Managing the information that drives the enterprise

STORAGE

ESSENTIAL GUIDE TO

Efficient Data Storage

Tips and tools on managing your data storage more efficiently are highlighted in this guide.

INSIDE

TechTarget

SYMANTEC IS

You can't stop the explosion of data across your enterprise, but you can change the way you back it up. NetBackup™ with integrated deduplication assures that data from your data centers, remote offices, and virtualized environments is stored only once — no matter what your storage hardware.

Upgrade to Netbackup 7 today. For more information and trial, visit www.netbackup.com

DEDUPLICATION.

Confidence in a connected world.

symantec™

Efficiency = Cost-Effective Storage

New tools now make it possible to use installed storage more efficiently than ever.

N INTERESTING THING happened as savvy storage managers started to see the light at the end of the recessionary tunnel. They not only realized that they had weathered one of the worst economic storms on record, but they did it with measly budgets while still managing to accommodate all that new data the business created.

But what's more impressive is how they did it. Not too long ago, the conventional wisdom was you met capacity challenges by rolling more big iron out onto the data center floor. Hardly an elegant solution; just parking a new array on the floor was kind of a blunt-instrument approach to the problem but there were few practical tools available to deal with the causes of skyrocketing capacity.

All that has changed—and fairly quickly, too. Storage managers learned a few valuable lessons as they struggled to scrape by, the most important of which was it's better to figure out how to use what's already installed more efficiently than to add even more gear to manage. To do that there had to be a shift in storage management from managing systems to managing data. Some of the tools to handle the new chores were already in place, like thin provisioning and storage tiering techniques. Practically every storage vendor offers thin provisioning as a system option, and more and more are joining the ranks of those now offering automated tiering tools to help ensure that data sits on the appropriately priced disk.

Using disk more efficiently and making sure data is in the right place are excellent steps, but still more control over data, on a more granular level is needed to be able to wring out the greatest efficiencies. Once just a niche application supported by only a few vendors, data reduction technologies have gained considerable visibility because they can help address the need for more effective management of stored data. These products address storage efficiency at the most basic level, and help conserve valuable disk space by compressing files to remove wasted space or by removing redundancies among all files in a group. The results are often dramatic, with required capacity often cut by half or more.

Thin provisioning, automated tiering and data reduction are perhaps the most talked about storage efficiency technologies, but there are others, too. And smart storage managers are building their own arsenals of efficiency tools to not only stretch their storage dollars but also to take some of the complexity out of managing storage. ⊙

Rich Castagna (rcastagna@techtarget.com) is the editorial director of TechTarget's Storage Media Group.

MANAGING ENTERPRISE DATA STORAGE MORE EFFICIENTLY: RECLAIM STORAGE AND CONSOLIDATE DATA

Enterprise data storage managers are using tiered storage, data deduplication, snapshots, thin provisioning and SSDs to reclaim storage and consolidate the amount of data in their firms.

By Beth Pariseau

TO RECLAIM STORAGE and consolidate the amount of data in their firms, enterprise data storage administrators are increasingly turning to technology tricks they can implement to get the most out of their storage arrays. Using techniques such as tiered storage, data deduplication, space-efficient snapshots, thin provisioning and wide striping, solid-state drives (SSDs) and energy efficiency, they can store as much data as possible in the smallest and least-expensive amount of space.

TIERED STORAGE

Tiered storage has been in vogue for a while, but tiering projects are becoming more popular as companies seek to avoid purchasing more disk capacity.

University HealthSystem Consortium (UHC), an Oak Brook, Ill.-based industry group for U.S. academic hospitals and medical centers that collects data and coordinates cooperation between healthcare facilities, has been aggressive with its tiered storage strategy. Instead of starting data at tier 1 and progressing down through lower tiers as it ages, all

data starts at tier 2, and the organization uses Hitachi Data Systems' Tuning Manager to evaluate what data needs to be moved up to tier 1 for performance reasons.

UHC purchased the Hitachi Data Systems' Universal Storage Platform V (USP V) storage virtualization controller in 2008 when it forecasted 400% data growth. Steven Carlberg, UHC's principal network administrator, said he chose the USP V system over systems from EMC Corp., Hewlett-Packard (HP) Co., IBM Corp. and Sun Microsystems Inc. because he anticipated virtualizing arrays behind the controller would make future migrations easier.

"Obviously, the main cabinet would take a forklift upgrade to replace, but I don't foresee having to swap it out for a long, long time. It can scale up to 247 petabytes," he said. "Going forward, on lower tiers, we can change out obsolete hardware without a lot of pain."

DATA REDUCTION MAXIMIZES TIERED STORAGE EFFICIENCY

Clackamas County in north central Oregon deployed tiered storage by using F5 Networks Inc.'s ARX file virtualization switch and Data Domain's DD565 data deduplication disk arrays in an effort to put off adding capacity to its tier 1 SAS-based iSCSI storage-area network (SAN) storage.

Similarly, Vancouver, B.C.-based Rainmaker Entertainment Inc. uses an ECO appliance from Ocarina Networks with network-attached storage (NAS) systems from BlueArc Corp. and NAS systems from Isilon Systems Inc. to reduce primary storage and keep more archival data online.

"Our tiered storage policy will free up about 10 TB of tier 2 storage. That, in turn, frees up SAS capacity used by database servers and our email archiving system," said Christopher Fricke, Clackamas County's senior IT administrator. "It helps us not have to chase capacity while we go through a budget crunch. We can focus on performance rather than capacity. In a couple of budget cycles, we should be able to look at replacing our current tier 1 with solid-state disk ."

"It helps us not have to chase capacity while we go through a budget crunch. We can focus on performance rather than capacity. In a couple of budget cycles, we should be able to look at replacing our current tier 1 with solid-state disk ."

—CHRISTOPHER FRICKE, senior IT administrator, Clackamas County

Data reduction technologies originally deployed for backup are making their way into primary and nearline storage. Cornell University's Center for Advanced Computing (CAC) and e-commerce site Shopzilla are among the early adopters of primary storage data reduction to consolidate their

storage and keep up with data growth. Both use compression appliances from startups. Ithaca, N.Y-based Cornell University runs Ocarina Networks ECOsystem appliances, and Shopzilla has Storwize Inc.'s STN-6000 device.

CAC is also trying to create economies of scale for the entire campus by getting other departments in the university to store research data on its storage. "We're hoping researchers will put their data on centralized storage devices as opposed to spending excessive amounts of money on terabyte USB drives that are deployed in silos," said David Lifka, director of CAC. "Siloed technologies cost the university money. They reduce scalability and cost more to maintain."

With Ocarina Networks, the effective cost of a terabyte is approximately $500, a price point that's appealing to CAC's clients. At the rate researchers are signing on, Lifka said, CAC may have to add another 100 TB to its DataDirect Networks Inc. arrays this summer.

SPACE-EFFICIENT SNAPSHOTS

Justin Bell, network administrator at Madison, Wis.-based engineering firm Strand Associates Inc., said space-efficient snapshots save valuable space on his storage arrays.

"We don't need deduplication," Bell said. Strand only sends changes to data on the company's servers at the main and branch offices to the backup server with FalconStor Software Inc.'s DiskSafe product. That reduces 12 TB total capacity, and 2 TB logical snapshot capacity (about four months worth of snapshots, according to Bell) takes up 660 GB with only the changes saved.

> "We have the space to keep a year-and-a-half's worth of backup snapshots on all production drives without deduplication."
>
> —JUSTIN BELL, network administrator, Strand Associates Inc.

"We have the space to keep a year-and-a-half's worth of backup snapshots on all production drives without deduplication," he said.

The company uses Riverbed Technology Inc.'s Steelhead wide-area network (WAN) optimization devices to centralize storage of production files among the branch offices, but Bell said Riverbed's data deduplication algorithms "don't touch the FalconStor stuff." Riverbed has its own compression and protocol optimization for sending data over a WAN, but "that wouldn't do much, since FalconStor is rarely sending the same data twice," Bell said.

THIN PROVISIONING AND WIDE STRIPING

UHC also uses Hitachi Dynamic Provisioning (HDP) to wide stripe and thin provision volumes. "We can run at 125% over-allocation fairly safely," UHC's Carlberg said. With the most recent release, volumes can be dynamically expanded.

"It's saved us countless hours of management," Carlberg said. "Instead of struggling to manage with almost two full-time engineers [FTEs], we're now at one-half of an FTE with 400% data growth."

He'd like to see Hitachi Data Systems go further, however. "I wish we had zero page reclaim when we were migrating volumes off [our older array]," Carlberg said. "We had to write the full volumes to the USP V whether they were really full or not."

Carlberg added that he hopes Hitachi Data Systems will add the ability to dynamically shrink over-allocated volumes if necessary, and he may get his wish.

"We're working with vendors like Symantec [Corp.] and industry standards groups to address automatic space reclamation and improve the ability for file systems to tell us when we can unmap storage and return it for better uses," wrote John Harker, Hitachi Data Systems' spokesman, in an email to SearchStorage.com. "We expect to share some news before the end of the year."

"We're working with vendors like Symantec [Corp.] and industry standards groups to address automatic space reclamation and improve the ability for file systems to tell us when we can unmap storage and return it for better uses."

—JOHN HARKER, spokesman, Hitachi Data Systems

SSDS AND ENERGY EFFICIENCY

The cost per gigabyte of solid-state drives remains a barrier to adoption for most users in the enterprise data storage market. But Bill Konkol, vice president of technology at broadcast advertising design firm Marketing Architects Inc. in Minneapolis, said he has met performance requirements more cheaply with Compellent Technologies Inc. Storage Center SAN than using short-stroked traditional hard disk drives.

To get the required performance, Konkol estimated he'd have to short-stroke 28 Fibre Channel disks, which would cost $75,000 more than the Compellent SAN he bought with less than 500 GB of STEC Inc. SSDs. Adding the solid-state drives has boosted performance enough so that searches that previously took hours now take approximately 15 minutes, he said.

Scott Stellmon, systems architect at Palo Alto, Calif.-based The Children's Health Council (CHC), said he's looking to deploy an Infortrend Inc. EonStor B12F modular storage system, which is scaled in 1U, 12-drive units. Infortrend's system can also accommodate SSDs in 2.5-inch drive slots, which Stellmon said he'd like to add for primary databases.

"We haven't made the final decision yet, but because some of our

data sets are small enough, it won't take much scale in SSD capacity," he said.

An 80 GB PCI Express solid-state drive costs approximately $3,000, and most of CHC's databases are between 5 GB and 8 GB.

This will be cheaper and draw less power than adding enough disk drives to the system to keep databases performing fast. "If I can justify saving money on electrical costs [with SSDs], I can turn around and use that money to buy more storage gear," Stellmon said.

Beth Pariseau is the senior news writer for SearchServerVirtualization.

MANAGING ENTERPRISE DATA STORAGE MORE EFFICIENTLY:

Reduce operating expenses using technology

Enterprise data storage managers can reduce expenses through capacity planning, writing scripts to automate storage, outsourcing and taking a long-term view on efficiency.

By Beth Pariseau

WHILE CAPITAL PURCHASES are clearly the most sensitive area for organizations with slashed IT budgets, operating expenses are also under pressure. Data centers face serious staffing shortages, and administrators say they're being asked to bridge the gap between capital budget cuts and the need to improve services by working harder and smarter rather than with new software or hardware. To accomplish this, enterprise data storage managers are using capacity planning, writing scripts to automate storage, outsourcing and taking a long-term view on efficiency.

Chris Cummings, senior director, data protection solutions at NetApp, said he's seen this pattern of behavior. "At the beginning of [2008], the thing on everyone's mind was power and cooling with oil nearing $200 a barrel," he said. "After September, those conversations tapered off—operational expenses were still important, but acquisition costs became paramount."

USE CAPACITY PLANNING TO MANAGE STORAGE

With the do-more-with-less attitude prevalent, it's no wonder capacity planning jumped to the top of TheInfoPro Inc.'s Wave 12 Storage Study Heat Index last April. The New York City-based market research firm's Heat Index measures the technologies people intend to implement. Capacity planning was 10th on the previous list and languished for years in the "nice to have" category among storage buyers.

Ed Sitz, IT manager at Med James Inc., a Shawnee Mission, Kan.-based management company, said he's quickly running out of space on his 15 TB EMC Corp. Clariion CX3-20 disk array. He can add expansion shelves, but at significant cost.

That's delaying Sitz's plans of adding Microsoft SQL and Exchange servers to the 20 or so servers he already virtualizes with VMware. Most of the hosts are attached to the SAN through iSCSI, but SQL performance requirements demand more expensive Fibre Channel (FC) attachment.

"Funding for Fibre Channel is a whole other issue," Sitz said. "We have storage capacity on there for adding a LUN [logical unit number] if needed, but putting virtualized SQL and Exchange on the same RAID group as other servers isn't something we want to approach right now."

In the meantime, Sitz said he plans to work with his EMC solution provider to analyze how the storage space is currently being used and how it can be better utilized. "It seems we're always trying to learn while implementing," Sitz said. "We really haven't stopped and analyzed the apps on particular RAID groups and how the LUNs are laid out. Are they managed efficiently? If not, is there a way to move them?"

AUTOMATE STORAGE WITH SCRIPTS

Enterprise data storage managers often search for ways to automate tasks. Currently, storage automation often means writing scripts to avoid having to buy tools that could make life easier.

Guy Gaylord is the IT manager in the Atlanta branch office of TTI Telecom, which is headquartered in Israel. Gaylord said he's conducting internal meetings with end users to see if service levels meet their expectations.

"The user of an application knows what needs to be done," he said. "We're paying attention to the end-users' requirements to see what we can do better."

"The user of an application knows what needs to be done. We're paying attention to the end-users' requirements to see what we can do better."

—GUY GAYLORD, IT manager, TTI Telecom

A lot of those improvements "could be made better by automating more, but I don't mean acquiring new tools," Gaylord said. "I mean just sitting down and writing scripts."

Scripting has also come in handy for Joe Tailleur, IT manager at James Western Star Sterling Ltd., a truck dealer in British Columbia, Canada. Tailleur oversees a centralized IT environment for three sites with approximately 1 TB of data spread over three physical servers running VMware.

Tailleur said even most small- to medium-sized business (SMB) products from mainstream storage vendors are outside his price range. “Their version of SMB isn’t our version,” he said.

James Western Star Sterling has instead leaned on Tailleur’s expertise in scripting and designing software applications to do work in-house that it would otherwise have to pay for. The company uses the free version of VMware and Tailleur scripts backup operations using open-source code. “We run batch files, shut down the virtual machine, copy it to a backup Windows server and from there dump it to tape,” he said.

"We run batch files, shut down the virtual machine, copy it to a backup Windows server and from there dump it to tape."

—JOE TAILLEUR, IT manager, James Western Star Sterling Ltd.

Adam Vogini, information services administrator at Canonsburg, Pa.-based Trigon Holding Inc., a holding company with subsidiaries in aerospace and medical manufacturing, also found himself taking on labor-intensive open-source storage projects to make ends meet while expanding storage.

Trigon cut costs by deploying network-attached storage (NAS) servers based on an open-source distribution called Openfiler rather than buying NAS from its primary storage vendor, EMC. “It saved us between $15,000 and $20,000” over the purchase of a Celerra NX4 low-end NAS system, Vogini said.

The project was undertaken with a bare-bones staff. Vogini said a plan to hire a full-time engineer was postponed last year when the economy went sour. “That got pushed down to a summer intern, and that was based on a grant available to companies in IT and engineering that hire interns in the Pittsburgh area,” he added.

OUTSOURCING TO THE CLOUD

Jon Krouskoff, director of instructional technology at Clarkstown Central School District in Rockland County, NY, uses a cloud data storage service, School Web Lockers, to facilitate communication between teachers and students while keeping costs down.

Allowing students to use webmail and portable devices to transport data across the school’s firewall would present a security risk, and the school lacked the infrastructure to offer online access to data inside and outside the school building.

The School Web Lockers service lets teachers store files and assignments online for students, and distribute files to the right students. “You can say “I want what I just uploaded to also go to these three class periods,” Krouskoff said.

Although this cloud service does save money, it has some downsides as well, mainly related to bandwidth issues. “Because it’s browser based, the file size is technically unlimited, but most browsers time out after 30 to 40 seconds,” he said. “You can get away with it when you know the network is relatively quiet, but if it’s busy, it might time out before the file is stored or retrieved.”

Krouskoff said he doesn’t feel he’s had much choice but to focus on cost-effectiveness. “Even a rollover of the budget keeping the status quo would have been an 8% to 9% increase from year to year,” he said. “We’ve had to make significant cuts to programs like notebook and netbook acquisition projects. We also haven’t issued a technology bond this year because of the economy. At a dollar a student, we look at School Web Lockers as an essential component of what we do.”

IMPROVE STORAGE EFFICIENCY FOR LONG-TERM CHANGE

Trigon Holding’s Vogini said the current downturn shows that IT has become more critical than the last time the economy cratered. “IT definitely hasn’t been hurt like it was in 2002,” he said. “A lot of my peers are actually implementing more projects and keeping more up to date on systems to produce cost savings. Without an efficient, well-run IT department, companies can be stuck in 1998—they’re the ones that are getting hurt most.”

Tailleur at James Western Star Sterling said he’s trying to use the downturn as a way to regroup and prepare for the next upswing. “I actually think it has been a positive thing,” he said. “It gives me time to investigate different solutions and gives me a base to work from. It’s nice to be able to catch a breath and say “What’s coming?”

Andrew Reichman, senior analyst at Cambridge, Mass.-based Forrester Research, said budget cuts give IT organizations a reason to focus on efficiency instead of the latest and greatest technology. “There isn’t just a technical fix,” he said. “It’s like cleaning out your garage—you have to look for stuff you don’t need and throw it out. No automated tool can make that judgment call.

“If the pressure’s back off in six months and people are going back to [focusing on] uptime and performance rather than the efficiency, who knows if this downturn will really have an impact,” Reichman continued. “Only if people really think about this long enough, will we see real significant long-term change.” ⊙

Beth Pariseau is the senior news writer for SearchServerVirtualization.

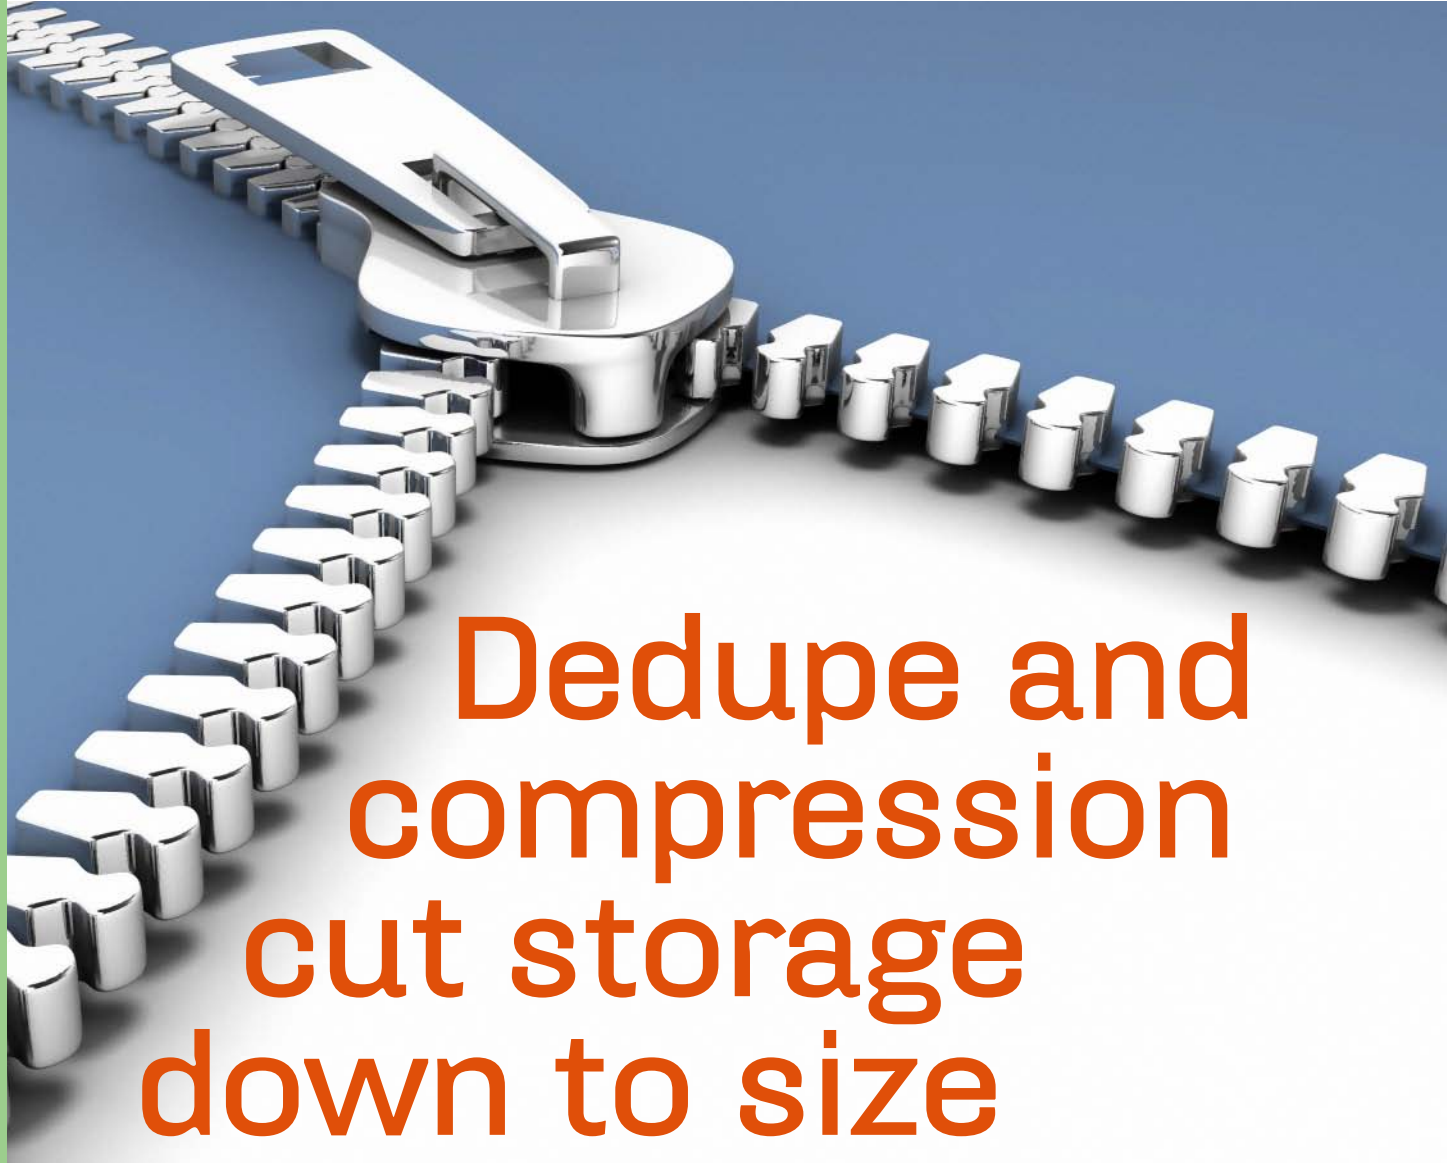

Dedupe and compression cut storage down to size

Data reduction technologies like data deduplication and compression have been well integrated into backup systems with impressive results. Now those benefits are available for primary storage data systems.

By W. Curtis Preston

USE LESS DISK, save more electricity. What's not to like? If you buy the right products, you can pare down the disk capacity your data needs and maybe even cut your electric bills by as much as 50%. That's the promise of primary storage data reduction, and while slashing utility costs is appealing, there's still plenty of skepticism about the claimed benefits of the technology. While there's little dispute that this new class of products can reduce the amount of disk your primary storage uses, uncertainty remains about whether the gains outweigh the challenges of primary storage data reduction.

The key questions about primary storage data reduction include the following:

- Why is it called "data reduction" rather than data deduplication?
- Disk is cheap. Why bother adding new technologies to reduce the size of the data it holds?
- What are the different types of data reduction for primary storage?
- How much disk space can actually be saved?

DATA REDUCTION DEFINED

In backup environments, data deduplication is a recognized and appropriate term for the technologies that eliminate redundancies in backup sets, but for primary storage, data reduction is a more accurate term because not all data reduction technologies use deduplication techniques. As an umbrella term, data reduction includes any technology that reduces the footprint of your data on disk. There are three main types of data reduction used today: compression, file-level deduplication and sub-file-level deduplication.

The main difference between primary storage and backups is the expectation of the entity that's storing or accessing the data.

Before we examine these different technologies—all of which were used for backups before they were applied to primary storage—let's look at how very different primary data storage is from backup data storage. The main difference between primary storage and backups is the expectation of the entity that's storing or accessing the data. Backups are typically written in large batches by automated processes that are very patient. These processes are accustomed to occasional slowdowns and unavailability of resources, and even have built-in technologies to accommodate such things. Backups are rarely read, and when they are, performance expectations are modest: Someone calls and requests a file or database to be restored, and an administrator initiates the restore request. Unless the restore takes an abnormally long time, no one truly notices how long it took. Most people have adjusted their expectations so that they're happy if the restore worked at all. (This is sad, but unfortunately true.) This typical usage pattern of a disk-based backup system means you could slow down backups quite a bit without a lot of people noticing.

Primary storage is very different. Data is written to primary storage throughout the day and it's typically written directly by real people who are entering numbers into spreadsheets, updating databases, storing

documents or editing multimedia files. These activities could occur dozens, hundreds or even thousands of times a day, and the users know how long it takes when they click "Save." They also know how long it takes to access their documents, databases and websites. Inject something into the process that increases save time or access time from one or two seconds to three or four seconds, and watch your IT help center light up like a Christmas tree.

This means that the No. 1 rule to keep in mind when introducing a change in your primary data storage system is primum non nocere, or "First, do no harm." Data reduction techniques can definitely help save money in disk systems, and power and cooling costs, but if by introducing these technologies you negatively impact the user experience, the benefits of data reduction may seem far less attractive.

The next challenge for data reduction in primary data storage is the expectation that space-saving ratios will be comparable to those achieved with data deduplication for backups. They won't. Most backup software creates enormous amounts of duplicate data, with multiple copies stored in multiple places. Although there are exceptions, that's not typically the case in primary storage. Many people feel that any reduction beyond 50% (a 2:1 reduction ratio) should be considered gravy. This is why most vendors of primary data reduction systems don't talk much about ratios; rather, they're more likely to cite reduction percentages. (For example, a 75% reduction in storage sounds a whole lot better than a 3:1 reduction ratio.)

If you're considering implementing data reduction in primary data storage, the bottom line is this: compared to deploying data deduplication in a backup environment, the job is harder and the rewards are fewer. That's not to suggest you shouldn't consider primary storage data reduction technologies, but make sure you properly set expectations before making the commitment.

Is archiving data reduction?

Some vendors consider archiving and hierarchical storage management (HSM) to be data reduction technologies. Both archiving and HSM systems can reduce the amount of disk you need to store your primary data, but they do so by moving data from one storage system to another. While they may save you money, they're not truly reducing the size of the data—they're just moving it to less-expensive storage. Therefore, while these are good technologies that companies with a lot of data should explore, it's not data reduction per se.

PRIMARY STORAGE DATA REDUCTION TECHNOLOGIES

Compression. Compression technologies have been around for decades, but compression is typically used for data that's not accessed very

much. That's because the act of compressing and uncompressing data can be a very CPU-intensive process that tends to slow down access to the data (remember: primum non nocere).

There's one area of the data center, however, where compression is widely used: backup. Every modern tape drive is able to dynamically compress data during backups and uncompress data during restores. Not only does compression not slow down backups, it actually speeds them up. How is that possible? The secret is that the drives use a chip that can compress and uncompress at line speeds. By compressing the data by approximately 50%, it essentially halves the amount of data the tape drive has to write. Because the tape head is the bottleneck, compression actually increases the effective speed of the drive.

Compression systems for primary data storage use the same concept. Products such as Ocarina Networks' ECOsystem appliances and Storwize Inc.'s STN-2100 and STN-6000 appliances compress data as it's being stored and then uncompress it as it's being read. If they can do this at line speed, it shouldn't slow down write or read performance. They should also be able to reduce the amount of disk necessary to store files to between 30% and 75%, depending on the algorithms they use and the type of data they're compressing. The advantage of compression is that it's a very mature and well understood technology. The disadvantage is that it only finds patterns within a file and doesn't find patterns between files, therefore limiting its ability to reduce the size of data.

File-level deduplication. A system employing file-level deduplication

PRIMARY STORAGE DATA REDUCTION VENDORS

VENDOR	PRODUCT	TYPE OF DATA DEDUPLICATION	COMPRESSION
EMC Corp.	Celerra NS-120, NS-480, NS-g60, NS-G8 Gateway	File-level	Yes
Exar Corp.	Hifn BitWackr B1605R	Sub-file-level	Yes
GreenBytes Inc.	GB-X Series	Sub-file-level	Yes
Microsoft Corp.	Windows Storage Server 2008	File-level	No
NetApp Inc.	Ontap operating system for FAS Filers	Sub-file-level	No
Nexenta Systems Inc.	NexentaStor 3.0	Sub-file-level	Yes
Ocarina Networks	ECOsystem and ECOvault	Sub-file-level	Yes
Oracle-Sun	ZFS file system	Sub-file-level	Yes
Storwize Inc.	STN Series appliances	None	Yes

examines the file system to see if two files are exactly identical. If it finds two identical files, one of them is replaced with a link to the other file. The advantage of this technique is that there should be no change in access times, as the file doesn't need to be decompressed or re-assembled prior to being presented to the requester; it's simply two different links to the same data. The disadvantage of this approach is that it will obviously not achieve the same reduction rates as compression or sub-file-level deduplication.

Sub-file-level deduplication. Sub-file-level deduplication is very similar to the technology used in hash-based data deduplication systems for backup. It breaks all files down into segments or chunks, and then runs those chunks through a cryptographic hashing algorithm to create a numeric value that's then compared to the numeric value of every other chunk that has ever been seen by the deduplication system. If the hashes from two different chunks are the same, one of the chunks is discarded and replaced with a pointer to the other identical chunk.

Depending on the type of data, a sub-file-level deduplication system can reduce the size of data quite a bit. The most dramatic results using this technique are achieved with virtual system images, especially virtual desktop images. It's not uncommon to achieve reductions of 75% to 90% in such environments. In other environments, the amount of reduction will be based on the degree to which users create duplicates of their own data. Some users, for example, save multiple versions of their files on their home directories. They get to a "good point", save the file, and then save it a second time with a new name. This way, they know that no matter what they do, they can always revert to the previous version. But this practice can result in many versions of an individual file, and users rarely go back and remove older file versions. In addition, many users download the same file as their coworkers and store it on their home directory. These activities are why sub-file-level deduplication works even within a typical user home directory.

The advantage of sub-file-level deduplication is that it will find duplicate patterns all over the place, no matter how the data has been saved. The disadvantage of this approach is that it works at the macro level as opposed to compression that works at the micro level. It might identify a redundant segment of 8 KB of data, for example, but a good compression algorithm might reduce the size of that segment to 4 KB. That's why some data reduction systems use compression in conjunction with some type of deduplication.

A SAMPLER OF PRIMARY STORAGE DATA REDUCTION PRODUCTS

The following vendors currently offer primary storage data reduction products (listed in alphabetic order):

EMC Corp. EMC introduced file-level deduplication and compression

of inactive files in its EMC Celerra filer. Administrators can configure various settings, such as how old a file must be before it's a candidate for this process, and what file sizes the process should look for. While deduplication and compression of older data obviously won't generate as much data reduction as compressing or deduplicating everything, EMC customers have reported significant savings using this data reduction implementation.

Exar Corp. Exar gained data deduplication technology with its April 2009 acquisition of Hifn Inc. End users may be unfamiliar with Exar, but they may already be using their products. Many high-end virtual tape libraries (VTLs) and data deduplication systems for backups use Exar hardware compression cards for data compression. Exar now has released a card, which is designed to be placed into a Windows or Linux server, that will deduplicate data as it's being written to any hard drive. Exar's Hifn BitWackr B1605R is a hardware and software product that offloads data deduplication and compression from a server's CPU and makes adding data reduction to a Windows or Linux server a relatively easy process.

GreenBytes Inc. GreenBytes is in a unique position; it's the first vendor attempting to make a single product to address the data reduction needs of both backup and primary data storage in its GB-X Series of network-attached storage (NAS) and storage-area network (SAN) storage devices. The firm uses a hash-based data deduplication technology, but the hash algorithm is different from that used by all other vendors: Instead of the widely used SHA-1, GreenBytes uses Tiger, which it says is more suited to general-purpose processors than SHA-1 and, therefore, offers significant performance advantages while not decreasing data integrity. Tiger's key space (192 bits) is significantly larger than that of SHA-1 (160 bits), which further reduces the chances of a hash collision. GreenBytes is also making extensive use of solid-state disk as a cache in front of SATA disk so that it can better meet the performance needs of primary data storage users.

Microsoft Corp. With its Windows Storage Server 2008, Microsoft offers file-level single-instance deduplication built into the operating system. A number of storage systems vendors are taking advantage of the built-in SIS, including Hewlett-Packard's StorageWorks X-series Network Storage Systems and Compellent's Storage Center with NAS. File-level deduplication alone will provide modest space savings for users of these systems.

NetApp Inc. NetApp was the first primary data storage vendor to offer deduplication, which leverages the company's existing write anywhere file layout (WAFL) file system technology. The WAFL file system

already computes a CRC checksum for each block of data it stores, and has block-based pointers integrated into the file system. (It's the secret behind NetApp's ability to have hundreds of snapshots without any performance degradation.) An optional process runs during times of low activity and examines all checksums; if two checksums match, the filer does a block-level comparison of those blocks. If the comparison shows a complete match, one of the blocks is replaced with a WAFL pointer. The result is sub-file-level deduplication without a significant impact on performance. NetApp's deduplication system has been tested by many users against multiple data types, including home directories, databases and virtual images, and most users have reported positive results in both reduction percentages and performance. As of this writing, NetApp uses only data deduplication and doesn't do compression.

Nexenta Systems Inc. Nexenta uses the Oracle Solaris ZFS file system in its NexentaStor family of storage system software products that are based on the open source OpenSolaris platform. However, the firm has added more than 30 additional features to its ZFS-based offering that are only available from Nexenta. Examples of these features include an integrated management console, LDAP integration, continuous data protection (CDP) and synchronous replication. The NexentaStor 3.0 offers deduplicated storage that's fully integrated with Citrix Systems Inc. XenServer, Microsoft Hyper-V and VMware Inc. VMware vSphere.

Ocarina Networks. Ocarina takes a very different approach to data reduction than many other vendors. Most vendors apply compression and deduplication without any knowledge of the data, but Ocarina has hundreds of different compression and deduplication algorithms that it uses depending on the specific type of data. For example, the company uses completely different techniques to compress images and Word documents. It also understands encapsulation systems such as the Digital Imaging and Communications in Medicine (DICOM) system. Ocarina will actually disassemble a DICOM container, examine and deduplicate the various components, and then reassemble the container. As a result, Ocarina can often achieve much greater compression and deduplication rates than other vendors can realize with the same data types.

Ocarina isn't a storage vendor; it works with existing data storage system vendors that will allow Ocarina to interface with their systems. Ocarina is currently partnering with BlueArc Corp., EMC, Hewlett-Packard (HP) Co., Hitachi Data Systems and Isilon Systems Inc.

Oracle-Sun. Oracle's Solaris ZFS file system also has sub-file-level data deduplication built into it. As of this writing, there's not much available information about how well it duplicates data or its perform-

ance in user production environments. However, the ZFS website does state that there shouldn't be a significant difference in performance between deduplicated and native data, as long as the hash table used for deduplication can fit into memory.

Storwize Inc. Storwize offers an appliance that sits in-band between your NAS (NFS/CIFS) filer and the systems accessing it. It works like the chips in your tape drives that compress data real-time as it's being written and uncompresses it real-time as it's being read. Like the tape drive chip, it doesn't hurt performance because it's compressing the data at the same speed it's arriving at the system. In fact, with certain applications it can even increase performance. One other interesting thing about their system is that files they compress simply appear as compressed files on the filer. Therefore, if your worst fears happened and Storwize disappeared one day, all you would need is enough space to uncompress the files using standard algorithms.

DATA REDUCTION IS STILL NEW AND GROWING FAST

A little over a year ago, there were virtually no viable options for reducing data in primary storage. Now there are half a dozen or so, with more on the way. Given the runaway growth in file storage that most companies are experiencing, it shouldn't take long for data reduction technologies to find their way into many of the products offered by data storage systems vendors. ⊙

W. Curtis Preston is an executive editor at TechTarget and an independent backup expert.

TIERED STORAGE: A LOOK AT INTERNAL AND EXTERNAL TIERED STORAGE MODELS

For organizations hoping to achieve information lifecycle management (ILM) goals, finding a solution that meets their specific policies, processes and practices can be difficult. Learn how tiered storage tools can help organizations meet these goals in a cost-effective and practical way. By Ron Scruggs

BY ASSIGNING DIFFERENT categories of data to different types of enterprise data storage media, tiered storage tools aim to reduce overall storage costs. While tiered storage is not a complete information lifecycle management solution, it does play an important role in a successful ILM hierarchy. As one vendor put it, "ILM is all about having the right storage at the right time at the right cost."

For tiering data within a storage array, 3PAR Inc., Compellent Technologies Inc., EMC Corp. and Hitachi Data Systems (HDS) Corp. all offer a LUN virtualization component, which allows data stored on high-performance disk to be migrated online to lower-performance disks and vice versa based on pre-set policies and I/O patterns.

All of these storage vendors support some sort of online LUN migration technique and all require a separate piece of hardware or a virtual machine (VM) that monitors the storage and initiates LUN migrations based on pre-set policies.

On the software solution side, 3PAR implements storage tiering through its Utility Data Lifecycle Management (DLM) product. Utility DLM uses templates, virtual copy and remote copy services to migrate data between tiers on the array both locally and remotely. The Utility DLM software can run on a Linux or Windows workstation, and all of the heavy lifting occurs on the array.

For EMC's Virtual LUNs, data can be seamlessly migrated between RAID levels and disk types. The Virtual LUN technology can be managed by EMC's Symmetrix Control Center (SCC) if the customer is using a Symmetrix V-Max. EMC's Fully Automated Storage Tiering (FAST), or its Navisphere Management Suite (NMS) for the CLARiiON line can be used to manage the virtual LUN migration based on policies and data access patterns for the Clariion CX-4 line. Both SCC and NMS require Windows hosts.

Hitachi also has a software-based tiering solution simply called Tiered Storage Manager, which supports LUN migration between tiers in a virtualized storage pool. The pool can include other vendors' arrays attached to the Hitachi USP storage array platform.

External tiered storage tools include IBM's SAN Volume Controller (SVC), EMC's Invista and Hitachi's USP-V.

SVC is an xSeries based "I/O engine" architecture that acts as a front end to all SAN storage. The SVC can front-end any Fibre Channel-attached storage array regardless of the vendor and can provide iSCSI connectivity on the front end as well as Fibre Channel. LUNs can be seamlessly migrated within and across back-end arrays, allowing customers to provide one vendor for tier 1, another for tier 2 and so on if their environment dictates heterogeneous storage.

The SVC can front-end any Fibre Channel-attached storage array regardless of the vendor and can provide iSCSI connectivity on the front end as well as Fibre Channel.

EMC's Invista also utilizes external nodes, but the twist is that it also uses an intelligent SAN switch (or linecard), and each LUN on the back-end array has a virtual address that is presented to the host. EMC considers this solution to be operating at Layer 2 in the Open Shortest Path First (OSPF) model and is touted as "network-based volume management."

Hitachi's USP-V is the company's high-end line of storage arrays; the LUN virtualization occurs within the storage controllers and does not require any external components. Typically, customers include the Hitachi

Hi-Command server, which is comparable to EMC's NMS or SCC. Hi-Command and the Tiered Storage Manager software are required for LUN migrations, and these migrations can be performed between arrays online with no disruption to hosts.

All of the tools discussed have been SAN infrastructure-level components and can certainly be considered ILM. All of the tools are policy-based, either by application requirements, age requirements or cost requirements; all are storage-centric; and all support live data migration in a transparent way. To be successful, storage tiering tools should have as little impact as possible on the user's experience, and the storage vendors are clearly showing that they can provide this service to aid IT in its pursuit of ILM. ⊙

Ron Scruggs was formerly an analyst with GlassHouse Technologies.

How to improve power efficiency in archive hardware and on primary storage

Storage efficiency is usually taken to mean capital expenditure or capacity savings, but recent research shows that some users are willing to spend for power efficiency. *By George Crump*

EVEN WITH A sour economy, power efficiency is still a hot-button issue among storage buyers. In fact, recent research shows that a sizable chunk of customers are willing to spend money on IT solutions that are power efficient. For example, the spring 2009 Purchasing Intentions survey conducted by *Storage* magazine showed that while only 3.5% of 566 storage buyer respondents said energy efficiency was the most important factor in the choice of a disk array, about 32% said it was a major factor in the decision. But even customers who understand the importance of power efficiency will need help to actually realize those power savings. So how do we convince them to go from Point A (a data center full of power-hungry hardware) to Point B (a much less ravenous room of equipment)? And can it be done on a range of data storage devices: backup, archive and primary storage hardware?

First, keep in mind this simple principle: If power-managed disks in a storage system are not being accessed for a period of time, they either spin down or power down, which saves power and cooling costs. The challenge, however, is that the data on these systems has to be inactive long enough that the drives will actually spin down. “Long enough” is a user-settable time period, often in minutes.

BACKUP, ARCHIVE POWER EFFICIENCY

In a backup-to-disk scenario, it's almost a no-brainer: Power-managed drives will often be inactive long enough so that they could go into a power saving mode. If the backup window is eight to 10 hours and there is minimal restore activity, a backup-to-disk system should be able to power down for four to six hours per day, adding up to a measurable power savings.

Beyond backup, power savings can also be found from hardware used for archiving. While some users in the storage industry argue that data archives don't belong on disk, partly owing to the high energy consumption of those disks compared to tape, it's relatively simple to decrease the power consumption of disk archives in power-managed drives by powering them down when idle.

And since the purpose of an archive is to store a customer's less frequently accessed data, there should be plenty of idle time on disk drives. To maximize the amount of idle time on a disk archive, solution providers should migrate data to the archive just once a week. More importantly, you should put tools in place (from companies such as Tek-Tools and Aptare) to monitor data being accessed from the archive to determine whether there's data that actually belongs on primary storage rather than on the archive disk.

To maximize the amount of idle time on a disk archive, solution providers should migrate data to the archive just once a week.

Occasional access of a file from archive is fine, but if a file shows up as very active, a de-migration back to primary storage might make more sense. By weeding out the files that are being frequently accessed from archive, you can increase the amount of idle time on the archive and thereby increase the power savings.

POWER EFFICIENCY ON PRIMARY STORAGE

In addition to backup and archive scenarios, customers can also decrease their power usage on primary storage. For instance, the layout of physical drives in array groups can be designed for optimal use of the power saving capabilities of these drives.

Primary storage will need to be monitored to assess access frequencies of data, with an eye toward grouping data with similar access frequencies together. That monitoring can be done using the tools mentioned for backup and archive efficiency.

When monitoring primary storage to find good candidates for power savings, look for folders and volumes that are not accessed for a few hours, especially overnight. An example might be user home directories. After the users go home for the night and the backups of those systems are complete, they may sit idle for hours at a time. Moving the volume

that contains the home directories from an array group with active drives to an array group that is on top of power-managed drives could save quite a bit of power. Identifying these types of poor associations between volumes and drives requires tools and storage experience but could save customers a significant amount of money because of reduced power consumption. And it allows the VAR to shine.

This leads to another way to fine-tune the power efficiency of your customer's environment: arranging the backup schedule so that power-managed areas of storage are backed up first at the end of the day, right after users go home, so those drives can go into idle mode and not be woken by the backup job.

An alternative way to achieve power efficiency is using a method called "Archiving in Place." If you provide only power-managed arrays to your clients when they buy new equipment, as they add new arrays to their environment, only move the most active data to the new environment. Instead of archiving the oldest 80%, promote the newest 20% to the new array. Then, over time, as the older data on the older array becomes accessed less frequently, the system begins to power down on its own and eventually will become extremely power-efficient.

While spin-down drives can have a big impact on a customer's power usage, in order to have the highest level of power efficiency possible, re-architect what data lives where within the data center, involving changes around backup, archive and primary storage. ⊙

George Crump is president and founder of Storage Switzerland, an IT analyst firm focused on the storage and virtualization segments.

Check out the following resources from our sponsor:

Data Center Optimization, Put Data Growth in its Place

Deduplication: How to Store it Only Once

Archiving and Backup Demos