure that storage infrastructure effectively supports teaching, learning, research and administration.

The goal of this white paper is to provide a guide to managing storage in the new digital education environment. The paper walks readers through steps that will help them implement strategies that improve efficiencies and cut costs by creating a solid and efficient storage program, including:

- Evaluating the current storage environment
- Considering a refresh while taking a look at the latest storage technologies
- Determining the possible role of service providers to help with modernization plans

Setting the Stage for Storage

Behind all the digital content, high-definition videos, sophisticated learning management systems (LMS) and data management solutions intrinsic to the education landscape, a quieter but equally significant transformation is playing out. The performance and reliability of these vital assets hinge directly on how successful the IT department is in building and managing the underlying storage infrastructure. Storage solutions typically aren't the technology systems that campuses focus on first and foremost — not in an age of hot new laptops, paper-thin tablet PCs and innovative social networking applications. But as digital content in all its forms explodes in volume and complexity, storage systems are becoming key factors in a campus' ability to use technology effectively and efficiently.

Consider this:

- Some universities report that current data volumes often reach hundreds of terabytes and that number is expected to double over the next 12 to 18 months.

- IT managers are struggling to absorb many types of new applications and data types, ranging from structured information that fits neatly into database systems to an exploding volume of free-form unstructured data associated with Web 2.0 and collaboration applications.

- There is an increasing amount of space-hungry high-definition video that permeates all levels of education, from interactive classroom resources to student services departments.

- Campuses face more stringent requirements for managing student records and generating reports that satisfy regulatory and funding requirements.

The pivotal role storage plays to meet information management goals will only increase in the years ahead as data growth burgeons and the requirements for greater access usher in additional challenges.

Fortunately, resolving these issues may become more manageable for two primary reasons. First, new storage technologies provide IT administrators with greater flexibility to manage data and use available capacities more effectively. Second, when necessary, campuses can selectively turn to third-party storage experts for help in reducing complexities ranging from assessing needs to managing day-to-day performance.

However, even with the proper tools and resources, IT administrators in both K-12 and higher education institutions need an overarching storage strategy that addresses today's new realities. Yesterday's "Band-Aid" approach of constantly adding new capacity is no longer viable or affordable.

Understanding the Challenge

One thing is becoming clear for campuses at all levels — storage systems are much more than just a place to house data. Today's challenge is to provide adequate capacity while simultaneously delivering performance rates high enough to serve sophisticated education applications and administrative needs. In short, just throwing more disk drives at the data explosion problem no longer works. In fact, that approach usually just adds unnecessary costs and complexity.

Instead, IT managers must look at storage needs and solutions in new ways to better manage data, take advantage of automated storage management innovations and ultimately buy proportionately less capacity than what exploding growth rates may indicate.

Why are storage considerations so different today than in the recent past? One reason is that the performance of servers, desktop devices and applications is directly linked to storage systems. IT resources must deliver high degrees of responsiveness as educational materials are being delivered online to administrators, teachers, students, researchers and other constituents. This means that each data file not only needs a home, but that home must be in the right place, at the right time and for the right cost.

The link between performance and storage systems becomes even more entwined in virtualized environments where effectively managing data in a centralized hub defines success. The underlying storage tier is critical to achieve server consolidation efficiencies. The increased demand on storage in a virtual data center also highlights limitations associated with traditional direct attached storage (DAS)

Instead of endlessly adding more capacity, educational organizations need a plan that points to long-term solutions rather than short-term fixes. The bottom line: K-20 campuses must develop a strategic view of storage that provides flexibility while meeting future demands.

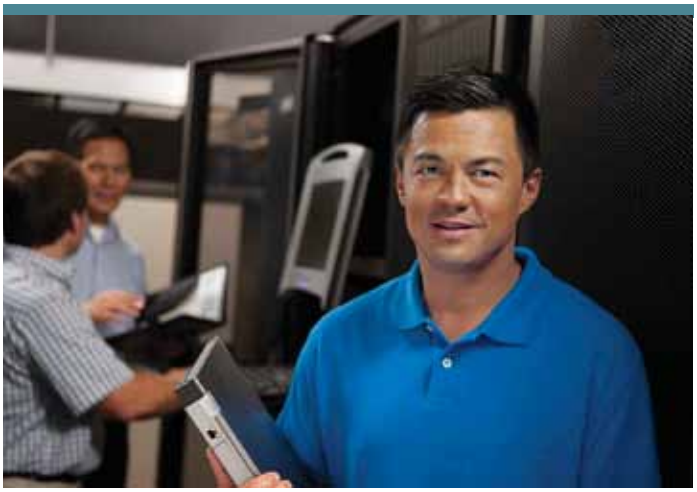

systems, where hard disks are dedicated to individual servers. These rigid ties challenge the IT department to quickly provision additional capacity when needed or scramble to keep operations running during load fluctuations.

Cost is another concern. Campuses can no longer afford a one-size-fits-all approach to storage. Environments where all data resides on a single storage platform don't provide sensible efficiency options to house older, less critical information on the most economical technologies.

What's the solution? Instead of endlessly adding more capacity, educational organizations need a plan that points to long-term solutions rather than short-term fixes. The bottom line: K-20 campuses must develop a strategic view of storage that provides flexibility while meeting future demands.

Assessments: A Valuable Baseline

For many education institutions, the first step in optimizing their storage systems is to perform a thorough assessment of their current IT operations. The resulting data provides a snapshot to help understand the existing environment while uncovering potential optimization opportunities for core storage functional areas such as:

- Service Delivery
- Staffing
- Financial Management
- Standardization
- Tools and Automation
- Tiered Archiving Strategies
- Data Deduplication
- Performance Metrics

However, a complete storage infrastructure assessment can require dedicated staff time for perhaps a month or more to gather the necessary information, analyze data and create a report of the findings. That's a commitment many resource-strapped IT organizations find difficult to make. An alternative is to hire outside technology experts to handle the data gathering phase, perform the analyses and use their expertise to recommend improvements.

Here are some must-have capabilities to look for from service providers:

- *Education expertise:* K-12 and higher education institutions share some unique requirements and resource limitations not seen in other vertical markets. To end up with a plan that's tailored for campuses, analysts should come from the education industry.

- *Storage specialists:* Since storage systems now play a pivotal role in overall IT performance and cost management, the assessment firm should have deep expertise in the latest storage technologies, networked architectures and implementation best practices.

- *A big picture perspective:* An end-to-end storage productivity assessment should also analyze improvement opportunities throughout the data center, including server and networking resources. For example, re-engineering storage systems may be done in conjunction with a server consolidation project that combines high-performance hardware with virtualization technology to reduce capital expenditures.

- *Multi-vendor experience:* Choose consultants with a broad knowledge about a variety of vendor platforms and a willingness to suggest integrated rather than point solutions.

To establish a solid and efficient storage program, IT managers need to take three key steps:

1. Evaluate the current environment
2. Consider a refresh while evaluating the latest storage technologies
3. Determine the possible role of service providers to help with modernization plans

Step 1: Evaluate the Current Environment

Campuses have several options available for assessing their storage. Options include using their internal resources, soliciting the help of third-party consultants and service providers, or combining the use of internal and external resources. The goal in each case is to profile current operations, quantify capacity levels and determine how effectively the school system is using existing storage resources.

There are various levels of storage assessments, but a good assessment should lead to a clear understanding of the storage environment, and provide recommendations for improvements in storage architecture and technology, backup, recovery, archiving, virtualization, utilization and management processes. Assessments look in detail at how data is being utilized, how often it is accessed and even how much data may be redundantly stored. (For more information, see the sidebar on the previous page "Assessments: A Valuable Baseline.")

Tracking Performance

Performance monitors built into storage area network (SAN) controllers can gather much of the information needed. For each class of hard drive, IT managers can determine the performance rates of their storage resources by comparing actual numbers to industry averages published by organizations such as PassMark Software[1] and Tom's Hardware.[2] Similarly, the information will show which resources are running at near capacity, as well as those that are underutilized because of ineffective load-balancing procedures. In many cases, the results may be surprising and could lead to optimization or cost-saving strategies.

For example, when American University ran numbers it gathered from a legacy SAN, it found the system was running at about 56 percent of capacity. In an unexpected resourcing paradox, there wasn't any available space for new data, even though only a little more than half of the theoretical capacity was actually housing files. How could this be? Much of the "capacity" was unused because it was needed solely to maintain performance levels. "We couldn't put anything on those drives in fear of slowing them down," says Chris Gehring, senior director of technology operations at American University, which manages about 300 terabytes of storage capacity.[3]

When Gehring and his staff compared the theoretical to the actual performance the school was deriving from its SAN, they came up with equally troubling numbers. "We were getting less than 40 percent of performance

When American University ran numbers it gathered from a legacy SAN, it found the system was running at about 56 percent of capacity. In an unexpected resourcing paradox, there wasn't any available space for new data, even though only a little more than half of the theoretical capacity was actually housing files.

utilization out of our SAN because of how space was carved out," he explains. This meant that about 10 percent of the disks were running at 100 percent of capacity, while about 75 percent were idle, and roughly 15 percent were somewhere between being maxed out and neglected.

Utilizing the Data

American University found flaws in how storage capacities were allocated for individual applications. In particular, IT administrators had ranked what were considered to be the top three systems with the most stringent storage requirements — the database sat at the top, followed by the school's ERP application and then the email platform. However, after gathering and analyzing performance numbers, the school found that email was the most demanding resource in terms of storage performance needs while the database and ERP solutions no longer even ranked in the top three.

The university then used the data to create requests for proposals that it distributed to storage vendors when it solicited ideas for new technologies. Gehring indicated that the university would tell potential partners, "This is our landscape today and here is what we think we are going to deploy in the future. Show us the system that could meet these requirements." He and his staff then asked vendors to explain how the university could change allocation parameters in their solutions, using the actual data as a reference. "It was a huge eye opener," he says. "Depending on the system, changing the landscape could be a huge challenge — and in some cases, not even possible." With the help of one primary storage vendor, American University used detailed performance data to design an environment that accommodates significantly higher data volumes while increasing utilization rates and reducing management overhead.

Nearly any organization can use similar data to optimize performance and utilization. Once an organization completes its storage system profiles and documents infrastructure characteristics, it has a strong foundation to identify current gaps in capacity and performance and estimate future requirements. The overview data will provide guidance for addressing gaps and setting the stage for future success. Options may include centralizing storage, consolidating disks and storage arrays, and using the latest automated management tools to balance available capacity and achieve practical drive utilization levels.

Despite the seemingly insatiable demand for more capacity, with the right storage systems, K-20 campuses can actually reduce their storage resources. The key is using available technologies that both optimize overall IT performance and enhance utilization levels of the underlying storage infrastructure.

Step 2: Consider a Storage Refresh

Despite the seemingly insatiable demand for more capacity, with the right storage systems, K-20 campuses can actually reduce their storage resources. The key is using available technologies that both optimize overall IT performance and enhance utilization levels of the underlying storage infrastructure.

Reduce Complexity and Cost with Centralization

Campuses should analyze how they implement storage resources across their entire organization.

- Are storage resources decentralized, so that each department essentially maintains its own systems?
- Are the resources direct attached storage (DAS), where PCs and servers use hard drives dedicated solely to that department and are expanded or swapped out to keep pace with rising capacity demands?

While DAS is a traditional and accepted management technique, it may no longer be flexible or cost effective enough in today's era of ballooning digital data. The DAS approach comes with significant management overhead required to keep all the disparate devices running. This often results in an impractical combination of unused and over-provisioned drives, where opportunities for sharing

storage capacity could be missed. In addition, adequate data back-up and recovery programs as well as proper security protocols are difficult, if not impossible, to maintain in a fragmented storage environment.

Centralized networked attached storage (NAS), SANs and virtualized storage solutions overcome these drawbacks and provide the cost-effective flexibility sought by most campuses today. With modern, affordable technology and tools that don't require specialized expertise, centralized SAN and NAS solutions are an appropriate choice for organizations of all sizes.

As DAS systems give way to networked SANs and NAS systems, campuses create consolidated pools of storage to meet the changing needs of end users. For example, Keller Independent School District — one of the fastest growing school districts in Texas — found that implementing new technology in education was putting severe pressure on its existing storage environment. "Teachers couldn't get their email, and we couldn't make files available for students," says Joe Griffin, chief technology officer for Keller ISD.[4]

As a result, Keller implemented a SAN solution that would scale readily and deliver greater performance and availability to its users. Now, adding new servers no longer results in additional disk costs. "With six servers running with less expensive storage hardware, we saved over $6,000 in a year," Griffin explains. "We're streaming video with no loss in performance, and we don't have to take servers down now. We just plug them in, add storage and we're done."

"With six servers running with less expensive storage hardware, we saved over $6,000 in a year," Griffin explains. "We're streaming video with no loss in performance, and we don't have to take servers down now. We just plug them in, add storage and we're done."
- JOE GRIFFIN, CHIEF TECHNOLOGY OFFICER, KELLER INDEPENDENT SCHOOL DISTRICT

What to Look for in a SAN

- **Determine the appropriate SAN technologies for your needs** — Implement hybrid multi-disk arrays and multi-protocol arrays — all of which help avoid performance bottlenecks and production downtime. Just as important, the arrays should be able to automatically perform many other essential management activities such as configuring RAID and load balancing data across spindles and arrays to help optimize performance. The right SAN can reduce complexity even as storage volumes continue to rise. "My storage administrators say that while we've doubled the number of SANs and quadrupled the amount of storage they need to manage, we've reduced their time managing it by somewhere around 300 percent," says Gehring, senior director of technology operations at American University.

- **Evaluate storage arrays on the quality of the software** — Automated management tools can eliminate the need to have a specialized storage administrator dedicated just to provisioning space. With the most advanced arrays, an administrator only has to provision how much space is needed and the SAN determines where to store the data in the virtual resource pool. The software can also find ways to increase performance. Instead of placing 10 terabytes of information on 10 individual 1-terabyte

drives, the tool may determine that access times will be shorter if it stores pieces of the file across 30 drives, which increases the number of spindles for faster input/output operations.

- **Consider built-in storage virtualization technology** — These capabilities will help campuses further improve performance and system efficiency. As with other virtualized technologies, storage virtualization works by cutting direct ties between physical devices and individual storage volumes, which further leverages the resource pooling efficiencies associated with SANs and NAS units. With some arrays, virtualized storage allows pools to span multiple disk trays or even storage arrays, which provides more options for dynamically allocating capacity on the fly to better balance workloads. The result is greater flexibility through faster performance and optimized utilization rates. To achieve these benefits without added management overhead, look for storage solutions that offer flexible, virtualized storage features that simplify administration and enhance performance.

  Done right, virtualized storage can have a significant impact on overall storage costs. The technology advisory firm Enterprise Storage Group estimates that some virtualized SANs can reduce total cost of ownership by 50 percent.[5] Keller ISD experienced similar benefits. "We had 100 servers in a room that was literally out of space. The combination of virtualized storage and servers shrunk our data center by 66 percent — reducing power costs and providing room to grow," says Vaughan Hamblen, Keller ISD's director of network services.

- **Ensure scalability for the future** — Ultimately, choose a SAN that is flexible enough to meet your campus' needs in the future. It is important to look for a solution that doesn't require a "forklift" upgrade down the road. When upgrading, there should be minimal downtime so that the campus is not negatively affected. Look for a solution with separate maintenance fees for software, or where software can be updated while still under warranty. The warranty structure should be simple so that the solution doesn't "break the bank" when it comes time to renew.

"My storage administrators say that while we've doubled the number of SANs and quadrupled the amount of storage they need to manage, we've reduced their time managing it by somewhere around 300 percent."

— CHRIS GEHRING, SENIOR DIRECTOR OF TECHNOLOGY OPERATIONS, AMERICAN UNIVERSITY

Increase Efficiencies with Storage Tiers

Advanced SANs and storage virtualization become even more effective when they are combined with automated storage tiering technology, an option that can add cost efficiency to data management activities while helping keep pace with constantly changing performance requirements.

Industry studies show that about 90 percent of data that organizations create is typically only accessed during a short period immediately after being created. In fact, large percentages of data are called upon only once and are never touched again. Storage tiering applications automatically take these realities into account and assign frequently

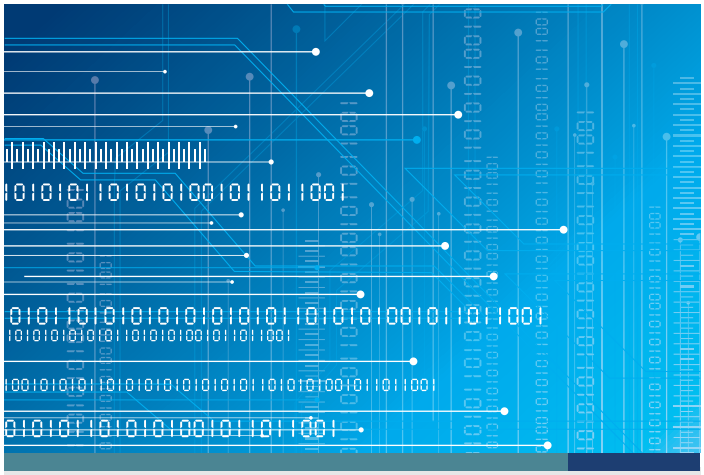

3 New Storage Technologies That Deserve Consideration

Thin provisioning — Allows IT managers to establish space allocation thresholds for services and applications. These resources then automatically get the capacity they need from the virtual pool up to a predetermined limit, all without having to carve out extra — and perhaps unavailable — space for possible future allotments. Keller ISD achieved optimal scalability thanks to thin provisioning and the ability to provision capacity as needed without having to pre-allocate disk space. "Now we can grow without micromanagement, unused capacity or downtime," says Griffin.

Data deduplication — Technologies that reduce the amount of storage capacity required to house unique data by automatically identifying and eliminating redundant data on either a file, block or bit level. A good deduplication system can assure that only one copy of an attachment ends up on a hard drive and uses slim, storage-sparing pointers to refer back to that file.

Data compression — When combined with deduplication, file-shrinking data compression technology reduces file sizes even more, which saves on hard disk investments and alleviates strains on networks by enabling them to move less data throughout the pipeline.

used files to high-performing devices (Tier 1 environments), such as 15,000 rpm Serial Attached SCSI (SAS) drives or solid state drives (SSD) when the highest throughput rates are required. Depending on usage patterns, only about 10 to 20 percent of the total storage capacity includes data that is accessed frequently enough to justify a home on SAS and solid state drives.

By constantly monitoring usage patterns, automated tiering can also identify formerly front-line information that's being called on less frequently and is appropriate for migrating to a second tier (Tier 2) of less expensive, moderate performance drives, such as 10,000 rpm SAS hardware.

Similarly, data that campuses rarely access but that must be retained for recordkeeping, such as student grades or information for regulatory compliance, move to a Tier 3 environment for archiving, which consists of the slowest and least expensive devices — such as SATA drives. Thus, with automated tiering in place, an individual file may start out on speedy Tier 1 disks and gradually work its way down the other two tiers as the information is used less frequently.

A tiering strategy can pay huge dividends for optimization. For example, because email systems require high performance from the underlying IT infrastructure, organizations often store all their mail files on the fastest disks. However, as much as 80 or 90 percent of that data may be old messages and attachments that aren't being viewed any longer, but nevertheless aren't being deleted. Tiering technology can identify the idle files and move them to inexpensive SATA drives, leaving most of the space on the high-performance disks open and available. "We find that when we need to do upgrades, we can buy additional cheaper disks as opposed to buying additional expensive disks," American University's Gehring says.

He also credits automated storage tiering with giving the university the flexibility it needs for its storage array to meet today's new data management challenges. He contrasts this flexibility with the demands of a legacy storage array that required the IT staff to make detailed upfront decisions about performance levels, the number of required disks and capacity allotments.

"In a traditional storage world, you have to go in and say, 'I'm going to carve off this piece of the SAN so that this application gets what it needs,'" Gehring explains. The downside was that those allocation decisions had to be made before the organization launched a new service, such as an additional email server. And the only way to judge how accurately the IT staff predicted allocation requirements was to gather performance data over a series of months and then analyze the performance information. This rearview mirror approach to storage optimization is becoming less and less practical as volumes balloon and performance requirements become stricter.

Contrast this with automated storage tiering, where embedded algorithms constantly monitor performance and usage rates and make appropriate adjustments down to sub-LUN (logical unit number) block levels (which allows a more granular control of data movement while maintaining a consistent LUN view) to constantly balance loads across the entire system. "With a true virtual storage environment we just carve out what we need for the system and the SAN makes sure that the data goes to the right tier and then moves the data if the environment changes," Gehring says.

The result is that large groups of disks aren't sitting idle while others are running at maximum capacity and delivering slow performance as a consequence. The university also runs technology that enhances automated tiering by dynamically moving the most frequently accessed data to the outer edges of the disks, which speeds access because the head that reads stored data needs to move smaller distances.

This level of control enabled American University to significantly reduce capital investments for storage systems. "When we moved to a SAN with automatic balancing we actually deployed a system that had half the number of disks than with legacy hardware, which traditionally would have meant half the performance," Gehring says. "But since we were only getting 40 percent performance out of the original SAN, we knew that the smaller number of disks would take us to where we wanted to go, and we were almost spot on."

Thanks to the latest storage technology, American University can restore a virtual server in its business continuity site in 30 minutes or less, for a 48-fold improvement in the university's recovery time objectives.

Bolster Reliability with Data Protection and Disaster Recovery Strategies

Many K-20 campuses — even some of the largest universities — maintain relatively small IT staffs that are already burdened with keeping the primary technology environment operating reliably. Thus, once the SAN is implemented, business continuity capabilities, such as data replication, should be automated as much as possible. The goal is to set up the replication system so that when a new file is stored on the primary system, a copy automatically appears in the backup resources.

Built-in support for RAID allows IT to store copies of critical data over a number of individual drives to provide fault-tolerance and increase I/O performance. SANs should also offer redundant internal components and enable swapping out faulty disks, controllers, power supplies and fans without storage system downtime.

Best practices in data protection also call for organizations to maintain duplicate storage environments at a

location geographically removed from the main data center to protect information against natural or man-made disasters that destroy or incapacitate the main facility. For example, American University uses replication technology to instantly mirror its full production environment to an offsite business continuity cluster that runs hardware identical to the production cluster, except for SSDs. The legacy system the university replaced had relied on tape-based disaster recovery, and bringing a system back to life often took 24 hours, Gehring says. Thanks to the latest storage technology, it can restore a virtual server in its business continuity site in 30 minutes or less, for a 48-fold improvement in the university's recovery time objectives.

When designing highly available, fault tolerant storage infrastructures, first set recovery goals for each application and data type based on how critical the individual resource is and how quickly the information will need to be recovered. Mission-critical systems with immediate recovery scenarios will justify spending for full redundancy and the highest performing solutions. For less critical systems, educational institutions may cope with having resources unavailable for hours, days or even weeks without harming operations, and can be served with more economical solutions. Paramount to any disaster recovery plan is the ability to create a regular schedule for testing and uncovering any gaps in the event a real crisis occurs.

Ensure Security and Compliance

The best storage systems work with complementary security technologies to keep data secure and tightly managed. In day-to-day operations, this means having the right tools in place to comply with government regulations and records retention rules. Advanced storage systems offer built-in encryption technologies to protect data when it's sitting in hard drives or traveling across network links to and from end users.

Automation tools for efficient archiving help campuses conform to state and federal data retention policies by applying information lifecycle management (LCM) rules set by IT administrators. LCM automatically moves primary data files from production systems to disks and tape libraries set up for long-term retention and preservation. The auditing software that can be built into storage systems maintains a constant record of who accessed data files and when — and if any edits were made, which helps campuses demonstrate the effectiveness of their data retention policies for government regulators.

Fortunately, the challenge of delivering cost-effective storage solutions is getting easier thanks to a new generation of technologies, emerging best practices and consulting expertise designed for today's digital education environment.

In addition, write once, read many (WORM) capabilities assure that records aren't altered after they're created — a benefit particularly important for universities and libraries to preserve and protect their research data and digital collections. Campuses must also have eDiscovery tools to dig through hundreds of terabytes of stored electronic records to quickly find data relevant to legal proceedings.

Step 3: How Service Providers Can Help

Designing and implementing a high-performance, cost-effective and reliable storage infrastructure requires the right mix of technologies and architectures. Rather than

following a well-planned IT roadmap, campuses are often forced to build out their IT environments, including storage systems, based on the types and timing of the grants they receive. Using this approach, institutions may acquire a variety of disparate technical solutions over time, only to realize the components have not been integrated in a way that allows efficient interoperability. Instead, the campuses try to manage disparate technology groups and struggle with evolving information management requirements. Fortunately, school systems don't have to develop winning IT strategies on their own — storage consultants are available to provide systems design and implementation expertise.

Look for storage implementation providers with expertise in streamlining each step of the implementation process, from the upfront design and project planning phases to the actual installation and provisioning of resources.

Conclusion: Storage, Front and Center

Storage systems are no longer behind-the-scenes IT resources important only to the technical experts running the campus data centers. Whether the administrative staff and students know it or not, their productivity and ability to experience the full impact of today's latest technologies hinge on how well disk drives, SSDs, arrays and other elements of the storage infrastructure support and fulfill performance and data demands. Fortunately, the challenge of delivering cost-effective storage solutions is getting easier thanks to a new generation of technologies, emerging best practices and consulting expertise designed for today's digital education environment. ■

Is Cloud Storage the Answer?

There's a lot to like — and a lot to consider — when it comes to moving storage infrastructures to the cloud. Many campuses will want to take a slow, incremental approach to evaluating cloud storage and whether a move makes sense for their organization.

Cloud Benefits

What's good about the cloud? First, efficient cloud technologies can provide the potential for school systems to reduce capital expenses devoted to their storage environment. By using cloud resources to contract for pay-as-you-go capacity funded by operating expenses, administrators purchase only what they need when they need it without sinking funds into expensive and depreciating capital equipment. That cost control flexibility delivers significant advantages in an era of exploding capacity requirements that would ordinarily push organizations to regularly provision new drives and storage arrays.

The cloud also relieves storage management and maintenance burdens for internal IT staff, who can then focus more time on strategic initiatives that improve education outcomes and teaching/learning activities.

Add to that the expertise of cloud providers that specialize in storage systems who can offer high levels of availability while containing costs. The result is that campuses can more easily meet peak load requirements at the beginning and end of semesters, while seeing higher service levels from cloud-based storage systems for equivalent or less cost.

Cloud Concerns

For most organizations, the biggest concern around cloud technologies is security. Shared (public) cloud environments operate under the principle of multi-tenancy, where cloud subscribers use the same storage systems, servers, applications, databases and other IT resources. While this model keeps costs affordable, organizations must carefully evaluate a potential service provider's security methodologies. The ability to securely partition shared resources is possible, but it happens only when providers have the right systems and staff in place to eliminate the risk of data breaches. Alternatives include a private cloud environment on the campus' premises, or only using a shared cloud for low-sensitivity storage, such as email back-up or long-term archiving of rarely accessed databases. It is also worth noting that in many instances, sophisticated service providers can actually provide a more comprehensive overall security program than a school system.

Campuses also must negotiate service level agreements that assure cloud providers will deliver needed performance and availability levels, and have detailed remediation plans in place in case service falls short of agreed-upon thresholds.

Finally, because storage technology continues to evolve and improve, campuses should understand a cloud provider's technology roadmap and receive regular updates showing how new innovations are helping improve performance.

A Storage Glossary

New technologies, tools and practices are helping campuses improve storage system performance and optimize resources.

Data compression — file-shrinking technology that helps reduce storage capacity requirements.

Logical unit number (LUN) — a unique identifier that is used to name individual or collections of hard disk devices addressed by the SCSI protocol or other similar interface.

Networked storage — an alternative to hard drives attached directly to servers and PCs; this approach creates consolidated pools of storage capacity for fast resource provisioning of storage space.

Data deduplication — process of identifying and eliminating redundant data to reduce overall capacity needs and relieve stress on backup and networking systems.

Solid state drives (SSDs) — devices that behave similar to hard drives, but are built on solid state memory enabling higher transactional performance and lower latency.

Storage area network (SAN) — a networked pool of disks that enables centralized administration of storage resources; often involves high performance and reliability technologies, including redundant and "hot swappable" components.

Storage tiering — tools that constantly monitor how often data is being accessed and automatically moves that data to the most appropriate storage tier based on these usage patterns. Rarely used data is housed on the slowest and least expensive devices.

Storage virtualization — technology that separates physical devices from individual storage volumes and makes it possible for IT managers to create pools of storage that span multiple storage tiers.

Thin provisioning — separates storage allocation from storage consumption eliminating the need to over-purchase and pre-allocate disk space. Storage space can be assigned to applications, but the space is not consumed until data is actually written to the drives.

Endnotes

1. http://www.harddrivebenchmark.net/
2. http://www.tomshardware.com/charts/hard-drives-and-ssds,3.html
3. Comments from Chris Gehring, American University, during a Center for Digital Education interview on December 9, 2011.
4. All information about Keller Independent School District from Dell Case Study, "Keller ISD exceeds cost, performance standards with Dell Compellent Storage Center SAN."
5. Dell EqualLogic TCO Analysis: The Economics of EqualLogic Virtualized iSCSI Storage

Dell is committed to helping students develop the knowledge and skills they need to learn and succeed in an increasingly digital world and a globally competitive workforce. As a top provider of technology and services to schools, Dell listens to and works with students, educators, administrators, communities and industry partners around the world to deliver innovative technology and services that give them the power to do more.

Learn more about Dell in Education at www.dell.com. Follow us on Twitter @DellEDU .

CENTER FOR DIGITAL EDUCATION

The Center for Digital Education is a national research and advisory institute specializing in K-12 and higher education technology trends, policy and funding. Along with its research services, CDE issues white papers and conducts the annual Digital School Districts and Digital Community Colleges surveys and award programs as well as hosting events across the K-12 and higher education arena. CDE also supports the Converge media platform comprised of the quarterly themed Converge Special Reports, Converge Online, and custom publishing services.

www.centerdigitaled.com

Acknowledgements:

Alan Joch specializes in technology best practices for the public sector, education, and industry. Areas of expertise include cloud computing, mobile applications, server and desktop virtualization, and enterprise storage. His feature articles appear in *The New York Times, Federal Computer Week, Engineering Inc.*, and other industry publications. Previously, Alan spent seven years as a senior editor at *Byte Magazine*, where he alternately ran the product testing lab and the features department. He also is author of the book "How to Find Money Online: An Internet-Based Capital Guide for Entrepreneurs."