Managing the information that drives the enterprise

STORAGE

Vol. 7 No. 12 January 2009

RECOVERY SITE OPTIONS

Hot, warm or cold

When the heat is on, the type of disaster recovery site you choose may determine how soon your company is back in business.

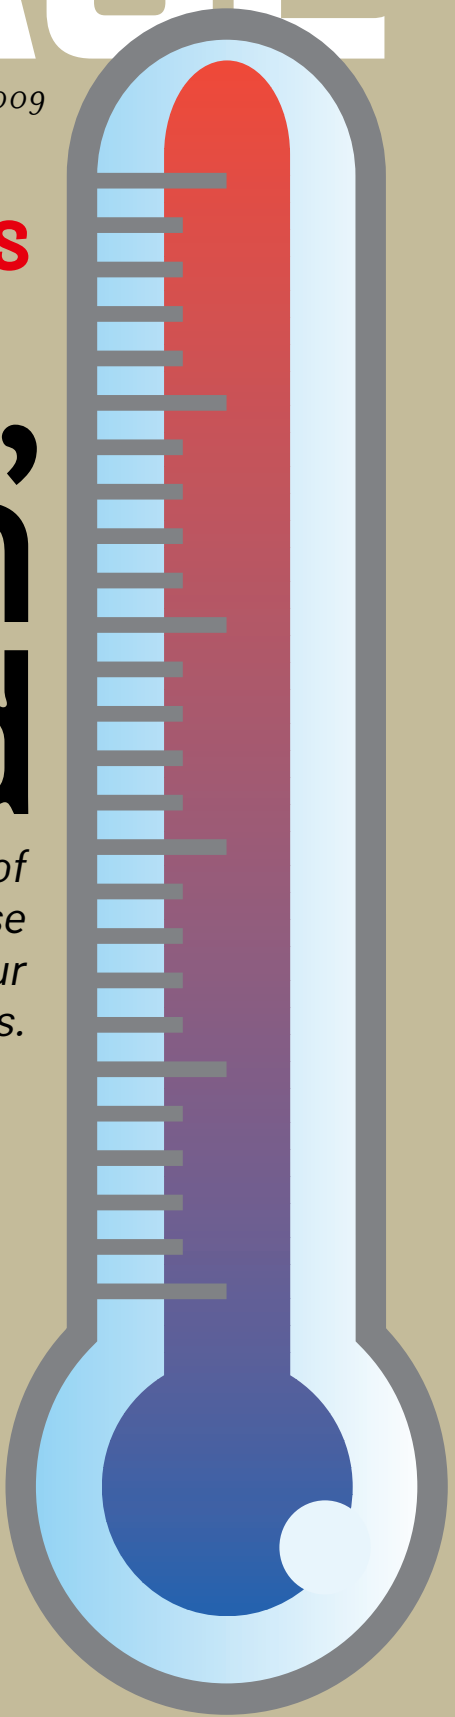

5 Get a grip on growth

8 DR site options

18 Consolidate archives

27 Rev up RAID rebuilds

31 Go green, save green

TechTarget
The IT Media ROI Experts

Solid state looks good on the showroom floor

I SWORE I'D NEVER USE another automobile analogy to describe storage equipment or computing gear after I committed the unforgivable faux pas of referring to a 1957 Chevy Impala in a previous column. More than a few readers were quick to correct me—Chevy introduced the Impala in 1958—and I worried that I might have lost credibility over an old car and one lousy year.

But comparing cars to computers has been a staple of high-tech journalism, and it's really hard to resist.

Solid-state storage has been getting a lot of attention lately, not just on these pages, but in scores of storage vendor announcements. I've said before that solid state is the future of storage, and I'm not backing down from that prediction. But right now, shopping for solid state is akin to fogging up the showroom window while ogling that Ferrari 599 GTB Fiorano with its 620 horsepower engine. For a few moments at least, it's easy to forget that the fire engine-red piece of art costs more than $300,000 and gulps fuel at a rate of approximately 10 mpg. While your heart says "Ferrari," your head knows that the $18,000 Mini Cooper—with less than a quarter of the horsepower but nearly three times the mileage—is the car of your future.

Shopping for solid state is akin to fogging up the showroom window while ogling that Ferrari 599 GTB Fiorano with its 620 horsepower engine.

And so it goes with solid-state disks (SSDs) right now. Sure, some companies with high-end, ultra-critical apps can justify shelling out the big bucks to buy the Ferraris of storage, but for 99.9% of companies and their applications toiling away out there, SSDs are still a future thing.

One argument says the future is now, as nearly all of the major storage vendors offer some sort of SSD options in their arrays. EMC was the first of the biggies to go solid state and others soon followed, with Hitachi announcing last month that it will offer STEC solid-state storage in its USP V and USP VM enterprise arrays.

There's nothing wrong with those vendors offering high-speed, low-energy storage options for their well-heeled customers but, at best, these are mostly stop-gap solutions—a way for vendors to say "We're ahead on SSDs," rather than having significant technology advancements. That's not to say customers of these SSD-adorned arrays won't

get a bang for their buck, it's just that these arrays represent short-term maneuvers.

A lot has to happen for SSDs to move into mainstream storage. For starters, the price has to come down to where it at least sits on the same budget sector with hard disk alternatives. But the technology itself needs to grow up, too. Although some of today's products make it look like SSDs are a plug 'n play deal, it might not be the best approach. Disk controller systems and software built to run hard disks are likely to need significant alterations to accommodate solid-state storage most effectively.

This isn't to say that storage and solid-state vendors are simply buffing the fenders of the Ferrari. Companies like Fusion-io are hard at work under the hood of solid-state storage, tinkering with the underlying engineering and coming up with more effective implementations. Micron Technology teamed with Sun to address the lifespan issue of SSDs, and recently announced that their joint effort achieved SLC NAND flash that can handle a million write cycles.

Hewlett-Packard is also bullish on solid state, predicting in a recent email announcement that the technology will gain momentum this year. This echoes IDC's prediction of 70% compound growth for SSDs by 2012. Pliant Technology, another solid-state engineering company working on technology specifically targeted at enterprise storage needs, recently announced its participation in SNIA's new Solid State Storage Initiative (SSSI). SSSI is a veritable "Who's Who" of memory and storage vendors, and while the SNIA stamp may not have all that much practical impact at this time, it does reflect its members' level of interest.

You may not see an $18,000 Ferrari anytime soon, but it's probably a good idea to keep your eye on developments in solid-state storage as it travels in the right direction.

Rich Castagna (rcastagna@storagemagazine.com) is Editorial Director of the Storage Media Group.

***Storage* magazine**
117 Kendrick Street,
Suite 800
Needham, MA 02494
editor@storagemagazine.com

Subscriptions:
www.SearchStorage.com

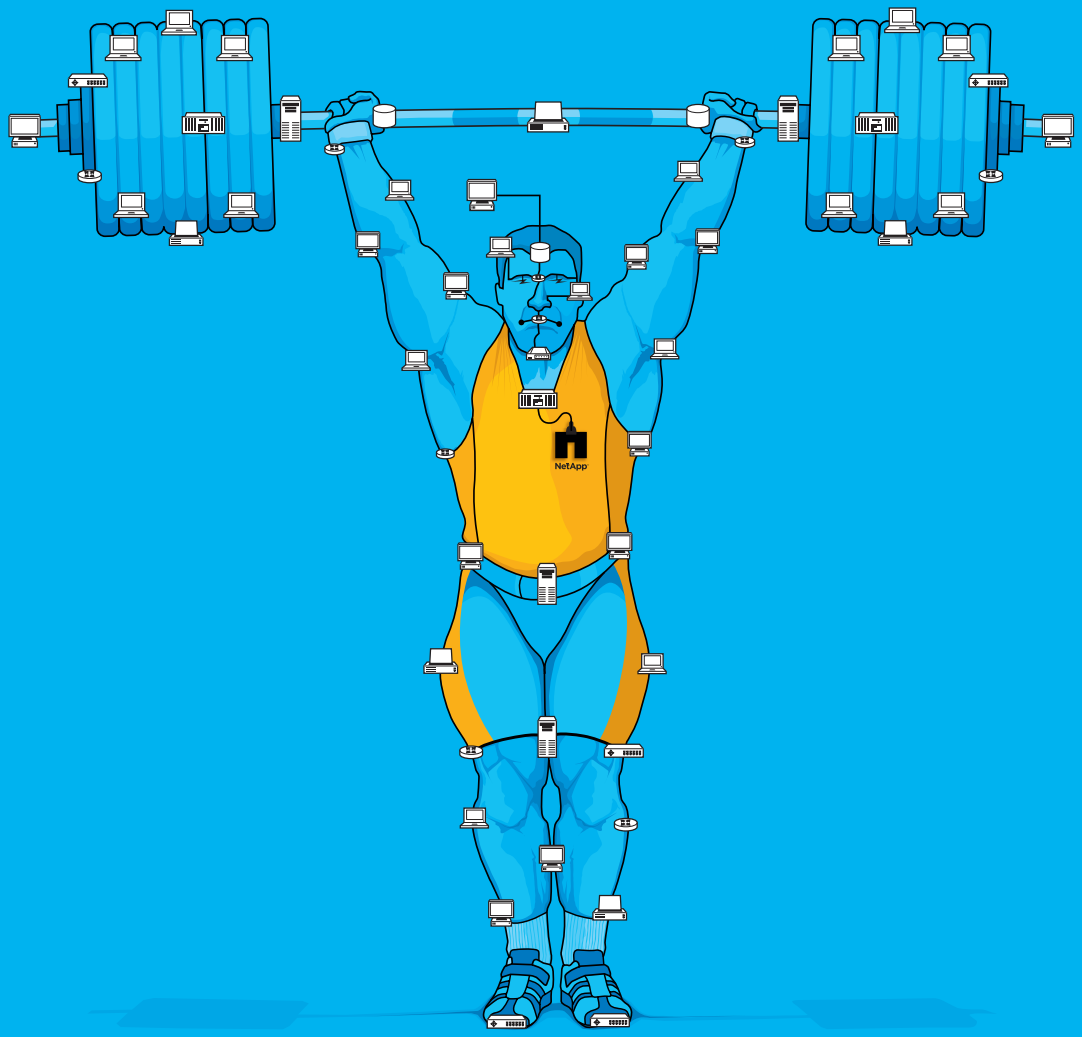

With NetApp® at the heart of your business, you can

EFFORTLESSLY HOLD YOUR DATA

at a fraction of the cost and footprint.

Imagine storage and data management solutions smart enough to support the data you need, and not a lot of dead weight. It's possible when you partner with NetApp. Our industry-leading solutions use deduplication and other space-saving technologies to help you store data efficiently and reduce your footprint by 50% or more. So you can manage exponential growth while minimizing your storage investment—all with the support of a team that will exceed your expectations. See how we can help your business go further, faster. Find out how you can use 50% less storage, guaranteed,* at netapp.com/efficiency.

NetApp™
Go further, faster

Killing off infectious data

Better start thinking about your data growth in deadly terms.

MANY OF THE PROBLEMS we face in our attempt to manage a data center are a direct result of data growth. Data growth is constant, and it sometimes seems intent on destroying everything in its path. Unaddressed data growth will wreak havoc on your file system, disk, system, network, protection plans, processes and life. If you're like a lot of people, you might try to stay ahead of this never-ending cycle of growth by buying more of whatever is going to break next.

I think it's time we address the cause and not the symptoms. There's new data generated all the time, but most of it is generated by our own processes. We have data sprawl, replicas, copies of copies, backup copies of copies, and backups of replicas of copies of copies. We don't have a capacity problem, we have a science problem.

There's a process in biology called mitosis in which one cell divides to produce two genetically identical cells. Left unchecked in the right environment, those cells will split again and again. Soon, the petri dish that stored a microscopic quantity of stuff is overflowing all over the table. If a scientist acted like an IT guy, they would address this issue by pouring (migrating) the contents of the petri dish into bigger and bigger containers before they overflowed.

Data growth is constant, and it sometimes seems intent on destroying everything in its path.

Originally, this science made sense. Scientists needed a bunch of exact replicas of a single cell to perform different tests or experiments on them. In IT the same holds true; we need a bunch of replicas of data to run different applications against them. We use these replicas to run tests, populate data warehouses, create backup and disaster recovery copies, and to send copies to other users. The difference is that scientists know up front how many replicas they want/need and plan for it. But IT processes seldom have the preplanning that exists in science labs. And that, my friends, is a huge part of our problem. When scientists are done with their experiments, they get rid of the replicas. Our answer to the challenge is to buy a bigger petri dish from our sales rep.

We know that Data Domain proved empirically that killing replicate data in the backup process is a very good thing. There are now a thousand dedupe stories to be told and they all share one theme: Killing data when it's no longer useful is a good thing.

So if killing off replicas at the end of the data lifecycle is good, killing them sooner would be even better. That's the next frontier.
If you get rid of replicas as soon as they're no longer valuable (and before they have a chance to cause problems), you eliminate problems associated with biological replication. Killing, compacting, deduplicating, eliminating or compressing replicate data as close to the point of conception as feasible will yield the greatest possible benefits downstream. It's only logical.

How will you do this? First, you'll have to address process and strategy requirements; i.e., actually know how many copies you need and for how long, as well as have an actual plan on how to deal with them. Second, you'll have to leverage technology that can wipe out copies before they take over. These multiple copies are like the cockroaches of IT. Eventually cockroaches win and you have to move out.

Dedupe in the backup target market has created more than $2 billion in value (and growing), so imagine what value will be generated by moving that function closer to the point of creation for all of the different data types we generate. We'd be green (less data is as green as it gets), rich (we wouldn't need to buy anything new for a while), calm (less things to manage equals less things to break) and might actually be able to take a few minutes to think about how we can add strategic value to our organization, as opposed to running around in a hazmat suit all day dumping out petri dishes. ⊙

Steve Duplessie is founder and senior analyst at Enterprise Strategy Group. You can see his blog at http://esgblogs.typepad.com/steves_it_rants/.

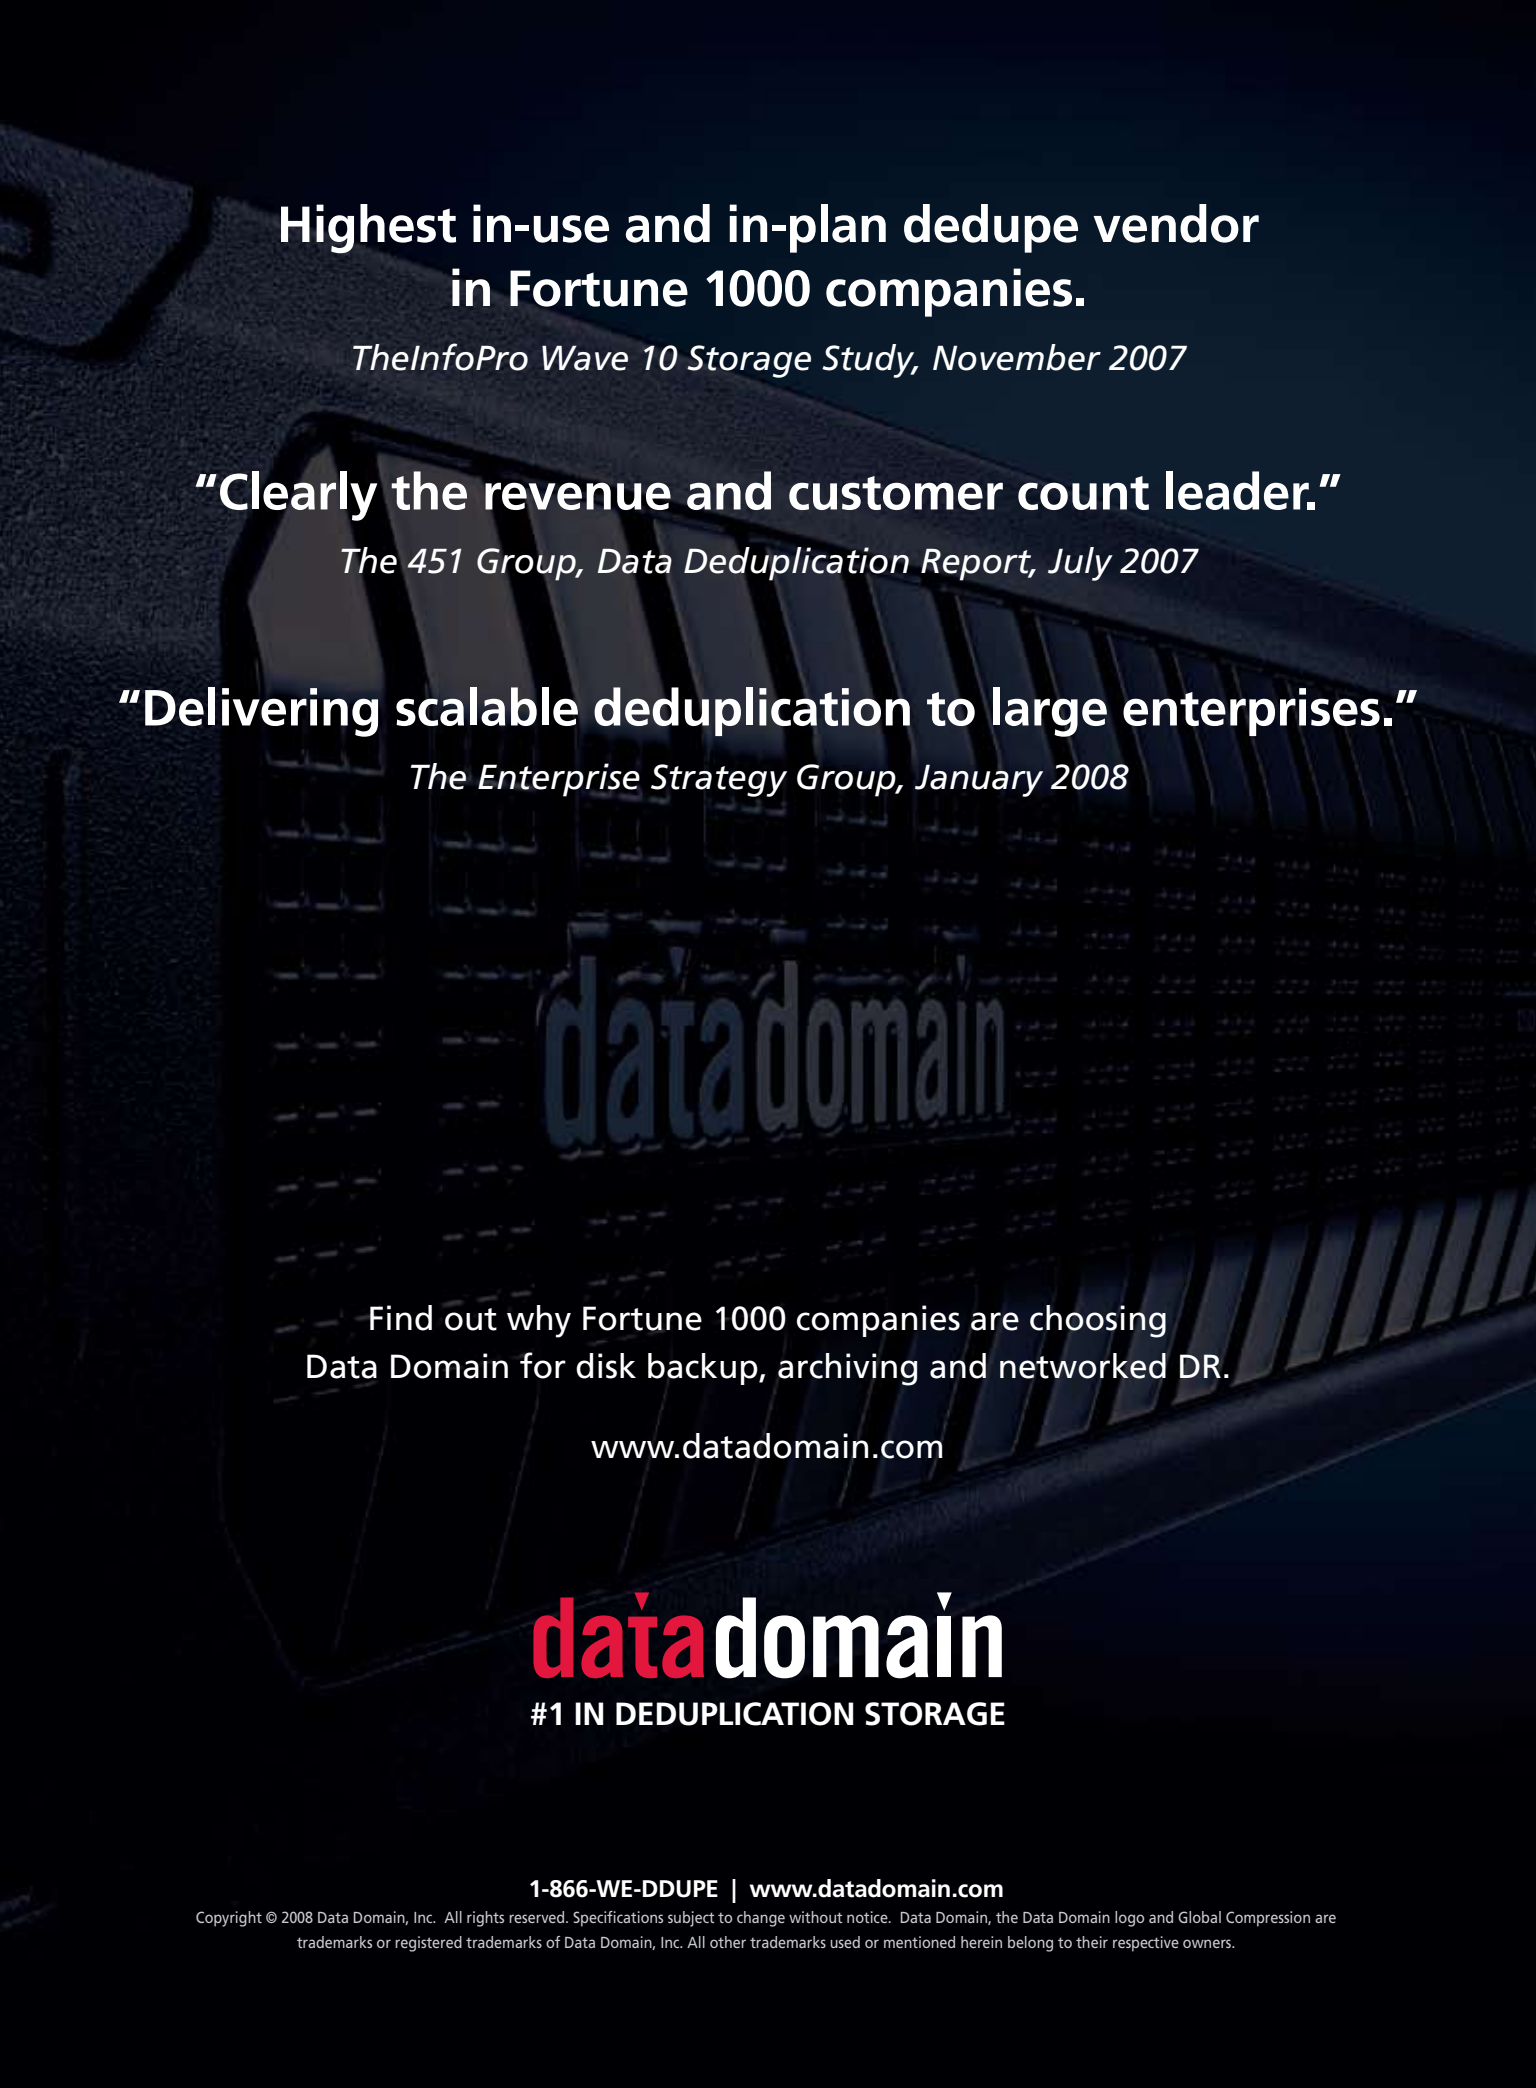

Highest in-use and in-plan dedupe vendor in Fortune 1000 companies.

TheInfoPro Wave 10 Storage Study, November 2007

"Clearly the revenue and customer count leader."

The 451 Group, Data Deduplication Report, July 2007

"Delivering scalable deduplication to large enterprises."

The Enterprise Strategy Group, January 2008

Find out why Fortune 1000 companies are choosing
Data Domain for disk backup, archiving and networked DR.

www.datadomain.com

datadomain

#1 IN DEDUPLICATION STORAGE

1-866-WE-DDUPE | www.datadomain.com

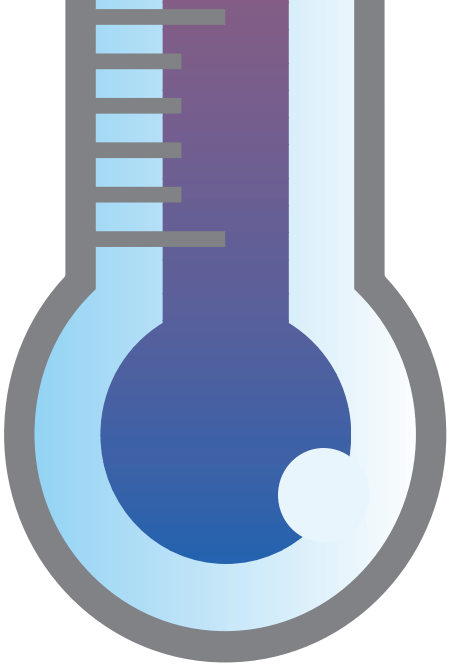

Disaster recovery site options

Hot site, warm site or cold site?

Here's how to figure out the best disaster recovery strategy for your company.

By Jacob Gsoedl

THE ABILITY TO RECOVER from a disaster in an acceptable period of time is a critical issue for companies with increasing dependence on information technology. Once thought to be a concern for only larger organizations, being able to recover mission-critical applications within a predictable timeframe is a mandate for any size company today. But some users see disaster recovery (DR) as a pricey insurance policy, and may take shortcuts to try and save a few dollars. To avoid becoming victims of budget cuts, DR provisions and sites must be built around a few basic principles that allow management to decide what's required while candidly showing the possible business impact and consequences of retrenchments.

GUIDING PRINCIPLES

Recovery time objective (RTO) and recovery point objective (RPO) are the key metrics to determine the DR level required to recover business processes and applications. They are reciprocally proportional to the cost of DR: The closer RTO and RPO need to be to zero, the more expensive DR provisioning will be. If recovery time can be days or even weeks, costs will likely be significantly less.

Determining the necessary RTOs and RPOs is the single most important exercise a business needs to perform to ensure the right level of DR without wasting money. RTOs and RPOs are derived through business impact analysis of business processes and applications to determine the value of business processes and the anticipated financial impact if they become unavailable. Obviously, this varies greatly by business process and application. "While for just-in-time manufacturing the critical threshold may be 15 minutes, it could be days for a marketing application," says George Ferguson, worldwide service segment manager for Hewlett-Packard (HP) Co.'s business continuity and recovery services.

Very likely, determining RTOs and RPOs will be an iterative process because of two competing forces: available budget and required recovery objectives. "The challenge of contingency services like disaster recovery is to find the right balance between available budget and what's required to sustain the business," says Greg Schulz, founder and senior analyst at StorageIO Group, Stillwater, MN.

Disaster recovery site terminology can be confusing—terms like hot site, warm site and cold site are common in DR parlance, but they're used inconsistently.

DISASTER RECOVERY OPTIONS

With a business impact analysis in hand and agreement on RTOs and RPOs, IT management can devise implementation options. Disaster recovery site terminology can be confusing—terms like hot site, warm site and cold site are common in DR parlance, but they're used inconsistently. A hot site in the U.S. typically comprises shared equipment, while "in Europe the term hot site is predominantly used for dedicated equipment," says Ferguson. The following definitions match the prevailing U.S. interpretations of these terms:

- **Hosted site.** A site with dedicated equipment; required whenever RTO and RPO need to be close to zero.
- **Hot site.** Uses shared equipment with dedicated storage and real-time replication; a typical RTO of a few hours.
- **Warm site.** Uses shared equipment without dedicated storage, but depends on data backup for recovery; RTOs can range from a few hours to days depending on the backup method in use.
- **Cold site.** Typically, dedicated space in a data center fully loaded with cooling, power and connectivity ready to accept equipment; RTOs are usually a week or more.

It's quite common for a DR site to serve various roles for different applications. For instance, a DR site may serve as a hosted site with close to real-time failover for a mission-critical e-commerce application,

and it may also serve as a low-end warm site with tape-based recovery for a less critical engineering application. Many DR sites are hybrids where the application determines the role of the site. As a result, disaster recovery companies that host DR sites typically offer their services in tiers that can be mapped to RTOs and RPOs required by applications (see "DR tiers," below).

HOSTED SITE

A tier 1 DR offering provides the highest level of DR protection, and is typically used for applications that require close to zero RTOs and RPOs. A characteristic of tier 1 DR is the use of dedicated equipment in

DR tiers

Comparison of disaster recovery (DR) tiering options for a 5 TB Microsoft Exchange 2007 environment using plans and pricing from Recovery Point Systems as an example.

DR tier	*DR method*	*Data protection method*	*Average RTO*	*Average RPO*	*Price per month*
Tier 1	Mirrored failover with history	Immediate failover to redundant production environment, including servers, and real-time and duplicated historical data	Minutes	Near zero	$25,000
Tier 1	Mirrored failover without history	Immediate failover to redundant production environment, including servers and real-time data only	Minutes	Near zero	$18,750
Tier 2	Hot site with real-time data replication	Real-time disk-to-disk replication with shared CPUs on standby	Four hours	Near zero	$12,500
Tier 3	Warm site V1 with point-in-time disk backup	Point-in-time data backup to remote disk with shared CPUs on standby	12 hours to 24 hours	Hours	$6,250
Tier 4	Warm site V2 with point-in-time tape backup	Point-in-time data backup to tape with shared CPUs on standby; tape-based recovery from physically delivered media	24 hours to 48 hours	One or more days	$3,750
Tier 5	Cold site	Empty data center fully loaded with cooling, power and connectivity, ready to accept equipment	Seven days or longer	Days, depending on the backup used for the restore	$1,250

Source: Recovery Point Systems

the DR site. As a result, it carries the highest price tag and is usually only for the most mission-critical applications. Because the equipment in the DR site is dedicated to a single client company, there are very few constraints on the equipment that can be used, even if the service is outsourced.

Among all DR options, it's best suited to be hosted in-house, where your company owns and maintains the site. Because of the need for specific DR equipment, it's typically less expensive to build tier 1 DR in-house than to outsource it. "Tier 1 DR can be done more cost-effectively in-house, especially if you have the facility and people," explains HP's Ferguson.

Because applications in the primary and DR site are closely coupled, production and DR equipment are commonly managed by one entity. If the DR site is hosted by a third party, it's not unusual for both the primary and DR equipment to be managed by the DR services provider. As an example, Citrix Systems Inc. decided to outsource management of both its primary HP XP12000 SAN and its DR site. While the production SAN physically resides in Citrix's primary data center in Miami, the DR SAN is hosted by HP. "Our SAN storage in Miami is outsourced with and managed by HP," says Michael Emerson, director, IT security, governance and business continuity at Citrix. "They own the SAN and manage it, including the replication from the production to the DR SAN at HP, using HP Continuous Access replication [HP StorageWorks XP Continuous Access Software]."

"Hot sites and warm sites can be implemented less expensively through outsourcing than doing them in-house because of shared equipment."

—George Ferguson, Hewlett-Packard Co.'s business continuity and recovery services

HOT SITE

If a recovery time of a few hours (instead of minutes) is acceptable, a hot site is likely appropriate. The biggest difference between a hosted site and a hot site is the use of shared equipment for infrastructure components like servers and peripherals. Storage is dedicated and real-time data replication is used to get data from the production site to the DR site. Because equipment in the DR site is shared by multiple customers, hot sites are significantly less expensive than hosted sites. "Hot sites and warm sites can be implemented less expensively through outsourcing than doing them in-house because of shared equipment," says Ferguson. "DR services providers rely on the fact that not all customers have a disaster at the same time."

On the downside, the use of shared equipment makes hot sites less flexible because customers are limited by the equipment the DR service provider offers. While some service providers may have a limited selection of equipment, others are more flexible. "About 90% of the time

we're able to use shared equipment, and the rest of the time we work with the customer to make it work," says Marc Langer, president at Recovery Point Systems, a provider of backup, storage and disaster recovery services. Larger service providers may be less flexible, so the nature of the shared equipment is likely to be a determining factor when selecting a hot or warm site provider.

Another consequence of using a site with shared equipment is the time limit on how long customers can use the shared gear in the event of a disaster. The limit varies among service providers, but typically ranges between 30 days and 90 days. "Customers can use the shared equipment for 60 days before they need to get out or before they get migrated to a cold site," says Langer. Service providers with a larger number of data centers, like IBM Corp., can be more flexible. "We're pretty open-ended because we can shift workloads to other data centers," says John Sing, senior consultant, business continuity strategy and planning at IBM's Systems and Technology Group. To avoid unpleasant surprises, a clear understanding of the terms, conditions and limitations of managed DR services is required prior to committing to an agreement that may span several years.

DATA CLASSIFICATION AND DR

DISASTER RECOVERY (DR) for files and folders is generally simpler than disaster recovery for applications because you don't have to consider issues like application consistency, transaction integrity and application dependencies. The challenge with DR for file-based content is mostly a problem of volume and size. Companies may have tens or hundreds of terabytes of file data, so determining what needs to be included in the DR plan can be a daunting task.

Some companies have turned to data classification tools to determine the value of data and its appropriate DR tier. Data may be classified using a variety of tools:

- **STORAGE RESOURCE MANAGEMENT (SRM) TOOLS** typically classify files by meta data such as file type, size and modification date. An example is the Hewlett-Packard (HP) Co. Storage Essentials File System Viewer module, which allows files to be grouped by various file properties.
- **ARCHIVING TOOLS** have built-in classification and tend to go beyond just meta data to include full content indexing. Symantec Corp.'s Enterprise Vault and archiving products from C2C Systems Limited are examples.
- **DATA-LOSS PREVENTION TOOLS** detect and prevent the unauthorized transmission of information and include data categorization capabilities. They're available from McAfee Inc., RSA (The Security Division of EMC Corp.) and Symantec, among others.
- **STANDALONE CLASSIFICATION TOOLS**, available from companies like Abrevity Inc., Kazeon Systems Inc., Njini Inc. and Permabit Technology Corp., can be used to categorize data to determine the appropriate DR tier.

WARM SITE

In contrast to a hot site, a warm site relies on backups for recovery. As a result, it doesn't require dedicated storage but instead can take advantage of less-expensive shared storage. In other words, all components of a warm site, including storage, are shared among multiple customers. Therefore, most of the considerations of hot sites also apply for warm sites.

In the past, there was a huge difference between hot sites and warm sites because backups were limited to tapes. As a result, warm site recoveries were typically measured in days. Warm sites that rely on tape-based backups for recovery are clearly at the lower end of the DR services spectrum.

Disk-based backups have narrowed the gap between warm sites and hot sites, and almost all DR service providers now offer an electronic vaulting option, which is essentially disk-based backup of production data over the network. RTOs and RPOs of warm sites with electronic vaulting are typically less than a day, which is very close to the recovery times offered by hot sites but at a fraction of the cost. "There has been about a 10x price difference between a replicated DR infrastructure and a shared infrastructure with electronic vaulting," explains HP's Ferguson. "Electronic vaulting is closing the gap between tape-based recovery and a replicated DR infrastructure, and customers need to look at it because of its price and reliability benefits."

Disk-based backups have narrowed the gap between warm sites and hot sites, and almost all DR service providers now offer an electronic vaulting option.

COLD SITE

A cold site is rented space with power, cooling and connectivity that's ready to accept equipment. With recovery times of a week or more, a cold site is only an option for business processes that can be down for an extended period. Cold sites are also used to complement hot sites and warm sites in case of disasters that last a long time. "Some of our customers sign up for a cold site as contingency to migrate equipment from the shared infrastructure to the cold site in case a disaster lasts more than six weeks," says Recovery Point Systems' Langer.

It's the customer's responsibility to provide equipment for the cold site during a disaster. A DR plan that relies on a cold site must clearly define the process of procuring and delivering equipment to the cold site when a disaster strikes. It's a risky strategy to rely on purchasing the equipment on the open market when it's needed as it may not be possible to get the equipment in a timely fashion. A better option is to consider subscribing to a quick-ship service available from companies like Agility Recovery Solutions. "You can rent equipment for as little as

$50/month with an option to buy it if needed," says Recovery Point Systems' Langer.

IN-HOUSE DR vs. OUTSOURCED DR

Whether to create a DR site in-house or to outsource it is a fundamental decision that needs to be made when putting a DR strategy in place. The in-house approach may be tempting, with the assumption that the work related to DR can be performed by existing staff. Unfortunately, experience shows that in-house DR is more likely to fail than outsourced DR services.

According to an IDC study, enterprises that didn't outsource lost on average $4 million per disaster incident across a variety of business

WHAT TO ASK WHEN SELECTING A DR FACILITY

1 **WHAT TYPE OF FACILITY SHOULD BE USED?**
- In-house using another office location
- A collocation facility
- Managed collocation space from the likes of Hewlett-Packard (HP) Co., IBM Corp. or SunGard Data Systems Inc.

2 **FOR IN-HOUSE DISASTER RECOVERY (DR) FACILITIES:**
- Is the facility equipped to deal with the increased load during a disaster (bandwidth, power, cooling, etc.)?
- Is designated DR staff available?
- Is equipment designated or at least ensured to be available in case of a disaster?
- Are resources available to periodically test failover?

3 **FOR COLLOCATION FACILITIES:**
- Is the collocation facility a far enough distance from the production site?
- Does the collocation facility have sufficient bandwidth options and power to scale and deal with the increased load during a major disaster?
- Who will manage the equipment in the DR site? If it's managed in-house, many of the considerations of in-house DR apply.

4 **FOR MANAGED COLLOCATION SPACE:**
- Based on recovery time objectives (RTOs) and recovery point objectives (RPOs), determine the type of site required (hosted, hot site, warm site or cold site)
- Ensure that DR testing is included in the proposal.
- As hot sites and warm sites typically limit how long they can be used during a disaster, clearly understand your options in case you need the DR site longer.

functions (e.g., sales/marketing, financing, e-commerce). In contrast, enterprises that outsourced to a third party lost an average of $1.1 million per incident. The study adds that companies that leverage an in-house model spend 32% more than those opting to outsource. It further shows that outsourcers can provide a shorter window of recovery, as measured by RTO over in-house operations by a reduced factor of 0.62. The study concludes that primary and DR data centers are more likely to get out of sync if DR services are performed in-house.

One of the primary reasons why in-house DR scores so poorly is the risk of taking shortcuts and burdening users already overloaded with other work. When a person's primary role is in conflict with their DR role, the primary role usually wins, to the detriment of the DR plan.

CALCULATING THE COST OF DR

Determining the cost of DR is company-specific, and the many variables make it difficult to devise a formula to calculate a DR cost for a given environment. In general, the cost of DR includes the cost for physical space, equipment, power, and network and professional services. But the cost of each of those components can vary greatly. "We have tried to put together a TCO tool, but data centers are too different and our DR options are so customized that it's very difficult to come up with a cost calculator," says David Palermo, vice president of marketing at SunGard Data Systems Inc.

Fujitsu Computer Systems Corp.'s Affordable Business Continuity (ABC) product is one of the few packaged DR kits that includes storage, hosting and bandwidth for a fixed cost of $190,000. The ABC kit includes two Eternus 4000s with 3 TB of raw storage each, replication software and one year of hosting with bandwidth. Fujitsu's professional services works with customers on customized bundles and assists with determining the required server infrastructure (servers aren't included in the bundle).

DR SITE OPTIONS

The prevailing options for DR sites are remote-office locations, collocation space and DR service providers' data centers.

Remote-office location and collocation space: Companies with multiple locations frequently use their remote data centers as DR sites. Leveraging existing facilities and infrastructure is a very cost-efficient DR option. For companies with multiple locations, but not multiple data centers, collocation space offered by providers like Equinix Inc., Savvis Inc. and telcos, may be a good alternative. Collocation facilities are relatively cost effective and usually provide first-class space with sufficient power, bandwidth and high facility standards.

Cost was the primary reason why Matt Blydenburgh, CIO at Tannenbaum Helpern Syracuse & Hirschtritt LLP, New York City, used collocation space in Connecticut for the firm's hot site. Blydenburgh uses Double-Take Software Inc.'s Double-Take to replicate data from the firm's New York City location to its hot site in Connecticut. "We looked at managed disaster recovery services from companies like SunGard, but it was very expensive," says Blydenburgh. "We now pay $1,800 for space and

another $1,600 for bandwidth for both sites."

Managed DR service providers: Managed DR services providers like HP, IBM, Recovery Point Systems and SunGard are dedicated to disaster recovery and are hard to beat in the quality of service they provide. But they're not cheap. To get a fair price comparison between a managed service and using in-house DR facilities, it's essential to take into account all cost components, including the cost of dedicated DR staff.

With 155 DR data centers worldwide, IBM is the largest managed DR firm. Similar to HP, IBM can source all DR components from within IBM. With 30 U.S. and 30 European data centers, and approximately 12,000 customers worldwide, SunGard is also a major player in the managed DR space. Prior to its acquisition of EDS, HP was focused mostly on providing managed DR for companies using HP equipment, but HP is now playing at the same level as IBM. Smaller DR services firms have the advantage of flexibility and are more willing to wheel and deal to win a contract.

Even in financially challenging times, you should never walk away from DR because you can't afford a certain DR tier. Instead, go with a lower, less-expensive tier that gives reasonable protection for the available budget. Not having a DR plan should never be an option. ⊙

Jacob Gsoedl is a freelance writer and a corporate director for business systems. He can be reached at jgsoedl@yahoo.com.

WITH GREAT POWER COMES GREAT RESPONSIBILITY.

MAXELL OFFERS THE PERFECT COMBINATION OF STORAGE & SUPPORT TO KEEP YOUR DATA SAFE.

On your sturdy shoulders rests the data of your entire company, but there's no need to worry. Maxell's years of experience creating superior backup tapes ensures that your data will be there when you need it. With Maxell's customized training sessions, free barcode labeling, computer tape testing and 24/7/365 technical service, you can relax knowing that retrieval will be a breeze. Of course, you don't have to let your office know that. **Let them treat you like the hero you really are.**

maxell
www.maxell.com

Consolidated archiving

By Stephen Foskett

There are two main approaches to archiving data:

1 ***independent architectures for individual applications; or***

2 ***one architecture that consolidates all archives on a single platform.***

We examine the pros and cons of each approach.

DATA ARCHIVING IS the vampire of the storage world. It promises to rejuvenate enterprise storage systems by sucking out debris so they can work as well as they did when they were young and not burdened by millions of files. But behind the benefits of archiving, there lurks a hidden detail: Getting archiving to work and the final cost of implementation can be terrifying.

But times are changing, and the benefits of archiving are becoming even more enticing. Once focused solely on the data lifecycle (moving content from expensive disk storage to cheap tape), archiving has transformed into something altogether different. Due to compliance and legal reasons, today's archives are increasingly serving as long-term data storehouses, and many implementations forego the old stub-and-delete hierarchical storage management concept altogether.

Archiving has become an indispensable tool to protect an organization, not just a few pieces of data here and there. Although most archiving efforts start with a single application, demand quickly grows to include multiple data types and diverse systems. The key question is whether

you should attempt to expand a single archiving system to include heterogeneous data or employ multiple, single-application archiving systems. But there are other ways to build a consolidated archive. You can leverage a single storage platform for multiple archiving systems or employ enterprise search technologies to put a unified face on a diverse set of systems.

EMAIL OFTEN STARTS THE PROCESS

Most archives start small, and many begin with one type of data like email (see "What to archive: Different data types," this page). IT decides it needs to control growth, so they go looking for a system that can stub out attachments or move data out of the Exchange server. Then the legal department demands a complete set of email messages for a litigation-related search. Later, records management needs to retain certain messages for compliance with Sarbanes-Oxley or other industry regulations.

WHAT TO ARCHIVE: DIFFERENT DATA TYPES

ALTHOUGH MOST business information is digital today, not all systems manage information equally. Manageability requires organization and structure, the ability to search for information and meta data to categorize content. We use these elements to classify data as structured, semi-structured or unstructured.

Structured applications are inherently organized, although the identification, description and relation between data can be highly customized. In the enterprise, structured databases are often core applications with specialized administrators managing the data and archiving.

Systems like email have some structure, but they weren't developed with information management in mind, and it shows. What structure they have is functional, designed to serve specific application needs rather than the higher goal of manageability.

Finally, there's the class of unstructured data that's so familiar in file systems. Although some basic systems are used to organize and describe these files, they can't be called truly structured as they lack information about their functional or organizational relationships.

	Structured	*Semi-structured*	*Unstructured*
Example	Relational database	Email, medical images, document repositories	File servers, multimedia files
Organization	Highly organized with intrinsic interrelation	Moderately organized by meta data	Basic location and name "tree" organization
Meta data	Complex custom attributes	Some intrinsic attributes	Basic attributes and rarely used extended abilities
Search	Built around query language	Keyword search and categorization by meta data	Third-party tools for search

Regardless of their original intent, the scope of these point solutions tends to expand over time. Email systems also include contact lists, calendar entries, to-do lists and notes, and these may not have been considered at first. What about attachments? Most corporate record-retention policies call for documents to be retained, but they might be duplicated on the file server or document management system. And once legal gets used to simpler search and retrieval of old email messages, they'll want to archive and search across a variety of data types beyond email, including document management systems, file servers and structured data systems.

Events like those often lead to a key turning point in the process of implementing consolidated archiving: Should the current system be expanded to include other record types or should another vertical solution be deployed for each new application?

Standardization and simplicity are often the primary reasons for expanding an existing archival system. Jason Beckham, director of IT at Payformance Corp. in Jacksonville, FL, sees the benefit in sticking to a single platform. "Our current Hitachi archiving platform is already in place and it's a known quantity," he says. "We plan to expand on the existing HCAP [Hitachi Content Archive Platform] when we add email archiving, since it's so simple to implement and will require much less management and training."

Standardization and simplicity are often the primary reasons for expanding an existing archival system.

THREE KEYS TO ARCHIVING

There are many ideas about what a consolidated or unified archive should look like, and preconceptions can clash when you're considering creating such a solution. There are three key elements to any archiving system:

- **Archiving software**, from companies like Autonomy Zantaz, EMC Corp., Mimosa Systems Inc. and Symantec Corp., manages the location, movement and disposition of data.
- **Storage hardware**, from companies such as EMC, Hewlett-Packard (HP) Co., Hitachi and NetApp, receives the data to be preserved and specialized platforms handle encryption, protection, retrieval and destruction of data.
- **Management software,** from vendors like Abrevity Inc., Attenex Corp., Autonomy Corp., Clearwell Systems Inc., i365 (a Seagate Company) and Kazeon Systems Inc., provides services like search, classification and e-discovery capability.

But most applications don't fall into these neat classifications. As they develop their products and the market matures, vendors continually add features like e-discovery support, search and data movement, blurring the lines among storage, archiving and management.

The variety of elements and overlapping features add complexity to the once-simple world of archiving.

Navigating the archiving market starts with an evaluation of your company's objectives for its archive. If the objective is to serve business demands like in-house e-discovery or retention to comply with regulations, it makes sense to let data management features drive product selection. But if IT needs a system to control data growth or enable lifecycle management, higher-level search and e-discovery features are less relevant. Regardless of the initial object, it's likely the archive will eventually serve both business and IT demands.

Not all archives will use all three archiving elements presented here. Some organizations send data to an archiving platform directly from a custom application, while others will use conventional storage systems rather than investing in a specialized device as their archive target.

Brian Greenberg, an independent IT strategy consultant based in Chicago, suggests a strategy is needed before expanding the archive environment. "Larger organizations, especially in regulated industries, are looking for federated search and management across data types, but smaller, less-regulated companies might be able to keep their data in silos," he says. "The key is the level of overarching management needed."

One practical solution to bring order to a diverse set of data archives is to leverage federated management software.

ARCHIVE FEDERATION

One practical solution to bring order to a diverse set of data archives is to leverage federated management software. These applications let you pick the best point solutions for email, databases, content management systems and file servers, but hide the complexity of having multiple archivers behind a unified interface. These are especially appropriate where search, rather than capacity management, is the primary reason for archiving.

Although many archiving platforms today include search and e-discovery features, none can match legal-focused data management tools. Simple Boolean text search can't hold a candle to the concept-clustering and fuzzy-search features in products like Attenex's Patterns, Clearwell's E-Discovery Platform or i365's MetaLincs. Tools like these also boast complex e-discovery features, including review and annotation, that are beyond anything found in the more IT-focused archiving applications.

Some of these tools can be used to search non-archived data as well. Autonomy and OpenText Corp. offer enterprise search platforms that can manage both production and archived data from a single interface and can be integrated into other enterprise applications for complex environments. CommVault even includes backup and remote site replication in its unified archiving and search platform, creating a one-stop data protection and management suite.

PLATFORM GAME

Many companies are also attempting to solve the archive consolidation puzzle from the bottom up, investigating unified storage platforms to retain data. This approach is most appropriate where capacity control, rather than search, is the key requirement, as storage unification can bring many advantages.

POPULAR ARCHIVING PLATFORMS

	Architecture	*Protocols*	*WORM*	*Data reduction*	*Search*
Caringo Inc. CAStor	Software for clustered nodes	HTTP, NFS, CIFS	File level, with hash	None	None
Copan Systems Inc. Copan 300A File Archiver	MAID integrated storage	Proprietary API, NFS, CIFS	File level	Dedupe (coming)	None
EMC Corp. Centera	Clustered nodes with storage	Proprietary API, XAM, NFS, CIFS, FTP, HTTP	Object level, with hash	Single-instance storage	None
Hewlett-Packard Co. Integrated Archive Platform	Clustered nodes with storage and integrated archiving software	Proprietary API, XAM	Object level, with hash	Single-instance storage	Full text
Hitachi Data Systems Content Archive Platform 300	Clustered nodes with storage	NFS, CIFS, HTTP, HTTPS, WebDAV, SMTP and NDMP	Object level, with hash	Single-instance storage, compression	Full text
Hitachi Data Systems Content Archive Platform 500	Clustered nodes using SAN storage	NFS, CIFS, HTTP, HTTPS, WebDAV, SMTP and NDMP	Object level, with hash	Single-instance storage, compression	Full text
IBM Corp. System Storage DR550	Integrated storage	Proprietary API, NFS, CIFS	File level	None	None
NetApp NearStore on FAS	Archiving features on modular storage	NFS, CIFS	File or volume level	Dedupe	None
Permabit Technology Corp. Enterprise Archive	RAIN clustered nodes with storage	NFS, CIFS, WebDAV	File level	Dedupe, compression	None
ProStor Systems Inc. InfiniVault	Integrated removable storage	NFS, CIFS	File or disk level, with hash	Single-instance storage	Keyword

Consolidated archiving doesn't have to be a great technical challenge. Although specialized protocols like EMC's Centera API and the Storage Networking Industry Association's new eXtensible Access Method (XAM) specification were developed specifically with archiving in mind, Thomas Savage, senior manager, product marketing data retention at NetApp, points out that "most archiving applications support a variety of storage devices over standard CIFS or NFS." Nearly any storage device can serve as a landing spot for archive data.

But specialized archive platforms like those from EMC, Hitachi, HP and NetApp bring special capabilities for storage of archival data (see "Popular archiving platforms," p. 22). Most offer native support for "objects" rather than files, and can manage these with custom meta data. Some can autonomously enforce retention policies, even securely deleting data once it has "expired." They may also support full-text indexing and search. Although higher-level archiving software will almost certainly duplicate some of the functionality of these devices, their presence at this lowest common layer can make it simpler to configure these features and consistently enforce policy.

Data protection is especially critical, so IT managers might feel more comfortable with a "belts and suspenders" approach. Even if the archiving software places some data off-limits by policy, a basic storage system used as the archive target may leave it unprotected. Although careful application of traditional security and access controls found on standard NAS systems can offset this risk, the enhanced features of specialized archiving storage systems go further.

Data protection is especially critical, so IT managers might feel more comfortable with a "belts and suspenders" approach.

Storage platforms specifically designed for archiving often include enhanced protection against modification or deletion, sometimes called write once, read many (WORM) storage. Contrary to popular belief, there are few regulations or laws specifically calling for WORM storage. But the concept of access control is as central to long-term data management as retention schedules and classification, and legal discovery regularly demands certification that data hasn't been accessed or modified. Therefore, no legal or compliance archive should be without WORM capability. Systems may also offer authentication that uses mathematical checksums to verify that data hasn't been modified, but use of those haven't yet been commonly presented in legal cases. Finally, make sure the system logs and reports all access attempts, as this data is critical for documenting compliance.

Many storage systems designed for archiving are now adding data-reduction technologies ranging from compression to single-instance storage (SIS) to advanced deduplication. These various technologies function similarly, using algorithms to reduce the amount of data that must be stored in a lossless fashion. Single-instance storage compares

whole files or objects with existing content, storing only a single copy of duplicates. Traditional data compression encodes files or objects, creating a "dictionary" of repeating patterns to shrink them. Finally, deduplication technology searches for duplicate data both within objects and across the entire data store, a complex computing task that can result in vastly reduced storage requirements. Each approach balances storage reduction against the computing power required to accomplish it.

These specialized archiving systems also include standard enterprise storage capabilities like scalability, high availability and data replication. Architecture varies from product to product, with some employing traditional storage array technology and others based on a cluster of redundant nodes. One differentiator is the extent to which the archive

CONSOLIDATED vs. UNIFIED STORAGE

AS ARCHIVES ARE IMPLEMENTED, it's tempting to create "stovepipes" of storage on the back end, with each application using its own storage system for content and indexes. "Of course a unified storage platform gives efficiency of resources and management, but consolidating archives has other benefits," says Rob Mossi, senior marketing manager for archiving at Hitachi Data Systems. "Creating a consolidated platform for the storage of archived content is a winning strategy, especially when leveraging advanced storage system features like duplicate elimination, compression and replication as found in a highly scalable, performance-enabled active archive solutions," he notes.

Although archive software can use a variety of storage platforms to store the archived content itself, there are other storage requirements for these applications. All archiving software products maintain an index of both the production and archived data, and that database is often stored on a conventional storage array. A unified, multiprotocol storage system with conventional block storage and archiving features allows both the index and content to share space, easing management and growth headaches.

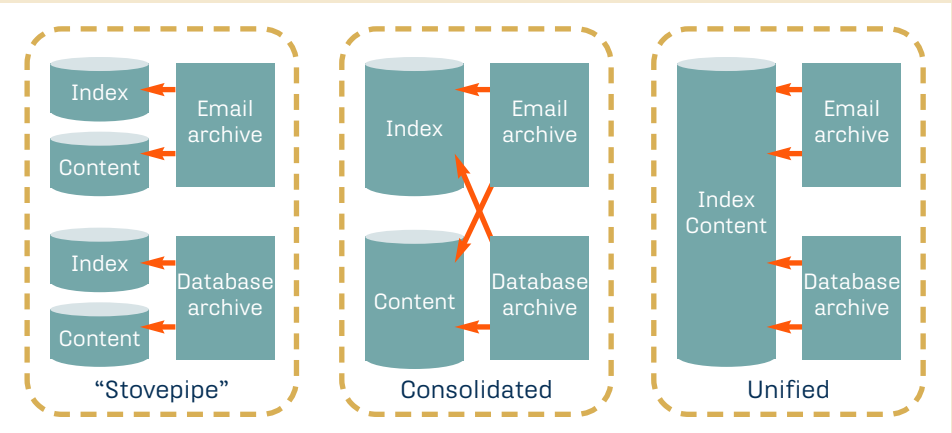

Source: Contoural Inc.

system can, or should, include non-archive data (see "Consolidated vs. unified storage," p. 24).

Financial considerations also come into play when comparing an archiving system to plain storage capacity. Payformance's Beckham recognized the cost differential, but says it was justified based on the added meta data and WORM functionality. "We weren't buying a storage system, we were investing in the value that an intelligent archiving platform brought to our business," he says. When evaluating a storage platform for archive consolidation, consider whether the advanced features of an archiving system are required. Although an existing storage device might be acceptable for retention, these capabilities might be worth the extra money. Note that technical issues like the scalability of deduplication and the manner in which protocols like NFS handle offline files sometimes crop up. In these cases, only a specialized archiving platform will do.

Finally, since archiving software often directly integrates with these capabilities, these specialized storage systems are much more likely to be configured correctly than basic storage devices. When purchasing archiving hardware or software, look for highly integrated combinations to minimize the risk and headache of management.

THE GLOBAL ARCHIVE

Where should one start when considering a consolidated archive? Like all IT decisions, the first consideration should be the business objectives to be served. Start by thinking about the goal: Will the archive deliver capacity control, compliance, business productivity or legal support, or a combination of these capabilities? Then compile a list of specific functional requirements: what the system must do, rather than how it must accomplish these things. Only then can products be evaluated fairly in this complicated corner of the storage market. ⊙

Stephen Foskett is director of data practice at Contoural Inc.

STOP PAYING

STORAGE TAXES

Save Over 90% on Storage Expenditures.

Disk capacities double every 18 months and the cost per gigabyte decreases 50% per year, but it's still not good enough. Ever wonder why? It's called storage tax.

For every terabyte of primary storage installed, you pay not only for the raw capacity, but also for replication, backup, maintenance, data migrations, and the personnel to make it all happen. Before you know it, you're paying 3,4 or even 5x the figure listed on your quote.

PowerFile eliminates these storage taxes and delivers savings of over 90% per terabyte on capital and operating expenditures within 12 months.

PowerFile®
storage evolved.

FIND OUT HOW...
DOWNLOAD A FREE WHITEPAPER AT **WWW.POWERFILE.COM/NO-TAX**

Faster disk drive rebuilds

RAID might not be the best choice for high-capacity drives. It's time to rethink your rebuild strategies.

HERE'S A LOT OF TALK about shortening rebuild times for large capacity disk drives in today's storage environments. Fast rebuild technology is widely deployed nowadays, but plenty of users still don't think in terms of hardware RAID and individual drive rebuild times. And here's a new angle on the discussion: perhaps the best way to shorten rebuild times is to not have to rebuild in the first place.

Roughly 50% of failed SATA drives returned to vendors result in a diagnosis of "no trouble found" and are returned to service as replacement drives that typically function like new. That's because SATA drives were originally designed for lightly loaded desktops and laptops, rather than high-performance enterprise arrays, and they occasionally experience slowdowns in performance that result in a disk being diagnosed as non-responsive. As a result, several vendors have introduced technology to diagnose these issues and determine if the disk is actually failing or if it's just an intermittent slow down. This technology is important to understand because it reduces the risk of data loss due to the potential of a second drive failing during rebuild.

Despite the standardization efforts of the now defunct RAID Advisory Board, most vendors developed protection schemes that met basic RAID definitions but varied widely in their implementations.

Before deciding on which approach or vendor best addresses your RAID rebuild challenges, let's look at how we got here. The term RAID, or redundant array of independent (or inexpensive) disks, was introduced in the late 1980s to describe a method of protecting disk drives in an array. Despite the standardization efforts of the now defunct RAID Advisory Board, most vendors developed protection schemes that met basic RAID definitions but varied widely in their implementations. No matter the strategy, RAID rebuild times across the board get longer as disk drive size increases. That's because there's more data to copy or rebuild from parity. In the event of a single disk drive failure in most RAID modes, data is left unprotected until the RAID rebuild is finished, and

RAID rebuilds suck up significant processing power.

However, there are ways to keep data protected in the event of a single disk failure. Users can implement dual-parity RAID 6—which keeps data available in the event of a dual drive failure in a single RAID group—or go as far as implementing remote mirroring technology to protect themselves not just from a drive failure but to keep data available in the event of a full site failure. But there are costs associated with each layer of protection added, and these need to be balanced against the value of the data to be protected; the overhead required to allocate capacity for data protection is, in some cases, three to four times the amount of data stored.

With the advent of high capacity TB-size Serial ATA (SATA) disk drives, the problem is compounded. SATA drives spin at less than half the speed of Fibre Channel (FC) drives, but hold up to 1 TB (twice the capacity of FC drives). The density of the drive doesn't make up for the slower rotation speed, however; average latency for a 7,200 rpm disk drive is more than two times the latency of a 15,000 rpm drive. With TB SATA drives, rebuilds could extend for multiple days, depending on how busy the system is, and become onerous to the point that they have an unacceptable impact on the business. There are significant cost advantages to storing data on large capacity drives: the price per MB is much less than high-performance FC drives and, thanks to their price advantage, SATA drives have been widely deployed in archive systems and scale-out storage architectures while higher performance FC drives have continued to hold court at the top storage tiers.

NEW DATA PROTECTION SCHEMES

Storage vendors are finally beginning to understand that it's not about protecting disks but protecting information, and their data protection schemes are evolving to reflect this. There are some novel approaches in the market to solving the problems produced by large, slow drives. Some technologies reduce the overall number of rebuilds a system performs. Some have shifted to information-based protection schemes in which, rather than mirroring a disk, they mirror information (files, chunks or objects). Some even do a little of each. So how does this impact rebuild times? When you think in terms of rebuilding information rather than a single disk, you can put the power of the system architecture to work, leveraging the massive parallelism opportunity presented by multidisk architectures.

There are several technologies in the market today that reduce the overall number of drive failures, and thus the number of rebuilds required. In some instances, vendors take unresponsive drives offline to diagnose problems and return them to service if no trouble is found. This is a great approach, as it eliminates the need to perform a full rebuild. When the drive goes offline, the system journals all writes that would have gone to that drive while attempting to recover the drive. After a successful recovery, only the data in the journal is required to be rebuilt, not the entire disk.

Some vendors have a two-pronged approach that reduces the overall number of rebuilds required and speeds rebuild time leveraging grid storage architectures. One approach kicks in when a drive doesn't respond immediately to an access request. The system responds by doing

a mini parity rebuild of the requested data and returning the rebuilt data while taking the non-responsive drive temporarily out of service. This drive then undergoes a brief diagnosis and is returned to service, thereby eliminating the need for a rebuild. Any data written while the drive is offline is written to other available space in the system.

This also speeds rebuilds by putting its grid architecture to work. Most grid-based architectures have capacity or storage nodes and separate processor nodes. Typically, all processor nodes can access all capacity nodes. When data is written, it's broken into a number of fragments. These fragments are then distributed across as many storage nodes as are in the system. Using a default of nine data fragments and three parity fragments (the exact number of parity fragments is user configurable), each of 12 storage nodes would get a fragment. If there are four storage nodes (the minimum configuration), each node gets three fragments. In the event of a drive failure, the data from that drive is rebuilt, just like in conventional hardware RAID. But unlike conventional RAID, data isn't rebuilt to a single drive; the data is redistributed across the storage nodes leveraging any available storage capacity. If an entire storage node fails, the data from those drives is rebuilt across the remaining storage nodes. We've seen this type of technology implemented for both parity-protected data and mirrored data. Thanks to protecting data rather than disk drives, as well as the power of a grid architecture, rebuilds happen in a fraction of the time it would take for a conventional drive rebuild. It's the information that's being rebuilt, not the exact drive layout.

Other vendors seek to leverage their architectures to speed rebuild time and reduce the risk of data loss if multiple drives fail. When a file is written, the data and parity is distributed across the available disk drives in the cluster. In the event of a drive failure, the data required for a rebuild is spread across multiple nodes in the cluster, so drives across the entire cluster are leveraged.

Shifting data protection strategies from a hardware-based approach to a software-based approach creates new possibilities. With a hardware-based protection scheme, the choice is often between protecting all of the data or none. Information-based protection opens the door to the possibility of more granular, policy-based information protection.

The bottom line is that different storage characteristics are required for various data types. Hardware RAID schemes continue to be a good solution for lower capacity, faster drives and won't go away any time soon. But it wouldn't be surprising to see information-based data protection schemes become more mainstream in tier 1 storage products over time, as vendors continue to simplify administration and build information-centric systems.

There are plenty of vendors offering information-based data protection schemes or rapid rebuild technology. Even in a tough economy, the number of vendors offering technology that accelerates or reduces the need for rebuilds seems to be growing. Remember that when you're evaluating technology that leverages high-capacity commodity disk drives, you should ask your vendor what they're doing to reduce your exposure to data loss during rebuilds.

Terri McClure is a storage analyst with the Enterprise Strategy Group, Milford, MA.

30-second recovery for your Exchange, SQL Server and Windows file system data

Get a free BakBone stopwatch just for attending a live demo!

See for yourself

Attend one of our weekly 30-minute LIVE online demonstrations and witness

- *Near-instantaneous Exchange recovery*
- *Exchange brick-level recovery without brick-level backups*
- *Fast and flexible any-point-in-time recovery for Windows file system data*
- *Easy set up for SQL Server protection*

Register today!

www.bakbone.com/30-second-restore

NetVault: Real-Time Data Protector™ gives you near-instantaneous recovery

In as little as 30 seconds, your users are able to regain access to their Exchange, SQL Server and Windows file system data, significantly reducing downtime and its impact on your business.

Near-zero data loss

Don't miss a thing! NetVault: Real-Time Data Protector is *always on* – capturing continuous history for granular any-point-in-time recovery of your most important data.

Maximize revenues and cut costs to survive (and thrive) in tough times

Downtime and data loss generate high recovery costs and strip away revenue. By significantly reducing downtime and ensuring near-zero data loss, NetVault: Real-Time Data Protector is a must have solution for those who need to protect their revenue streams while reducing capital expenditures.

www.bakbone.com

BakBone®
The Power of Simplicity™

Green is still the color of money in storage

Vendors are quick to tout their products as environmentally friendly these days. But for respondents to a recent *Storage* survey, "green" is more likely to mean the color of money. Only one-third of respondents say their company has commitments from management to address green storage issues. When asked which technologies they've implemented, 67% of those surveyed have virtualized servers, 43% have virtualized their storage and 34% use compression—all considered green technologies. (Note: Respondents could choose more than one option.) However, the top reasons noted for implementing these technologies are saving money on equipment power-consumption costs and saving money on equipment cooling costs. While these are more economic than ecologic reasons, they still yield green results.

—Christine Cignoli

Does your company have initiatives or commitments in place from corporate management to address green storage issues?

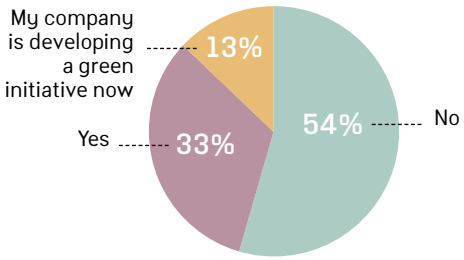

Which of these technologies have you implemented?

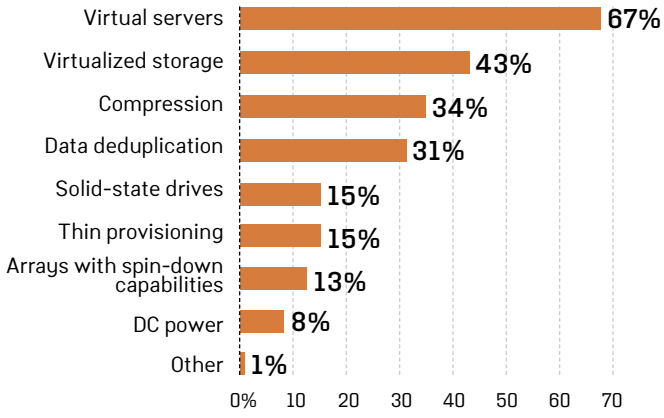

*Respondents could choose more than one option.

56%

Respondents who would pay "a little more" for a green product

What are the reasons you've implemented these technologies?

63%	To save money on equipment power-consumption costs
46%	To save money on equipment cooling costs
40%	We needed to upgrade systems for more capacity
31%	We needed to upgrade systems for performance
23%	To lessen environmental impact

*Respondents could choose more than one option.

"Revolutionary things haven't happened yet or, if they have (in the case of SSDs), they're out of reach of most businesses. The ROI simply isn't adding up yet."

—Survey respondent

Check out the following resources from our sponsors:

BakBone Software Inc., page 30

BakBone Software

Using Business Continuity to Ensure Uninterrupted Operations

Data Domain Inc., page 7

Data Domain Data Deduplication Center

Deduplication School

Maxell Corporation of America Inc., page 17

With Great Power Comes Great Responsibility

NetApp, page 4

PowerFile Corp., page 26

Storage Tax Whitepaper

Understanding and Avoiding the Storage Tax