

An interactive eBook

STORAGE FOR TODAY AND TOMORROW

No matter what way you look at it, the volume of data needing storage in the enterprise data center is skyrocketing. Research firm IDC tracks the storage growth at 60% annually – a hard figure to cope with for even the savviest of storage administrators. Smart storage starts with an understanding of an enterprise's needs matched to the right technologies for meeting those demands. But that's easier said than done, given changing enterprise data types, regulatory provisions surrounding data storage and the rise of storage options. In these articles, *InfoWorld* and its sister publications *CIO*, *Computerworld* and *Network World* explore the latest storage challenges and the technologies for addressing them.

Custom publishing from

InfoWorld
GET TECHNOLOGY RIGHT

2. Struggling with Supersized Storage

Shrinking budgets and swelling data demands send IT managers on a search for more efficient storage.

5. How to Avoid 5 Common Storage Mishaps

Blindsided! These companies thought they had their stored data locked tight, but they were wrong. Here's how you can avoid a similar fate.

10. QuickStudy: Storage Virtualization

Managing disk storage was once simple: If we needed more space, we got a bigger disk drive. But data storage needs grew, so we started adding multiple disk drives. Finding and managing these became harder and took more time, so we developed RAID, network-attached storage and storage-area networks. Still, managing and maintaining thousands of disk drives became an ever more onerous task.

12. Managing the Complexities of Storage Virtualization

Storage virtualization is hot, and for good reason. But its benefits bring added layers of complexity.

16. Going Virtual Raises Storage-Management, Procurement Issues

Five things to think about before virtualizing storage

19. Thinking Outside the Storage Box

Unbridled growth in data storage and the rise in Web 2.0 applications are forcing a storage rethink. Is this the end of the storage-area network as we know it?

Sponsored by



InfoWorld

Struggling with Supersized Storage

Avoid 5 Common Storage Mishaps

QuickStudy; Storage Virtualization

Managing Complexities of Storage Virtualization

Going Virtual Raises Storage-Management, Procurement Issues

Thinking Outside the Storage Box

Analysis

Struggling with Supersized Storage

Shrinking budgets and swelling data demands send IT managers on a search for more efficient storage. **By John Edwards**

TASTY BAKING CO. produces more than 4.8 million cakes, doughnuts, cookies and pies each day. And the Philadelphia-based snack food giant also manages to generate another crucial commodity — computer data — in equally impressive amounts. “It’s something to be concerned about,” says Brendan O’Malley, Tasty Baking’s vice president and CIO.

Like a growing number of businesses worldwide, Tasty Baking is facing a data explosion. The use of e-mail and rich media ap-

plications plus the need to stay on top of regulatory compliance are stretching storage resources to the limit at a time when budgets are shrinking. Handling rapidly spiraling storage needs without spending tons of money is a challenge that’s facing just about all IT managers. O’Malley crystallizes the need into a single phrase: “More space for the money.”

To get a handle on storage demands, IT managers need to carefully examine internal storage practices, use specialized soft-

ware tools and find appropriate storage systems for various kinds of data. Andrew Reichman, an analyst at Forrester Research Inc., says that they must look first at the two best ways of cutting storage costs. “One is using more-dense drives, and the other is reducing your footprint,” he says.

Footprint reduction is relatively easy, says Reichman, asking, “Can you use the Delete key?” But David Hill, founder of The Mesabi Group, a storage management consultancy in Westwood, Mass., notes that many businesses, particularly those in regulated industries, are scared stiff of deleting anything other than spam. “There are a lot of legal reasons,

tax reasons,” he says. “You have to be careful about how you go about getting rid of data.”

Fixed vs. Fluid

Although managers may be reluctant to flush files, they can bring greater efficiency and cost savings to their storage processes by organizing data more efficiently. Hill recommends separating fixed content — e-mails, boilerplate documents and other rarely modified files — from active archives that contain more fluid types of data, such as transactional data.

Hill notes that many businesses unthinkingly use costly, high-performance Fibre Channel storage environments for both types of

data. Yet splitting up fixed and active data makes it possible to use less-expensive storage technologies for less volatile information. “You can use what’s called capacity storage instead of performance storage,” Hill says.

Capacity storage usually comes in the form of Serial ATA drives, which have higher capacity but cost less than other options. “So you get more bang for your buck,” Hill says. “That’s one thing that businesses can do: move some of their data to these lower-cost drives.”

O’Malley agrees that organizing content is a relatively easy way to cut costs. “We have some high-performance applications that run

Sponsored by



Analysis

on Fibre Channel, and other applications that don't need that kind of performance that we run on regular SATA types of disks, where we have a better price for performance," he says.

Data de-duplication — the process of eliminating redundant files — can also save precious storage space. Backup products that de-duplicate data at the host level include Asigra's Televaulting, EMC's Avamar and Symantec's Veritas NetBackup PureDisk. Disk libraries from Data Domain, Diligent Technologies, EMC, Quantum and Sepaton de-duplicate data at the target level, such as the desktop. Both methods determine whether data segments or files are identical by running a file-level comparison or using a hashing algorithm.

Even relatively insignificant files can pile up to create a big storage headache, says Dave Richerson, manager of infrastructure services at Republic Bank in Louisville, Ky. "If we send a Word doc-

ument out, that doesn't make a huge impact on the e-mail system," he says. "But 15 people might store that Word document on their home directory as something to refer to down the road

and, well, that's 15 copies."

To create space-efficient backups without endangering critical data, Republic uses Avamar for its branch office backups and EMC Disk Library for data center back-

ups. "We've had this system in place for about four months and have seen an immediate, positive improvement in storage efficiency," Richerson says.

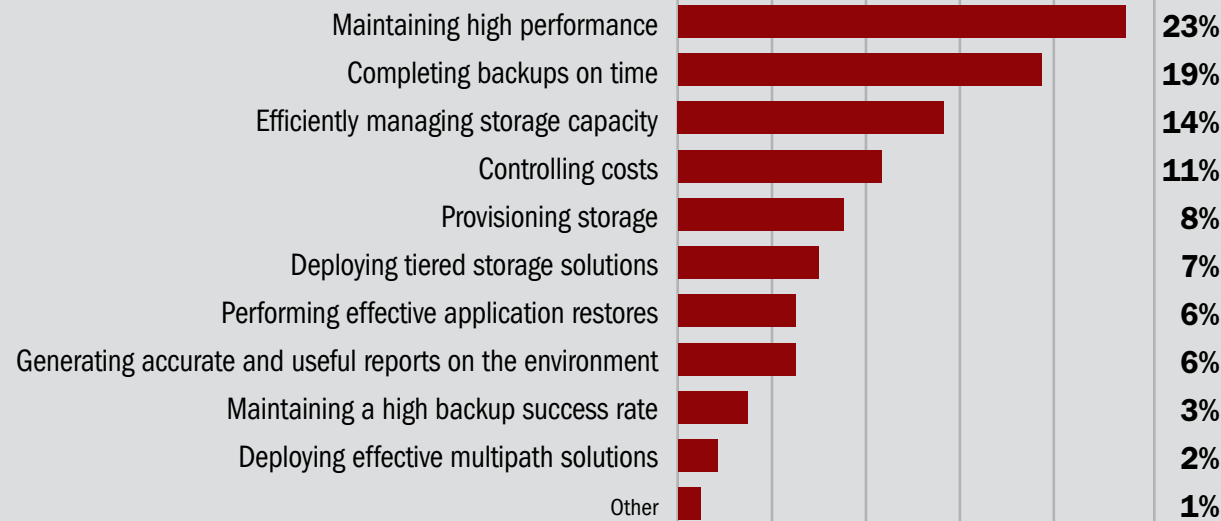
De-duplication vendors often claim that their tools can provide a specific ratio of data reduction. Actual data de-duplication ratios can vary widely from customer to customer, of course, but Richerson says he was "floored" by the ratio his deployment achieved. "We see a de-dupe of 8-to-1, or basically 24TB of data is de-duped down to 3TB," he says.

Virtual Efficiency

At the cutting edge of storage, virtualization promises to help businesses archive data with unprecedented efficiency. Greg Thompson, director of IT at North-

Virtual Storage Headaches

WHAT IS THE TOP CHALLENGE RELATED TO STORAGE FOR YOUR VIRTUAL SERVER ENVIRONMENT?



Source: Forrester Research Inc. survey of 124 global IT decision-makers currently using x86 server virtualization technology, September-October, 2008

Sponsored by



Analysis

east Bank, is relying on the Equal-Logic PS5000 Series virtualized storage-area network from Dell Inc. to store corporate data. The systems are based on a virtualized modular storage architecture that enables Lewiston, Maine-based Northeast to purchase storage as it is needed, preventing overprovisioning.

For Thompson, the new technology marked a welcome departure from local server-based storage. “Local storage on servers gets unruly,” he says. “When you run out of space, what do you do, buy another server? So we decided central storage was the answer for us.”

Thompson says he opted to tie central storage together with vir-

tualization for server and storage consolidation. “The whole process has worked really well for us,” he says, noting that storage can now be scaled easily and efficiently whenever the need arises.

Thompson also appreciates the virtualized environment’s inherent redundancy. “If we do have a critical, fatal error happening here at the primary data center, we can have access to just about all our network data within a matter of hours,” he says.

Like Thompson, O’Malley says that virtualization can give a big boost to businesses striving for better storage efficiency. Tasty Baking now uses a FAS3020 system from NetApp Inc. to provide

Fibre Channel-based storage resources to 40 virtual servers running on 10 physical servers. “We moved away from having any local storage on servers or having any sort of direct-attached storage or anything of that kind,” O’Malley says. “Now everything runs off NetApp.”

Before O’Malley and his team moved to the virtualized environment, adding storage was both a challenge and a puzzle. “If we were out of storage on a particular server, we wondered if we should find more slots where we could put more disks in, or maybe attach some storage directly to that box,” he recalls. “You might have storage space on another

server, but you couldn’t really share it over to that machine.”

The new environment allows O’Malley to consolidate and optimize storage resources. “We’ve got one place where everybody is going for storage,” he says. “With the thin provisioning, and the ability to move storage, we have one bucket that we’re pulling from, and we can get by with a lot less overhead.”

The virtualized arrangement allows for more planning flexibility. “We don’t have to determine what the maximum possible storage might be for an application over the next three years,” O’Malley says. “We don’t have to ask that question before we start and ad-

vised the server to hold that.”

Virtualization also frees storage from constraints imposed by a server’s physical configuration. “It’s now just a question of finding more storage from the single pool on NetApp,” O’Malley says.

He notes that storage efficiency and simplicity meshes well with Tasty Baking’s corporate philosophy. “When you think about it, we’re just like a corner bakery, only a thousand times bigger,” O’Malley says. “What’s key for us is keeping our environments as simple as we can.”

Edwards is a freelance writer in Gilbert, Ariz. Contact him at jedwards@gojohndwards.com.

Sponsored by



Guide

How to Avoid 5 Common Storage Mishaps

By Mary Brandel

Blindsided! These companies thought they had their stored data locked tight, but they were wrong. Here's how you can avoid a similar fate.

Think you can guess the No. 1 threat to the security of your stored data? If you said hackers, or even trouble-making insiders, you'd be wrong.

While malicious threats are an ongoing concern, it's your well-meaning employees who are more likely to unknowingly expose your company's stored data through, say, a file-sharing network or a misplaced laptop.

In fact, a recent Ponemon Insti-

tute study found that negligent insiders are by far the biggest threat to data security, accounting for 78% of all breaches.

In this special report, you'll learn the latest techniques for protecting stored data within company walls as well as stored data that flows freely in and out of your organization on laptops, tapes and other movable media.

Data breaches, unfortunately, have become a way of life for corporate America. According to the Identity Theft Resource Center (ITRC), 2008 saw a 47% increase in documented data breaches from the year before. And those are just the ones that made the news, says Craig Muller, an identity theft expert and founder of Identity Doctor in Irvine, Calif. "I

get e-mails constantly telling me of breaches," he says.

The public is definitely feeling the pain. In a 2008 study by the Ponemon Institute in Traverse City, Mich., over half (55%) of 1,795 adult respondents across the U.S. said they'd been notified of two or more data breaches in the previous 24 months, and 8% said that they'd received four or more notifications.

But companies are still not sure how to protect themselves. In a Ponemon survey released last month, only 16% of the 577 security professionals who responded said that they were confident or very confident that current security practices could prevent the loss or theft of customer or employee data.

One way to gain confidence is to examine actual breaches and learn from them. Here's a look at five com-

mon types of breaches, with advice about how to avoid similar mishaps.

1. Stolen Equipment

In May 2006, personal data on 26.5 million veterans was compromised when a laptop and a storage disk were stolen from the home of a subcontractor working for the U.S. Department of Veterans Affairs. Both items were recovered, and arrests were made. The FBI claimed that no data had been stolen, but the incident prompted sweeping reform at the VA. However, in January 2007, another breach occurred when a laptop was stolen from an Alabama medical facility, exposing personal data on 535,000 veterans and more than 1.3 million physicians.

COSTS: By June 2006, the VA was burning through \$200,000 a day

Sponsored by



Guide

to operate a call center to answer questions about the breach. It also spent \$1 million to print and mail notification letters. It was given permission to reallocate up to \$25 million to pay for those costs. Class-action lawsuits were also filed, including one demanding \$1,000 in damages for each person affected. After the 2007 breach, the VA set aside an additional \$20 million for breach-related costs. And the department recently agreed to pay \$20 million to current and former military personnel to settle a class-action lawsuit.

BLINDERS: Lost or stolen equipment accounts for the largest portion of breaches – about 20% in 2008, says the ITRC. According to Bart Lazar, a partner in the Chicago office of law firm Seyfarth Shaw LLP, incidents

involving lost or stolen laptops make up the majority of data-breach cases he works on.

EYE-OPENERS: Lazar recommends restricting the placement of personal identifying information on laptops. For instance, don't tie customer or employee names to other identifiers, such as Social Security or credit card numbers; alternatively, you can truncate those numbers. Also, consider creating your own unique identifiers by, for example, combining letters from an individual's last name with the last four digits of his Social Security number.

Second, require personal information on laptops to be encrypted, despite the potential cost (\$50 to \$100 per laptop) and performance hit that involves, says Lazar. This needs to be accompanied by con-

sciousness-raising, says Blair Semple, storage security evangelist at NetApp Inc. and vice chairman at the Storage Networking Industry Associ-

ation's Storage Security Industry Forum. "I've seen situations where people had the capability to encrypt but didn't," he says. "Scrambling the bits

is the easy part; it's the management and deployment that's hard."

Third, Lazar recommends policies requiring very strong passwords to protect data on stolen devices.

2. Insider Theft

In November 2007, a senior database administrator at Certegy Check Services, a subsidiary of Fidelity National Information Services, used his privileged access to steal records belonging to more than 8.5 million customers. He then sold the data to a broker for \$500,000, and the broker resold it to direct marketers. The employee was sentenced to over four years in jail and fined \$3.2 million. According to company officials, no identity theft occurred, although affected consumers received marketing solicitations from the companies that bought the data.

Breaches on the Rise

Since 2006, the number of documented data breaches* has risen by over 40% annually.

	Documented breaches	Records exposed
2006	315	20 million
2007	446	128 million
2008	656	36 million

**To qualify, breaches must include personal identifying information that could lead to identity theft, especially the loss of Social Security numbers. Five categories of data-loss methods were tracked, including breaches of data on the move, accidental data exposure, insider theft, subcontractor breaches and hacking.*

SOURCE: Identity Theft Resource Center, San Diego

Sponsored by



Guide

In another high-profile case, a 10-year veteran scientist at DuPont downloaded trade secrets valued at \$400 million before leaving the company in late 2005 to join a competitor in Asia. According to court records, he used his privileged access to download about 22,000 document abstracts and view about 16,700 full-text PDF files. The documents covered most of DuPont's major product lines, including some emerging technologies. The scientist did this while in discussions with the competitor and for two months after accepting the job. He was sentenced to 18 months in federal prison, fined \$30,000 and ordered to pay \$14,500 in restitution.

COSTS: In DuPont's case, the estimated value of the trade secrets was more than \$400 million, although

the government pegged the company's loss at about \$180,500 in out-of-pocket expenses. There was no evidence that the confidential information was transferred to the competitor, which cooperated in the case.

According to Semple, theft of customer information is nearly always more costly than theft of intellectual property. In Certegy's case, a 2008 settlement provided compensation of up to \$20,000 for certain unreimbursed identity theft losses for all class-action plaintiffs whose personal or financial information was stolen.

BLINDERS: Nearly 16% of documented breaches in 2008 were attributed to insiders, says the ITRC; that's double the rate of the year before. One reason for this increase is that employees are being recruited

by outsiders with ties to crime — a trend that accounts for half the insider crimes committed between 1996 and 2007, according to the CERT Coordination Center at Carnegie Mellon University.

Insiders commit crimes for two reasons, CERT says: financial gain (as in the Certegy case) and business advantage (as in the DuPont case). In the latter, criminal activities usually start when the employee resigns, CERT says, but the thefts typically occur after they depart, having left secret access paths to the data they want.

Insider threats are among the hardest to manage, Semple says, especially when the workers use privileged access.

EYE-OPENERS: A good precaution is to monitor database and network

access for unusual activity and set thresholds representing acceptable use for different users, CERT says. That makes it easier to detect when an employee with a particular job designation does something beyond his normal duties. For instance, DuPont discovered the illegal activity because of the scientist's unusually heavy usage of its electronic data library server.

If you suspect that a breach has occurred, CERT says it's important to act quickly in order to minimize the chance of information being disseminated and to give law enforcement agencies a chance to start investigating the case.

Companies should also implement role-based access-control tools to maintain a high level of accountability over who is accessing valuable assets, Lazar says. Databases con-

taining customer or employee information should allow very limited access. "How many people, on a daily basis, need to review Social Security numbers and addresses without permission?" he says. "Personal information should be protected at the same level as trade secrets."

Muller recommends using data loss prevention tools to restrict personal data from being e-mailed, printed or copied onto laptops or external storage devices. Some of these tools provide alerts that inform administrators when someone tries to copy personal data and create a log file of such an event. "In a lot of cases, companies don't have proper audit trails in place," he says.

It's also important to strengthen internal controls and audit measures by, for example, implementing iterative checks on network and database

Sponsored by



Guide

activity logs, Semple says. It's not enough to keep detailed logs; you also need audit measures in place to see if anyone has modified a log or illegally accessed it. "Unless there's some way to verify the log information wasn't tampered with, it's hard to know it's of value," he says.

But in the end, technology isn't enough. "You need to find a way to ensure users you trust are worthy of that trust," Semple says.

3. External Intrusion

In January 2007, retailer The TJX Companies Inc. reported that its customer transaction systems had been hacked. The intrusions — which occurred between 2003 and December 2006 — gave hackers access to 94 million customer accounts. Stolen information was found to have been used in an \$8 million gift-

card scheme and in a counterfeit credit card scheme. In the summer of 2008, 11 people were indicted on charges related to the incident, which was the largest hacking and identity theft case the U.S. Department of Justice has ever prosecuted.

COSTS: TJX has estimated the cost of the breach at \$256 million. That includes the cost of fixing computer systems and dealing with litigation, investigations, fines and more. It also includes payments to Visa (\$41 million) and MasterCard (\$24 million) for losses they incurred. The Federal Trade Commission has mandated that the company undergo independent third-party security audits every other year for the next 20 years.

However, others expect that costs may rise to \$1 billion, which would include the costs of legal settlements

and lost customers. According to an April 2008 Ponemon study, 31% of a company's customer base and revenue source terminates its relationship with an organization following a data breach. And in its recently released annual "Cost of a Data Breach" study, Ponemon found that breaches cost companies \$202 per compromised customer record last year, compared with \$197 in 2007. Costs associated with lost business opportunities represented the most significant component of the increase. The average cost of a data breach in 2008 was \$6.6 million, compared with \$6.3 million in 2007.

BLINDERS: According to a 2008 Ponemon study, data breaches by hackers rank a distant fifth in terms of security threats. Indeed, about 14% of documented breach-

es in 2008 involved hacking, according to the ITRC. That doesn't mean companies shouldn't be wary, however. In TJX's case, hackers infiltrated the system by "war driving" and hacking into the company's wireless network. TJX was using subpar encryption, and it had failed to install firewalls and data encryption on computers using the wireless network. This enabled the thieves to install software on the network to access older customer data stored on the system and intercept data streaming between handheld price-checking devices, cash registers and the store's computers.

EYE-OPENERS: According to Muller, the WEP encryption that TJX used on its wireless network was insufficient — weaker even than what many

home users have. "If from the parking lot you can gain access to the database, you need a higher level of data security and data encryption," he says. TJX had also stored old account information instead of permanently deleting it, Muller says.

4. Negligent Employees

The spouse of a telecommuting Pfizer Inc. employee installed unauthorized file-sharing software on the worker's company laptop, enabling outsiders to gain access to files containing the names, Social Security numbers, addresses and bonus information of about 17,000 current and former Pfizer employees. An investigation revealed that about 15,700 people had their data accessed and copied by people on a peer-to-peer network, and another 1,250 may have had their data ex-

Sponsored by



Guide

posed. Because the system was being used to access the Internet from outside of Pfizer's network, no other data was compromised.

COSTS: Pfizer contracted for a "support and protection" package from a credit-reporting agency, which includes a year's worth of free credit-monitoring service for those affected and a \$25,000 insurance policy covering costs that individuals might incur as a result of the breach.

BLINDERS: Careless insiders — not malicious ones — are the No. 1 threat to data security, according to a recent Ponemon study, in which IT professionals said 88% of all breaches involved negligent insiders. "If there were more employee awareness about security, the num-

ber of breaches would come way down," Muller says. In Pfizer's case, the employee's spouse had configured the software so that other users of the file-sharing network could access files the spouse had stored on the laptop, but that gave people access to Pfizer files, too.

Combine negligent users and file-sharing software, and you've got a dangerous mix. Although most companies have outlawed P2P file sharing on their corporate networks, according to a 2007 study by Dartmouth College, many employees install it on their remote and home PCs. The study found, for example, that employees at 30 U.S. banks were sharing music and other files on peer-to-peer systems and inadvertently exposing bank account data to potential criminals on the network. Once business data is ex-

posed, it can spread to dozens of computers around the world.

EYE-OPENERS: First off, IT needs to either ban P2P software entirely or set policies for P2P usage and implement tools to enforce those policies. "[Pfizer] should have done a better audit of their systems to stop employees from loading any software," Muller says. "You can take away their admin rights so they can't install anything." Also important is training, he says, so users understand the dangers of P2P, what makes a good password and other standard security practices.

"There's a huge need for education so employees understand we're not trying to make things difficult but that bad things could happen," Semple notes. "It's having them understand, 'I can't do this, and here's why.' "

5. Subcontractor Breaches

In November 2008, the Arizona Department of Economic Security had to notify families of about 40,000 children that their personal data may have been compromised following the theft of several hard drives from a commercial storage facility. The drives were password-protected but not encrypted. The agency says no information was used to commit fraud.

COSTS: Subcontractor breaches are more costly than internal incidents, averaging \$231 per record compared with \$171, according to Ponemon.

BLINDERS: According to Ponemon's annual cost study, breaches by outsourcers, contrac-

tors, consultants and business partners are on the rise, accounting for 44% of all cases reported by respondents last year. That's up from 40% in 2007. In the ITRC study, 10% of breaches were associated with subcontractors in 2008.

EYE-OPENERS: Companies need to create service-level agreements that are airtight and specific, and then ensure that subcontractors are in compliance and penalize them if they aren't. In cases that involve the use of backup tapes or disks, Semple says, insist on encryption and password protection.

Brandel is a *Computerworld* contributing writer. You can contact her at marybrandel@verizon.net.

Sponsored by



QuickStudy

QuickStudy: Storage Virtualization

By Russell Kay

MANAGING DISK STORAGE was once simple: If we needed more space, we got a bigger disk drive. But data storage needs grew, so we started adding multiple disk drives. Finding and managing these became harder and took more time, so we developed RAID, network-attached storage and storage-area networks. Still, managing and maintaining thousands of disk drives became an ever more onerous task.

The latest answer to this dilemma is storage virtualization, which

adds a new layer of software and/or hardware between storage systems and servers, so that applications no longer need to know on which specific drives, partitions or storage subsystems their data resides. Administrators can identify, provision and manage distributed storage as if it were a single, consolidated resource. Availability also increases with storage virtualization, since applications aren't restricted to specific storage resources and are thus insulated from most interruptions.

Also, storage virtualization generally helps automate the expansion of storage capacity, reducing the need for manual provisioning. Storage resources can be updated on the fly without affecting application performance, thus reducing downtime.

Technical Matters

Because virtualization operates as an intermediate layer, it becomes the primary interface between servers and storage. Servers see the virtualization layer as a single

storage device, while all the individual storage devices see the virtualization layer as their only server. This makes it easy to group storage systems — even devices from different vendors — into tiers of storage.

This layer shields servers and applications from changes to the storage environment, letting users easily hot-swap a disk or tape drive. Data-copying services are also managed at the virtualization layer. This means that data replication, whether for snapshot or disaster recovery, can be handled entirely by the virtualization system, often in the background, with a common management interface. Because data can be moved at will, lightly used or outdated data can be easily moved to slower, less-expensive storage devices.

Storage virtualization can be structured in three ways:

- **Host-based.** Here, physical drives are handled by a traditional device driver, while a software layer above the device driver intercepts I/O requests, looks up metadata and redirects I/O.
- **Storage-device-based.** In this type of setup, virtualization can be built into the storage fabric; for example, newer RAID controllers allow other storage devices to be attached downstream. A primary storage controller (usually a dedicated hardware appliance, though some systems now use switches) handles pooling and man-

Sponsored by



QuickStudy

ages metadata, allowing the direct attachment of other storage controllers. Such systems may also provide replication and migration services

across different controllers.

■ **Network-based.** In this configuration, storage virtualization is viewed as a network-based

device, generally using Fibre Channel networks connected as a SAN. Here, too, an appliance or switch-based implementation is most common.

Experienced users agree that all three approaches can work well. But although virtualization promotes cross-vendor storage utilization, most implementations lock

you into a specific vendor.

Kay is a *Computerworld* contributing writer in Worcester, Mass. Contact him at rususkay@charter.net.

Sponsored by



Analysis

Managing the Complexities of Storage Virtualization

Storage virtualization is hot, and for good reason. But its benefits bring added layers of complexity.

By Gary Anthes

THERE'S AN AGE-OLD CHOICE in IT — whether to adopt a “best of breed” strategy for the power and flexibility it can bring, or go with a single vendor for accountability and simplicity. J. Craig Venter Institute Inc. (JCVI) believes in best of breed. The genomic research institute runs Linux, Unix, Windows and Mac OS in its data center. For storage, it draws on technology from EMC, NetApp, Isilon, DataDomain and Symantec.

“It’s quite a heterogeneous envi-

ronment,” says computer systems manager Eddy Navarro. “Thankfully, we have a very talented staff here.”

And a talented staff was just what was needed to master the many flavors of storage virtualization, which can make multiple physical disks look like one big storage pool. Like JCVI, many organizations are enjoying the lower costs and added flexibility of storage virtualization. But the benefits can come with some headaches.

Here, five IT managers who have

led successful storage virtualization projects offer advice for relieving the pain.

Headache 1: Managing Multiple Vendors

For several years, JCVI had employed software-based virtualization in the form of Red Hat’s Linux Logical Volume Manager, which allows logical partitions to span multiple disk drives. More recently, the not-for-profit research institute added hardware-based virtualization in the form of NetApp’s V Series system to create a single virtual pool of storage consisting of EMC Symmetrix disks and legacy Clariion disks.

The Clariion drives, which came

into the data center from a corporate merger, were being poorly utilized, Navarro says. Now, the NetApp V system reformats data going to and from the EMC disks, “and then you carry on just as if it’s another NetApp system,” Navarro says. That enabled JCVI to wring better performance from the legacy disks.

Each of JCVI’s vendors makes its own unique contribution to a powerful and cost-effective storage architecture, Navarro says. But the diversity comes at a cost. “When you are talking about multiple vendors’ hardware — and they compete with each other — it may not be the easiest thing to get support when something goes wrong,” he says.

“So you have to ensure compatibil-

ity first and foremost, and you have to know in advance something is going to work.”

How to cope: Study the documentation, do your homework, and ensure that your approach has been tried before and is certified by the vendors, says Navarro. And if you don’t have experienced technical staff, he adds, be prepared to hire some outside professional help.

Headache 2: Dealing With Extra Technology Layers

Even companies with less-complex environments report that although virtualization can ultimately simplify storage administration, putting it in place

Sponsored by



Analysis

and tuning it is a demanding job.

Lifestyle Family Fitness, a rapidly growing chain of 60 health clubs based in St. Petersburg, Fla., is a Microsoft shop built around SQL

Server and .Net development of Web applications. For storage virtualization, it uses IBM's SAN Volume Controller (SVC), disk arrays from IBM and EMC, and IBM Bro-

cade SAN switches. IBM DS4700 disks provide 4Gbit/sec. Fibre Channel connections for the company's online transaction processing applications, while the Clariion

drives handle less-demanding jobs like backups.

The IBM SVC was brought in to resolve an I/O bottleneck. The high-speed Fibre Channel drives and cache on the SVC appliance opened up the bottlenecks almost like an I/O engine would, says Mike Geis, director of IS operations. Moreover, the setup allowed Lifestyle Family Fitness to use its new IBM-based SAN while continuing to use its old EMC SAN. "In the past," he says, "you'd bring in a new SAN and have to unload the old one."

Geis says the SVC architecture promises vendor independence. He says he has a "great relationship" with IBM, but if that ever changed, he could easily bring in drives from another supplier and quickly attach them directly to his storage network. "We aren't held hostage by the vendor," he adds.

But the advantages come with some difficulties, Geis notes. "You are adding complexity to your environment. You add overhead, man-hours of labor, points of failure and so on. You have to decide if it's worth it."

How to cope: "Pick strong partners – both vendors and implementation partners – and make sure you are not their guinea pig," Geis advises.

Headache 3: Scheduling Maintenance/Backups

Ron Rose, CIO at travel services company Priceline.com Inc., takes a holistic view of virtualization. In fact, he speaks of a "virtualization sandwich" consisting of integrated technologies for server virtualization, storage

How Am I Doing?

ONE OF THE IDEAS behind virtualization is to "abstract" the physical layer in IT from the software layer, to in essence mask hardware boundaries from the application and the application's users.

But the benefits of hiding the physical resources – greater flexibility, better utilization and potentially easier administration – come at a price, says Eddy Navarro, computer systems manager at J. Craig Venter Institute.

"So you have this abstracted area of storage, and you have a perfor-

mance issue," he says. "In the traditional model, it's a straightforward deduction to say that this area maps to these disks so that must be where the hot spot is. But with virtualization, if things are running slowly and there's this amorphous pool of storage, where exactly is the problem? You want to make sure you have the proper tools to tell you where the problems are."

Storage virtualization vendors have tools for performance monitoring and troubleshooting. "But with

these enterprise tools, it's a matter of installing agents everywhere, and it can balloon out of control," Navarro warns. "The agents themselves can cause this giant admin task. So is it really worth it to have this huge application monitoring things, or do you want a little bit of smarts and do some in-house work to write some custom scripts to tell you what's going on?"

JCVI has chosen to apply carefully targeted smarts via some home-grown software. "Fortunately," Navarro says, "we have the technical expertise to do that. If you don't, it's not easy to set that up."

Sponsored by



Analysis

virtualization and server provisioning. He uses 3PAR InServ S400 and E200 tiered disk arrays for storage, BladeLogic tools for provisioning, and 3PAR Thin Provisioning and other products for virtualization.

Rose says most companies could reduce their server and storage footprints by 20% to 40% using a virtualization sandwich. “And not only are there cost savings; there are green benefits. It’s good for the planet,” he says.

But like most practitioners of storage virtualization, Rose says there is no free lunch. “You have to plan your architecture more thoroughly and look at all your applications. The more systems you have running on the box, for example, the more challenging it is to schedule maintenance. If you have 10 applications running on that chunk of infrastruc-

ture that you are going to do maintenance on, you have to schedule it and move the apps to other machines in an orderly manner.”

He says 3PAR has powerful tools that can hide much of the complexity of virtualization, but the kind of maintenance scheduling needed “is not a system or tool issue; it’s a process and discipline issue.”

Similarly, ensuring reliability requires extra care, Rose says. “As with maintenance, you don’t want to get too many eggs in each basket,” he explains. Priceline keeps critical files on three machines – what it calls “tri-undancy.”

How to cope: “Think of your entire virtual environment, not just storage,” Rose advises. “You will get better ROI in aggregate if you think through all three layers of the virtual

sandwich. And getting a little consulting from real experts early in the process will help you anticipate the entire environment.”

Headache 4: Setting Up Management Tools

Like Rose, Jon Smith takes a very broad view of virtualization. “For me, a server is no different from a hunk of data storage, and I can move it wherever I want,” says the CEO of ITonCommand, a hosted IT services provider. “Whether it’s running the operating system or it’s just data, it’s all storage.”

Smith says that eventually virtualization technology will enable any data to go anywhere – on direct-attached storage when high performance is needed, or somewhere on a SAN when speed is less critical

and a higher level of redundancy is required.

ITonCommand uses HP BladeSystem c3000 disks for direct-attached storage, and LeftHand Networks Virtual SAN Appliances and LeftHand’s SAN/iQ software on an HP StorageWorks array for storage virtualization on its iSCSI SANs.

The company is now standardizing on Microsoft’s Hyper-V hypervisor, part of Windows Server 2008, for server virtualization and on Microsoft’s System Center Virtual Machine Manager for administration.

The glue that holds everything together, Smith says, is Microsoft’s new Virtual Machine Manager for provisioning and managing physical and virtual computers.

“With VMM on a display, a system admin can look at all the virtual servers’ hypervisors across my whole en-

vironment, all in one spot, and adjust them,” he says. “It’s pretty cool stuff.”

It’s cool when it’s set up, but getting there isn’t so easy, he acknowledges. “System Center is new, and so is [Hyper-V]. It took us a while to figure out how to connect all our old virtual machines into the hypervisor. It’s not the easiest setup out of the box.”

Smith says continued virtualization at ITonCommand will result in a true “utility computing” model for his clients. “It will take a while, but people will stop thinking of physical boxes running one operating system. Hardware will be nonexistent to the end user. It’s just going to be, ‘How much horsepower and storage do you want?’”

How to cope: “Find an expert who knows virtual technology and knows Microsoft System Center,” says Smith.

Sponsored by



Analysis

Headache 5: Getting the Right Gear

Babu Kudaravalli, senior director of business technology operations at National Medical Health Card Systems Inc., gives this definition of storage virtualization: “The ability to take storage and present it to any host, of any size, from any storage vendor.” He’s pursuing those goals with three tiers of storage, each supported by a different HP StorageWorks product. The technology used in each tier is chosen for the mix of cost, performance and availability it offers.

Kudaravalli uses high-end HP XP24000 disk arrays for the most demanding and mission-critical applications, lower-cost Enterprise Virtual Array 8000s for second-tier applications, and Modular Smart Array 1500s for archiving, test systems

and the like. His five SANs hold 70TB of data, of which about 35TB in the EVA and MSA tiers is virtualized, he says.

Kudaravalli says there are several things to be careful about when buying storage virtualization products. First, be aware that vendors typically certify their products to work with the latest versions of other vendors’ products. If you don’t have those exact versions, your interfaces might not work. He says this is a good reason to think about replacing your old gear when you go to a heterogeneous storage environment — or at least to keep current on the latest releases.

Second, Kudaravalli says that although virtualization should ultimately simplify storage management, setting up a virtual system is complex. Careful planning and an understanding of

the limitations of products is crucial.

A few years ago, vendors had very different definitions and standards for virtualization, says Kudaravalli. “But now they seem to be coming together,” he says. “They are trying to offer similar features and capabilities, but it is not completely mature.”

How to cope: Although storage virtualization is often undertaken to better utilize existing resources, it may have a perverse impact, says Rick Villars, a storage analyst at IDC. “The whole point of virtualization is to make it easier to provision or move a resource, to create a new volume or another snapshot, or to migrate data from one system to another,” he says. “But when you make something easy to do, people are induced to do it more often.”

According to Villars, volumes, snapshots, data sets and even applications can needlessly proliferate. “You can go from being more efficient to more wasteful. It’s just what can happen with virtual server sprawl.” Preventing that is a matter of policies, procedures and good business practices, not technology, he says.

Users agree that there are many technical details to master when pursuing storage virtualization. But Navarro suggests starting with a basic question: Why am I doing this? “Virtualization is a hot word, a big thing. But is it really necessary? There are benefits, but ask yourself if you are doing it for the right reasons, or just because you want to be on the cutting edge. It’s very easy to get swept up in these groupthink movements.”

Virtualization Pain Relievers

- 1. Be realistic: This is going to be complicated.**
- 2. Assign someone on staff who really knows the technology, or hire a consultant, at least at the beginning.**
- 3. Do your homework. Read the documentation and understand the pieces and their interfaces.**
- 4. Be sure that your gear and their interfaces are certified by your vendor(s) for the versions/releases that you have.**
- 5. Consider upgrading your old storage gear when you go to storage virtualization.**
- 6. Make sure you have thoughtful policies and procedures for maintenance and backups.**
- 7. Guard against “virtual sprawl” (of both storage and servers).**
- 8. Ask yourself: Do I really need this?**

Sponsored by



Analysis

Going Virtual Raises Storage-Management, Procurement Issues

Five things to think about before virtualizing storage **By Jon Brodtkin**

IF YOU'RE AN IT EXECUTIVE,

chances are you're already thinking about storage virtualization. Nearly one-quarter of companies with at least 500 employees have deployed storage virtualization products already, and another 55% plan to do so within two years, a recent Gartner survey found.

Storage virtualization is an abstraction that presents servers and applications a view of storage that is different from that of actual

physical storage, typically by aggregating multiple storage devices and allowing them to be managed in one administrative console.

The technology is emerging fast onto the enterprise scene for good reasons: In many cases, it can reduce the management burdens associated with storage; and offer better models for data-center migrations, backup and disaster recovery.

Enterasys Networks reaped these benefits recently when it moved a

data center from Boston into its headquarters in Andover, Mass.

"In days gone by, before storage virtualization, that might have been an all-day, if not an all-week kind of process," says Enterasys vice president of marketing Trent Waterhouse. "Because of the storage virtualization technologies, the entire move happened in less than 30 minutes."

There are still common pitfalls that storage administrators should ponder, as well as questions they should ask before they roll out a storage-virtualization project. Here's a look at some of the top issues.

Managing capacity

With storage virtualization, allocating storage is easy — perhaps too easy.

"You have the ability to affect more systems in the whole for-

est if you do something," says Jonathan Smith, CEO of ITonCommand in Denver, Colo., who cautions fellow IT shops to pay close attention to both the storage and performance needs of each application. "You just didn't have that power before. Now all of a sudden you can do whatever you want."

Smith, who is using LeftHand Networks virtualization on HP storage, says an IT pro might see a lot of empty space in a given storage volume and be tempted to fill it up. Overusing a resource, however, can decrease performance if the storage is allocated to a database or some other I/O-intensive application.

"Make sure you size it correctly and really understand how much horsepower [your applications need]," Smith says.

These concerns are especially

true when it comes to thin provisioning, a component of virtualization technology that lets an IT administrator present an application with more storage capacity than is physically allocated to it. This eliminates the problem of storage over-provisioning, in which storage capacity is pre-allocated to applications but never used.

With thin provisioning, more than 100% of storage capacity can be allocated to applications, but capacity remains available because it won't be consumed all at once.

You can play it safe by allocating small volumes that never exceed the physical storage, or allocate as much as you want to each application, then monitor your systems closely, says Themis Tokkaris, systems engineer at Truly Nolen Pest Control in Tucson, Ariz. It's best if

Sponsored by



Analysis

you can find a happy balance between those two extremes.

“You have to monitor your pool so you don’t run out of space, because that would really crash everything,” Tokkaris says.

How server virtualization fits in

A common question is whether it makes sense to virtualize storage if you’re not also using server virtualization. The short answer is yes — though it’s true you won’t get as much flexibility as IT shops that virtualize both servers and storage.

“If you virtualize both, then you have the maximum flexibility when deploying new applications,” says Chris Saul, IBM’s storage-virtualization marketing manager.

Nevertheless, there are benefits to just virtualizing storage.

Improved disaster recovery, availability and data migrations can all be gained without having virtual servers, says product marketing manager Augie Gonzalez of storage virtualization vendor DataCore Software. In addition, storage virtualization by itself can provide thin provisioning, as well as the simplified management structure that comes with pooling storage devices and managing them from a central console.

On the flip side, virtualizing servers without virtualizing storage is problematic. It doesn’t make sense to have multiple virtual servers on a physical machine that aren’t able to share data, says Enterprise Strategy Group (ESG) analyst Mark Peters.

“You can gain tremendous benefits from storage virtualization,

even without server virtualization. It’s harder the other way around,” Peters says.

Single-vendor virtualization

Given that virtualization is designed to combine multiple storage devices, it’s not immediately obvious why it makes sense to virtualize your storage if it all comes from a single vendor.

There are compelling reasons, however, says storage analyst Arun Taneja. “A lot of people think storage virtualization has a prerequisite of heterogeneity, that it only comes into play when storage from three companies is involved,” he says. “I say, forget it, it has value even if you are stuck with a single vendor.”

The storage market is more pro-

prietary than just about any other IT space, and this creates problems even if you have just one storage vendor, Taneja says.

Say you’re an EMC customer with two Symmetrix DMX boxes, and “you just want to combine the power of those two boxes and manage it as one,” Taneja says. “[Without storage virtualization] you can’t do it. That’s how ridiculous the world of storage is.”

This “ridiculous” level of exclusivity in the storage market obviously takes on a new dimension when you’re managing storage from multiple vendors. That leads to the next issue.

Choosing a vendor

Enterprises’ primary procurement dilemma is whether to purchase storage-virtualization products from a storage vendor or a third party.

If your true objective is flexibility, especially if you’re planning major data migrations, a third party is the way to go, Taneja says. Such vendors as FalconStor Software and DataCore are capable of managing storage from multiple vendors simultaneously, whether they are EMC, HP, IBM or Hitachi.

Truly Nolen chose a third party, DataCore, even though the company uses only HP storage. The company evaluated virtualization vendors including HP, EMC, and Dell EqualLogic, but settled on DataCore because it was less expensive and offers the flexibility of using whichever hardware vendor it likes, Tokkaris says.

The major storage vendors promise to be able to manage a heterogeneous environment. Examples include IBM’s SAN Volume

Sponsored by



Analysis

Controller, NetApp's V-Series, and EMC's Invista. As a general rule, though, vendors support their own storage products first and others second, if at all.

"They always support their own systems first," Taneja says. "That means EMC's Invista supports DMXs and Clariions, and they might support some other foreign devices; but the support for foreign devices always lags, and support for foreign devices is always incomplete. The whole idea is don't support your enemies' boxes."

Peters predicts that as storage virtualization becomes more common, market pressure will force vendors to do a better job supporting their rivals' technology.

If you get storage from just one vendor, however, the solution is simple.

"I say to the IT people I talk to, if you're a Hitachi Data Systems customer and you like working with them and you're stuck with them, just buy their virtualization to make life more manageable within Hitachi product," Taneja says.

Sifting through the hype

By most accounts, storage virtualization is a no-brainer. Who wouldn't want to manage multiple storage devices from a single console, and gain data mobility that makes disaster recovery a breeze?

Storage virtualization will be about as common as automatic transmissions in automobiles within a couple of years, ESG's Peters thinks. "There are certain technologies that are just smarter and better than people doing

it manually," he says.

Even storage virtualization vendors, however, can admit there are instances when the technology isn't a fit.

Storage virtualization is not for everyone, says Kyle Fitze, an HP director of storage marketing. Virtualization actually adds a layer of complexity, he argues. You have to manage the individual storage devices, as well as the virtualization layer, he notes. Despite virtualization, you still have to perform such tasks as reconfiguring devices after adding physical disks to

storage arrays, he adds.

As a general rule of thumb, the more complicated your storage environment, the more benefit virtualization brings.

"There's a complexity/benefit tradeoff," Fitze says. "If their current environment is difficult to manage and complex . . . adding a virtualization layer can simplify that complexity. If it's a small, efficiently managed environment without data-protection challenges, then virtualization just for virtualization's sake is probably not a good idea."

Sponsored by



Analysis

Thinking outside the storage box

Unbridled growth in data storage and the rise in Web 2.0 applications are forcing a storage rethink. Is this the end of the storage-area network as we know it? **By Joanne Cummings**

WITH STORAGE GROWTH tracking at 60% annually, according to IDC, enterprises face a dire situation. Throwing disks at the problem simply doesn't cut it anymore. Andrew Madejczyk, vice president of global technology operations at pre-employment screening company Sterling Infosystems, in New York, likens the situation to an episode of "House," the popular medical drama.

"On 'House,' there are two ways to approach a problem. You treat the symptoms, or you find out what the root cause is and actually end the problem," Madejczyk says. "With

storage, up until now, the path of least resistance was to treat the symptoms and buy more disks" — a method that surely would ignite the ire of the show's caustic but brilliant Dr. Gregory House.

Were the doctor prone to giving praise, he'd give a call out to enterprise IT managers who are rethinking this traditional approach to storage. He'd love that technologists are willing to go outside their comfort zones to find a solution, and he'd thrive on the experimentation and contentiousness that surround the diagnosis.

House probably would find an

ally in Casey Powell, CEO at storage vendor Xiotech. "Everybody acknowledges the problem and understands it, but nobody's solving it. As technologists, we have to step back, look at the problem and design a different way," Powell says.

Optimizing the SAN

Today most organizations store some combination of structured, database-type and unstructured file-based data. In most cases, they rely on storage-area network (SAN) technologies to ensure efficiencies and overall storage uti-

lization, keeping costs down as storage needs increase.

In and of themselves, SANs aren't enough, however. Enterprises increasingly are turning to technologies that promise to provide an even bigger bang for the buck, including these:

- **Data deduplication**, which helps reduce redundant copies of data so firms can shrink not only storage requirements but also backup times.
- **Thin provisioning**, which increases storage utilization by

making storage space that has been overprovisioned for one application available to others on an as-needed basis.

- **Storage tiering**, which uses data policies and rules to move noncritical data to slower, less expensive storage media and leave expensive Tier 1 storage free to handle only the most mission-critical applications.
- **Storage resource management software**, which helps users track and manage storage usage and capacity trends.

"In the classic SAN environment, these tools don't just provide a partial solution. They al-

Sponsored by



Analysis

low you to make fundamental improvements,” says Rob Soderbery, senior vice president of Symantec’s Storage and Availability Management Group. Clients that have pursued these strategies even have been able to freeze storage spending for a year at a time, he says. “And when they get back on the storage spending cycle, they get back on at about half the spending rate they were at before,” he adds.

Although few IT executives report such dramatic reductions in storage spending, many are pursuing such strategies.

At Sterling, for example, moving from tape- to disk-based backups via Sepaton’s S2100-ES2 virtual tape library reduced the time it takes for nightly backups from 12 to just a few hours, Madejczyk

says. Sepaton’s deduplication technology provides an added measure. In addition, he has virtualized more than 90% of his server environment, “reducing our

footprint immensely” and implemented EMC thin provisioning and storage virtualization.

Still, his company’s storage needs keep growing, Madejczyk

says. “In this economy, Sterling is being very responsible and careful about what we spend on,” he says. “We’re concentrating on the data-management part of the problem, and we’re seeing results. But it’s a difficult problem to solve.”

Tom Amrhein, CIO at Forrester Construction in Rockville, Md., has seen similar growth. The company keeps all data associated with its construction projects in a project management database, so the vast majority of that stored data is structured in nature. Regulatory and compliance issues have led to increased storage needs nonetheless.

“Most companies need to keep their tax records for seven years, and that’s as long as they need to keep anything,” Amrhein says. “But tax records are our shorter-


cycle data. Depending on the jurisdiction, the time we could be at fault for any construction defect is up to 10 years – and we’re required to have the same level of discovery response for a project completed nine years ago as we would for a project that’s been closed out two weeks.”

Forrester Construction has cut down a bit on storage needs by keeping the most data-intensive project pieces – building drawings, for example – on paper. “Because the scanning rate is so high and paper storage costs so low, retaining those as physical paper is more cost-effective,” Amrhein says.

The real key to keeping costs in check, however, is storage-as-a-service, Amrhein says. IT outsourcer Connectria hosts the company’s main application servers, including

Got files?

THE EXPLOSION OF FILE-BASED DATA WILL SEE UNSTRUCTURED STORAGE NEEDS SURPASS STRUCTURED STORAGE DEMANDS BY NEXT YEAR, IDC REPORTS.



Sponsored by



Analysis

Microsoft Exchange, SQL Server and SharePoint; project management, finance, CRM and Citrix Systems. It handles all that storage, leaving Forrester Construction with a predictable, flat monthly fee.

"I pay for a set amount of gigabytes of storage as part of the SLA [service-level agreement], and then I pay \$1 per gig monthly for any excess," Amrhein explains. "That includes the backup, restore and all the services around those. I'm paying \$25K a month to Connectria, plus paying for about 10GB over my SLA volume. That overage is a wash."

For the firm's unstructured data, Forrester Construction uses Iron Mountain's Connected Backup for PCs service, which automatically backs up all PCs nightly via the Internet. If a PC is not connected to the Internet at

night, the user receives a backup prompt on the next connection.

"With 60% of the people out of the office, this is perfect for us," Amrhein says. "Plus, Iron Mountain helps us reduce the data volume by using deduplication," he says. "Even for people on a job site with a wireless card or low-speed connection, it's just a five- or 10-minute thing."

Still, the unstructured side is where the construction company sees its biggest storage growth. E-mail and saved documents are the biggest problem areas.

The rise in Web 2.0 data

Forrester Construction is not alone there. In the enterprise, IDC reports, structured, transactional data will grow at a 27.3% compounded annual rate over the next three to five years. The rise

in unstructured, file-based data will dwarf that growth rate, however. IDC expects the amount of storage required for unstructured, file-based data to increase at an unprecedented 69.4% clip. By 2010, enterprises for the first time will find unstructured storage needs outstripping traditional, structured storage demands.

The rub here is that although SANs are extremely efficient at handling structured, transactional data, they are not well optimized for unstructured data. "SANs are particularly ill-suited to Web 2.0, scale-out, consumer-oriented-type applications," Symantec's Soderbery says. "No. 1, the applications' architecture is scale-out, so you have hundreds or thousands of systems working on the same problem instead of one big system, like you would have with a da-

tabase. And SANs aren't designed that way. And No. 2, these new applications — like storing photos on Facebook or video or display ads or consumer backup data — are tremendously data intensive."

Symantec hit the wall with this type of data in supporting its backup-as-a-service offering, which manages 26 petabytes of data, Soderbery says. "That probably puts us in the top 10 or 20 storage consumers in the world. We could never afford to implement a classic Tier 1 SAN architecture," he says.

Instead, Symantec went the commodity path, using its own Veritas Storage Foundation Scalable File Server software to tie it all together. "The Scalable File Server allows you to add file server after file server, and you get a single namespace out of that

cluster of file servers. This in turn allows you to scale up your application and the amount of data arbitrarily. And the software runs on pure commodity infrastructure," Soderbery explains. Plus, the storage communicates over a typical IP network vs. a more expensive Fibre Channel infrastructure.

Symantec's approach is similar to that of the big cloud players, such as Google and Amazon.com. "We happen to build packaged software to enable this, whereas some of the early adopters built their own software and systems. But it all works the same way," Soderbery says.

The prudent approach to storage as it continues to grow, Soderbery says, is to optimize and use SANs only for those applications that merit them — such as high-transaction, mission-critical

Sponsored by



Analysis

ERP applications. Look to emerging commodity-storage approaches for more scale-out applications, such as Web 2.0, e-mail and interactive call-center programs.

Does that mean enterprises need to support SANs and new cloud-like scale-out architectures to make sure they're managing storage as efficiently as possible? Perhaps.

Eventually, however, the need to support unstructured, scale-out data will trump the need to support structured, SAN-oriented data, IDC research shows. With that in mind, smart organizations gradually will migrate most applications off SANs and onto new, less expensive, commodity storage setups.

A new enterprise approach

One interesting strategy could

provide an evolutionary stepping-stone in the interim: using Web services. Championed primarily by Xiotech, the idea is to use the Web-services APIs and standards available from such organizations as the World Wide Web Consortium (W3C) as the communications link between applications and storage.

"The W3C has a nifty, simple model for how you talk between applications and devices. It includes whole sets of standards that relate to how you provision resources in your infrastructure, back to the application," says Jon Toigo, CEO of analyst firm Toigo Partners International. "All the major application providers are Web-services enabled in that they ask the infrastructure for services. But nobody on the storage hardware side is

talking back to them."

Nobody, that is, except Xiotech.

Xiotech's new Intelligent Storage Element (ISE) is the first storage ware to talk back, although other vendors quickly are readying similar technology, Toigo says. ISE, based on technology Xiotech acquired with Seagate Technology, is a commodity building-block of storage, supporting as many as 40 disk drives plus processing power and cache that can be added to any storage infrastructure as needed. Xiotech claims ISE can support anything from high-performance transactional processing needs to scale-out Web 2.0 applications.

All storage vendors should work to Web-services-enable their hardware and software so they can communicate directly with applica-

tions, Xiotech's Powell says. This would preclude vendor lock-in and let enterprises build storage environments using best-in-breed tools instead of sticking with the all-in-one-array approach. "They'd be able to add more storage, services or management, without having to add everything monolithically to a SAN," Powell says.

Eventually Web services support will have virtualized storage environments realizing even greater efficiencies, to the point where applications themselves provision and de-provision storage. "Today, when we provision storage, we have to guess, and typically, we either over- or underprovision," Powell says. "And then, when the user is no longer using it, we seldom go back and reclaim the storage. But the application knows exactly

what it needs, and when it needs it. Via Web services, it can request what it needs on the fly, and as long as its request is within the parameters and policies we set up initially, it gets it."

Web services already have proved an efficient storage tack at ISE user Raytown Quality Schools in Missouri, says Justin Watermann, technology coordinator for the school system. The system went with Xiotech shortly after it moved to a new data center and created an all-virtual server infrastructure. A big plus has been Xiotech's Virtual View software, which uses Web services to communicate with VMware's Virtual-Center management console for its ESX servers, Watermann says. He can manage his virtualized server and storage infrastructure

Sponsored by



Analysis

from a single console.

“When you create a new data store, Virtual View shows you what port and LUN [logical unit number] is available to all of your ESX hosts in that cluster,” Watermann says. “And when you provision it, it uses Web services to communicate with VirtualCenter, and says, ‘OK, this is the data store for these ESX hosts.’ And you automatically have the data store there and available to use. You don’t even have to refresh or restore.”

That bit of automation saves on administration, but enabling the application to do the provisioning and de-provisioning would be an even greater boon, Watermann says. “It’s really hard to get more staff, and you only have so many hours in the day. If you don’t have to tie your staff up with the repeti-

tive tasks of carving up space and assigning it, so much the better.”

If they build it . . .

Now Watermann is using a Web service from Xiotech partner Eagle Software to improve the school system’s backup times.

Eagle, a storage reseller, provides Raytown with backup software from CommVault Systems. “The Web-services tool lets us mirror our data, pause that mirror, attach that to the backup server, back it up, then disconnect it from the backup server, unpause the mirror, and re-sync that data so we don’t have to push it across our Ethernet network,” Watermann says.

Like Eagle, companies are starting to develop pieces to put Web-services-enabled storage all

together, Watermann adds. Such small, point approaches are the norm today, but experts say that in five or 10 years, every application and every device will use some kind of software glue, such as Web services, to provide a range of storage and IT services in an automated, efficient manner.

“It would have to be levels of software that create services that include computing resources, network and storage infrastructure, and the retention and reliability metrics associated with all of those components,” Symantec’s Soderbery says. “It will be a combination of the point solutions we see now, like VMware and SAN virtualization via Cisco and Brocade, plus thin provisioning, replication and deduplication. We’re going to require all of those things to work

in concert with a level of software that ties them together cohesively to provide those appropriate levels of service to the application.”

Others, including Ken Steinhardt, CTO for customer operations at EMC, are less optimistic. “If someone could write something magical that does things that we’d love to have, wouldn’t that be great?” he asks rhetorically.

That would take a miracle, Steinhardt says. “The tools to write software are out there, but the catch is it’s just not that simple. You need to be able to have a solution that works broadly across a range of applications as well, and typical, highly consolidated environments run a mix of broad, diverse apps, not just a single application. I don’t see it happening,” he says.

The Web services model is too much of a stretch to Steinhardt: “From a storage perspective, Web services are a completely separate issue. We’re talking about storing zeros and ones on a storage device. That’s pretty agnostic to the individual application and always has been,” he says.

Not so, analyst Toigo asserts. Web services provide a common framework, so by default, they can support every application. “People need to tell their vendors, ‘Look, I’m not buying your junk anymore if I can’t manage it using this common [Web services] metaphor,’” he says. “That puts you in the driver’s seat.”

Cummings is a freelance writer in North Andover, Mass. Reach her at joanne@cummingslimited.com.

Sponsored by

