

# STORAGE

Vol. 7 No. 13 February 2009

PRODUCTS OF THE YEAR

Innovation was the byword for storage vendors in 2008. See which products came out on top.

# The Best!

# 08

## ALSO INSIDE

- 4 Four ways to cut costs
- 9 Dedupe isn't just for backups
- 12 Products of the Year
- 30 10 GbE: Ready for storage?
- 35 Multitiered DR
- 39 Scale-out NAS
- 42 Storage in a down economy

## Survival Guide for Storage Managers

- 4 **EDITORIAL** Extreme circumstances call for extreme measures. We list four ways to trim or hold down storage costs this year. by RICH CASTAGNA

## Deduping: The Rest of the Story

- 9 **STORAGE BIN 2.0** Deduplication is great for paring backup data, so it should be even better for primary data. But where does it make the most sense to start dedupe? by STEVE DUPLESSIE

## 2008 Products of the Year

- 12 Traditional and not-so-traditional products and services earned honors this year, addressing such issues as data deduplication, solid-state storage and storage cloud services. by ANDREW BURTON, RICH CASTAGNA, TODD ERICKSON, BETH PARISEAU, DAVE RAFFO AND CAROL SLIWA

## Timetable for 10 Gigabit Ethernet

- 30 The next generation of Ethernet is likely to have a profound effect on storage; pumped-up iSCSI performance may challenge Fibre Channel's tier 1 dominance. by CHRISTINE CIGNOLI

## The Under-Over on DR

- 35 One size doesn't fit all, especially when it comes to disaster recovery planning. Learn how to build a multitiered DR services capability. by JIM DAMOULAKIS

## Scale-Out NAS Poised for Growth

- 39 **HOT SPOTS** The compelling economic benefits of deploying scale-out NAS have the technology increasing its footprint in the general storage space. by TERRI MCCLURE

## Storage Shops Holding Up Under Economic Stress

- 42 **SNAPSHOT** Whether budgets are up or down, most storage managers are doing some belt tightening, ready to forego some features or performance in favor of lower price tags. by RICH CASTAGNA

## Vendor Resources

- 43 Useful links from our advertisers.






**You didn't just lose backup tapes.  
Your customers' trust is gone, too.**

**Protect your high-value data and your reputation  
with the Fujifilm Tape Tracker™ System.**

Data loss adds up to millions of dollars a year. But there's no reason you have to become another statistic. Not when you have Fujifilm's Tape Tracker System on your side.

With Tape Tracker, powered by LoJack InTransit™ Technology, you'll know precisely where your data tapes are—24/7. Utilizing our Fujifind security application, geo-fencing and monitoring center, you get true visibility and security for backup tapes when they're traveling outside your data center.

Fujifilm's Tape Tracker System. Stop costly data loss—before it happens.



**Download the White Paper—***AGPS Surveillance of Data In-Transit*,  
or get more information on our website at [www.tapepower.com/tapetracker](http://www.tapepower.com/tapetracker)

© 2009 FUJIFILM Recording Media U.S.A., Inc.

**FUJIFILM**  
No one backs you up like we do.



# Survival guide for storage managers

*Extreme circumstances call for extreme actions. Here are four ways to trim or hold down storage costs this year.*

**M** **EWS FLASH:** The economic situation is grim right now and the short-term outlook isn't much better. Money's tight—if it's there at all—and the chance that you'll be able to get help for your beleaguered staff is less than nil. It would be nice to just hunker down and do your best to weather the storm, but for most shops there's still too much happening to just assume a defensive posture and hope for the best.

Lousy economy or not, your disk capacity will grow, and now isn't the time to put critical things like disaster recovery (DR) planning on hold. While your business pulls itself up by its proverbial bootstraps this year, there will likely be new applications and more demands on storage systems. With the business world turned on its ear, new thinking is required because "same old" ain't gonna cut it.

Here's a basic survival kit for storage managers, with four suggestions for cutting costs while not sacrificing all that much.

- 1. Deep-six support.** Support contracts can cost as much as or more than the equipment they're supposed to protect, and the older the gear the higher cost. You can save a bundle by cancelling support contracts for older, less-critical storage systems. The risk-adverse among you probably think I'm crazy, but I didn't make it up—I've heard from many storage managers over the last couple of years who have eschewed support to give their budgets a break. They stock up on the parts that are likely to fail (based on years of experience with the systems) and take a DIY approach to equipment maintenance. Some of them have even said that their uptime improved because they could fix the problem themselves in less time than it took to wait for a service rep to show up.

- 2. Think thin.** If thin provisioning is an option for your arrays, buy it. If it's already there but you're not using it, turn it on. Thin provisioning isn't a panacea for storage growth because sooner or later you're going to have to add capacity. But why not make it later? Later, when the economy

We're so far from business as usual right now that even a far-out idea or two might be worth a try.

isn't so pathetically anemic. Industry statistics are all over the place, but it's a good bet that you're not using even half the storage you have installed. With thin provisioning's legerdemain, you're likely to improve usage and put off having to buy more disks.

**3. Jettison junk.** It's a lot easier to talk about getting rid of old, useless data than it is to actually dump the data. But it's also about the easiest way to reclaim capacity without having to shell out too much cash. There are a number of ways to seek and destroy all the junk clogging your arrays. Some, like data classification, might require purchasing software or an appliance to help you sort the junk from the jewels. But there are other ways, like expanding your backup app's exception list so you're not making more copies of useless files and compounding your capacity problems. Or you can take this approach with employees: "Everything will be deleted unless you tell us not to." That's effective sometimes, but be careful—it can lead to what backup guru W. Curtis Preston calls an RPE (resume producing event) or, at the least, unhappy campers on your staff.

**4. Shed some tiers.** If you've already tiered your storage, you know how effective it is for controlling expenses and minimizing management. But even if your tiered storage system has been working fine, it might be a good time to rejigger the tiers. The goal is to get even more data off expensive tier 1 storage. The criteria you developed to determine what data goes where might have worked in healthier economic times but, let's face it, those standards might not hold up as well during the biggest bust in 70 years. Make the criteria even tougher so that the justification for parking an app's data on tier 1 has to be totally bullet-proof. You'll be surprised at how much data you'll be able to move to cheaper storage just by changing the names of your service levels. For example, instead of a "gold" service level, call it "the most expensive storage we have and you'll have to convince the CEO if you want to use it."

I'm sure a lot of you have thought of trying at least one or two of these ideas. We're so far from business as usual right now that even a far-out idea or two might be worth a try. Drop me a line and let me know how you're coping. ☺

---

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

*\* [Click here for a sneak peek at what's coming up in the March issue](#)*



# De-boxed in

When it comes to data de-duplication, most companies only offer one kind of solution. But with Quantum, you're in control. Our new DXi7500 offers policy-based de-duplication to let you choose the right de-duplication method for each of your backup jobs. We provide data de-duplication that scales from small sites to the enterprise, all based on a common technology so they can be linked by replication. And our de-duplication solutions integrate easily with tape and encryption to give you everything you need for secure backup and retention. It's this dedication to our customers' range of needs that makes us the smart choice for short-term and long-term data protection. After all, it's your data, and you should get to choose how you protect it.

**Find out what Quantum can do for you.**  
For more information please go to [www.quantum.com](http://www.quantum.com)

**Quantum**  
BACKUP. RECOVERY. ARCHIVE.

# STORAGE COMING IN MARCH

## The Best Midrange Arrays Revealed

In our fourth annual Quality Awards for Midrange Arrays, *Storage* readers rank the leading vendors and their midrange storage systems for service and reliability. Readers evaluate their products in five categories: sales-force competence, initial product quality, product features, product reliability and technical support. Read next month's issue to see the surprising results.

## Removing Tape from Remote-Office Backup

Getting remote offices to properly back up their data can be a challenge, especially if you rely on tape backup systems at each remote location and have limited local tech expertise. Disk-based technologies, along with online services, can make remote-office backup more predictable and reliable.

## Dedupe Redux

2008 was a big year for data deduplication products, but this year, you can expect the pace of new product introductions to continue unabated. Key vendors, including some storage heavyweights, will roll out new or enhanced dedupe products this year. Check out this overview to see who's doing what on the dedupe front in 2009.

**And don't miss our monthly columns and commentary, or our Snapshot reader survey.**

### TechTarget Storage Media Group

#### **STORAGE**

**Vice President of Editorial**  
Mark Schlack

**Editorial Director**  
Rich Castagna

**Senior Managing Editor**  
Kim Hefner

**Senior Editor**  
Ellen O'Brien

**Associate Editor**  
Christine Cignoli

**Creative Director**  
Maureen Joyce

**Contributing Editors**  
James Damoulakis,  
Steve Duplessie, Stephen  
Foskett, Jacob Gsoedl



**Site Editor**  
Ellen O'Brien

**Senior News Director**  
Dave Raffo

**Senior News Writer**  
Beth Pariseau

**Features Writer**  
Carol Sliwa

**Senior Managing Editor**  
Kim Hefner

**Associate Site Editor**  
Christine Cignoli



**Site Editor**  
Andrew Burton

**Associate Site Editor**  
Heather Darcy

**Assistant Site Editor**  
Chris Griffin

**Features Writer**  
Todd Erickson

**Executive Editor and  
Independent Backup Expert**  
W. Curtis Preston

#### **Storage Decisions**

**TechTarget Conferences**

**Sr. Editorial Events Manager**  
Lindsay Mullen

**Editorial Events Associate**  
Nicole Tierney

**Storage magazine**  
117 Kendrick Street,  
Suite 800  
Needham, MA 02494  
[editor@storagemagazine.com](mailto:editor@storagemagazine.com)

Subscriptions:  
[www.SearchStorage.com](http://www.SearchStorage.com)

# 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](http://www.bakbone.com/30-second-restore)

[www.bakbone.com](http://www.bakbone.com)

## 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.

 **BakBone®**  
The Power of Simplicity™





## Deduping: The rest of the story

*Where does it make the most sense to start dedupe?*

**IN MY LAST COLUMN**, I explained the biology around how and why we end up with so many data replicas, as well as why deduping in the backup process is such a great idea. So when do we apply this concept up the food chain? If it's good in backup, it should be great in the primary infrastructure. But different types of data are created in primary: records, files, objects, blogs and so on. Data lives in the primary infrastructure, but it goes through different lifecycle stages. So where and when does it make the most sense to kill data copies? Read on, my friends.

**STAGE 1.** All data—Word docs, PowerPoint presentations, trading data, arbitrage, video, MP3s—is born dynamic or transactional. Everything is dynamic for some time. At this early stage, data tends to matter the most and has the highest degree of protection. If we lose data—whether it's a document being written on a laptop or big money transactions being processed on a massive system—the biggest impact is here. This is where we normally make our first replica; at a minimum, we probably mirror data at this point.

**STAGE 2.** According to the universal data lifecycle, which I'm perpetuating because it's correct, simple and obvious, all data becomes fixed or persistent after some time. Not at the same time, but eventually. It's subjective, not objective (not truly, but it makes people feel better if I say that). At some point, data stops changing and simply "is."

The second stage is what we term "persistent active data," which is data that no longer changes, but is still very active. That doesn't mean access to that data is automatically less important; usually, it's more important at this stage. But this is where we tend to make the most primary copies of data. We replicate for disaster recovery by making backup copies and snapshots. We replicate to test/development systems and data warehouses. We email copies to our suppliers, partners and cousin Chuck. Then we back up the copies of the copies and make more copies. We need to provide these copies for as long as disparate systems and applications require them. We probably don't need to keep backing up all 87 copies but, to repeat the only phrase my 16-year-old daughter ever utters to me, "Whatever."

**STAGE 3.** The third stage of life is when data enters the “persistent inactive” state: non-changing data that’s rarely accessed. This is where 90% of all commercial data sits in its lifecycle and, thus, is where 90% of the capital and operational gains can be made from both process and technology. Why would anyone back up this data? It never changes, and you’ve already backed up copies of copies of it. It’s the same with disaster recovery. At this stage, you want to be thinking about treating this data much differently than in previous stages. It needs to be on cheap, write-once, read-seldom-if-at-all, power- and cooling-efficient gear that preferably a monkey can manage. This is the stage where we want to massively reduce the copies of data we have. It’s still primary storage, but by applying dedupe here we can probably chop at least 50% or more of our overall capacity off at the knees. If you couple that with some common sense backup/disaster recovery policy changes, you might get a free weekend or two.

**STAGE 4.** The fourth stage is the “Who cares? I’m quitting if we ever actually need to go to this stage to recover” stage. It’s the offsite deep archive or doomsday play. You have to do it, but you don’t have to do it with 9,756 copies of the same non-changing data, do you? Three or four copies seems OK to me.

The inevitable next step is to figure out how to slide the dedupe lever closer to the point of creation, and the biggest value point will be at Stage 3. Eventually, it will go right up to the actual creation point itself, but for that to happen we’re going to need data virtualization, which is a different topic. We also have to recognize that crushing backup data (which is brilliant, by the way) means deduping files, but in primary capacity we don’t just have files. We need to dedupe blocks, records, objects and so on. Doing it all at backup is cool because we can take all of the data types and amalgamate them into files and deal with them, but we’ll have to get smarter when we move upstream. There are only a handful of people talking about squashing the database, for example. Talk about a big money-play potential. The ROI of squishing data on the most expensive, most complex, most visible transaction systems will be huge. Backup is a pain in the backside for sure, but if dedupe in the backup process has created a few billion dollars of value, imagine what it could do in the transactional world. Video and multimedia will also be huge because of the sheer volume it will consume. Object-based stuff was born to hash, but it’s still not a mainstream play outside of compliance. If you think you’re done hearing about dedupe, it’s about to replicate. ☺

The ROI of squishing data on the most expensive, most complex, most visible transaction systems will be huge.

---

Steve Duplessie is founder and senior analyst at Enterprise Strategy Group. You can see his blog at [http://esgblogs.typepad.com/steves\\_it\\_rants/](http://esgblogs.typepad.com/steves_it_rants/).

# IS YOUR INFORMATION AT RISK?



## LTO Cartridge Memory Analyzer knows!

See the LTO Cartridge Memory Analyzer at work.

**maxell**  
[www.maxell.com](http://www.maxell.com)





*Traditional and not-so-traditional products and services earned honors this year, addressing such issues as data deduplication, solid-state storage and storage cloud services.*

*By Andrew Burton,  
Rich Castagna,  
Todd Erickson,  
Beth Pariseau,  
Dave Raffo and  
Carol Sliwa*

# S

**TORAGE MANAGERS DIDN'T** get much of a break in 2008 as the struggle to keep pace with growing capacity demands continued unabated, but there was still plenty of good news to hearten them. Scores of excellent products were rolled out in the past year, introducing new technologies or adding significant enhancements to tried-and-true storage technologies. As with most technological advances, “bigger” and “faster” were key themes; however, this year they were joined by “efficiency” and “cost savings,” terms that provided even more encouragement to beleaguered storage pros who are quickly learning how to do more with less.

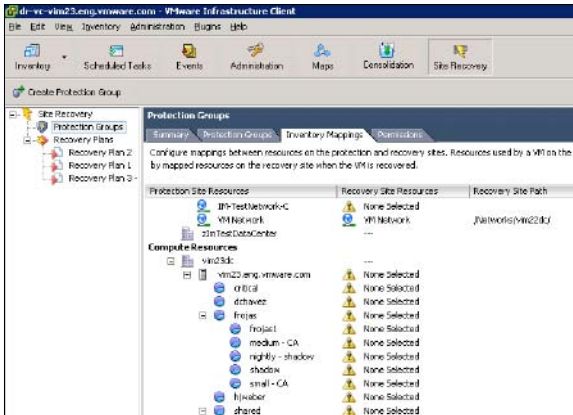
The editors of *Storage* magazine and SearchStorage.com, as well as a panel of judges—including users, consultants and analysts—pored through this bumper crop of new or improved products and selected the 15 we feel stood out based on innovation, performance, ease of integration, functionality and value. Congratulations to all of the storage Product of the Year winners!

# WINNERS!



## BACKUP AND DISASTER RECOVERY SOFTWARE AND SERVICES

# GOLD VMWARE SITE RECOVERY MANAGER 1.0



IT DIDN'T TAKE long for the server virtualization juggernaut to roar through data centers, forever altering the server landscape and even the meaning of the word "server." Led by VMware, virtualization applications have proven to be the most effective consolidation tool yet. But despite all of the benefits server virtualization brought, it's also responsible for causing a few headaches along the way. Just ask any storage administrator.

Virtualization has had a profound effect on the storage side of the shop, posing new challenges related to capacity allocation and data

protection. Traditional methods can be used to back up virtualized servers for business continuity, but while the old methods work, they lack the efficiency and resiliency required in virtual environments.

There's no shortage of vendors addressing this problem, but it has taken the company most responsible for the storage shakeup—VMware Inc.—to come up with some truly elegant solutions, like VMware Site Recovery Manager (SRM) 1.0.

VMware SRM ranked high in all of our evaluating categories and topped all backup and disaster recovery (DR) software and services for value. "Increasing use of VMware makes this software incredibly valuable to every organization," notes one judge. While other judges point out that implementing SRM takes some planning, all agree that the effort is worth it.

With VMware SRM you can automate much of the DR process. You use it to create a DR plan that defines how virtual servers will failover based on the apps they host and their criticality (essentially documenting your DR process and creating a runbook). SRM then lets you test the automated failover scenarios without disrupting the production environment.

The automated recovery plans you build with SRM can be as sophisticated as necessary; you can control which servers are recovered first, the recovery sequence and which servers don't need to be recovered.

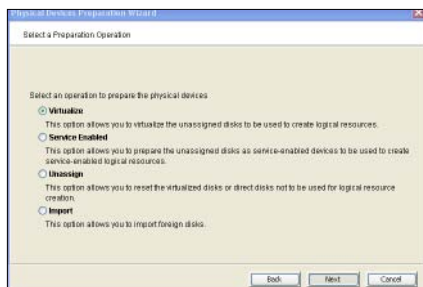
VMware SRM runs on a server at each site involved in the DR plan. It works hand in hand with replication apps from other vendors that tap into VMware "adapters" to integrate with SRM. EMC Corp.'s SDRF, FalconStor Software Inc.'s Continuous Data Protector and Hewlett-Packard (HP) Co.'s StorageWorks Continuous Access are just a few of the popular replication products that have been certified by VMware for SRM. VMware vCenter Server is required to run SRM.

Pricing for VMware Site Recovery Manager starts at a little more than \$2,000 for a single-processor physical server and includes one year of support.



## BACKUP AND DISASTER RECOVERY SOFTWARE AND SERVICES

# SILVER **FALCONSTOR NETWORK STORAGE SERVER (NSS) 6.0**



IT'S NEARLY IMPOSSIBLE to understate the popularity and importance of server virtualization today, and FalconStor Software Inc.'s Network Storage Server (NSS) 6.0 offers a variety of functionality designed to ease storage management in virtual server environments. NSS 6.0 is a software-based SAN solution that acts as a storage virtualization and replication gateway to storage arrays from any vendor, provisioning virtual disks to VMware ESX.

NSS 6.0 is the first software-based storage virtualization and replication gateway to be certified by VMware as a storage virtualization device (SVD), and it enables VMware Site Recovery Manager capability on any VMware-certified storage hardware. NSS 6.0 also supports Microsoft Windows Server 2008 Failover Clustering and Hyper-V virtual server environments.

The product offers integrated support for Fibre Channel (FC), iSCSI and InfiniBand, and allows users to manage virtualized and non-virtualized storage from a single console. It uses thin provisioning to improve storage resource allocation and capacity management.

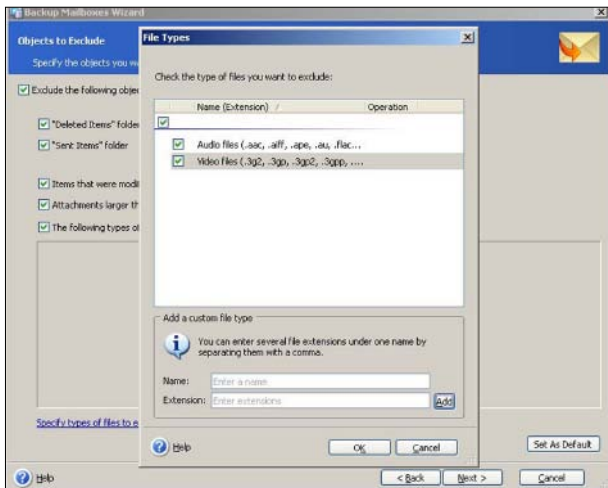
NSS enables VMware Site Recovery Manager with any VMware-certified storage array platform, and works with heterogeneous storage hardware at the primary data center and the disaster recovery site, which could translate into significant savings for many users. NSS 6.0's Application Snapshot Director for VMware complements VMware Site Recovery Manager to ensure that active application data can be replicated to a remote disaster recovery site with full transactional integrity. Finally, FalconStor's MicroScan Thin Storage Replication technology optimizes WAN utilization for more efficient replication and faster failback operations.

Our judges rank NSS 6.0 highly for its breadth of functionality, as well as for its innovation and performance. FalconStor NSS 6.0 is available as software only or as a turnkey appliance. Pricing starts at \$2,000.



## BACKUP AND DISASTER RECOVERY SOFTWARE AND SERVICES

# BRONZE ACRONIS RECOVERY FOR MICROSOFT EXCHANGE



**FOR MOST COMPANIES**, email systems are communications lifelines linking them with their suppliers, customers and contractors; for many firms, doing business as usual means having fully functional, reliable email service.

Plenty of data protection vendors say they can protect and sustain this valuable corporate asset, but few do it as well and as easily as Acronis Recovery for Microsoft Exchange. Although not widely known in the enterprise realm, Acronis Inc. has established a worldwide presence, based largely on the success of its True Image line of bare-metal recovery apps. With Acronis

Recovery for Microsoft Exchange, the company builds on that expertise and incorporates its trademark ease of use.

Acronis Recovery for Microsoft Exchange protects Exchange data stores at the brick or database level with near continuous snapshots of mail transaction logs. A management console lets you remotely manage all instances of the application running on your network. From the console, you can install agents, set options on Exchange servers, and launch recoveries for any Exchange server running the agent.

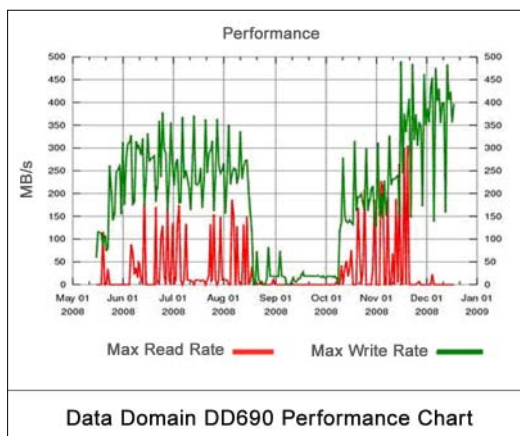
Restores are essentially point-and-click affairs, allowing recovery of entire mail databases, or individual mailboxes or messages. You can also easily pick the point from which you want to recover. One of Acronis Recovery for Microsoft Exchange's most impressive features is its ability to bring the email system back up and available to users within minutes while it's busy recovering the email data in the background. Acronis calls this dial-tone service, providing users with at least basic mail service during recovery, rather than making them wait until the operation completes.

A host of wizards makes deploying, configuring and using Acronis Recovery for Microsoft Exchange easy enough for even those businesses strapped for IT resources to deploy it. One judge calls it "a solid SMB backup product." We agree, especially with a price (\$999 per server) that puts it within reach of most companies. But we suspect plenty of enterprises will give it a long, hard look, too.



## BACKUP HARDWARE

# GOLD DATA DOMAIN DD690 DEDUPLICATION STORAGE SYSTEM



**FOR STORAGE PROS** charged with backup, data growth remains the biggest challenge. Given the implications exploding capacity has on the cost of storage and energy, management best practices, staffing considerations and even data center floor space, users tell us that data deduplication has become a core value proposition when it comes to backup hardware products. But users also say data deduplication vendors will have to push the scalability and performance of their products to keep up with their data growth.

Data Domain Inc.'s data deduplication products have achieved the broadest acceptance, with the company boasting more than 1,800 customers to

date. The DD690 continues the company's scaling approach, which follows the curve of the commodity processor market; the DD690 adds quad-core processors along with 10 Gigabit Ethernet connections. The company's Stream Informed Segment Layout (SISL) architecture means throughput performance isn't dependent on disk I/O availability, so the throughput of Data Domain systems increases with each Intel processor performance improvement, according to the company.

Thus the DD690 boosts the single-stream throughput of Data Domain's dedupe to 600 GB per hour. The DD690 can be configured as a DDX array (up to 16 DD690 controllers in 20U of rackspace) to deliver up to 22 TB per hour of aggregate throughput and 28 petabytes (PBs) of usable capacity. One of our judges calls it "an evolutionary advance for Data Domain."

The DD690's flexibility adds to its appeal. It supports both virtual tape library (VTL) and NAS interfaces, as well as multiple, simultaneous protocols, including Ethernet and Fibre Channel. NAS for disk backup is gaining popularity, most notably among small- to medium-sized businesses (SMBs), which was one of the hottest storage markets in 2008. The NAS interface and enhanced performance also mean the DD690 can be used for nearline or secondary storage, rather than strictly for backup; Data Domain's Retention Lock software also allows it to be used for compliance archiving.

In addition, Data Domain's system addresses some of the other key storage issues facing storage managers this year, including disaster recovery, wide-area networking (WAN), and storage for remote and branch offices. Data Domain offers distance replication software that takes advantage of the data reduction to reduce bandwidth consumption when sending data over a WAN, and the boosted capacity of the box means the fully configured DDX version can support up to 60 remote sites.





## BACKUP HARDWARE

# SILVER

## PERMABIT ENTERPRISE ARCHIVE DATA CENTER SERIES



**REACTING TO THE** significantly amended Federal Rules of Civil Procedure (FRCP), many companies spent much of 2008 crafting their data-retention policies and deploying products to help them comply with the new guidelines for the legal discovery of electronic information. Other organizations deployed the same gear to archive inactive data and ease the impact of data growth on strained back-up and management processes.

Typically retained for long periods, the growth of archived data can become as much of an issue as the one it's trying to solve.

Maintaining data-retrieval performance and keeping ongoing ownership costs in check are critical as the archive repository grows. The ability to migrate data within the repository to take advantage of new technologies and economies of scale is also important. Toss in regulations like the Sarbanes-Oxley Act (SOX) and Health Insurance Portability and Accountability Act (HIPAA), and the situation is further complicated by the need for heightened security and data integrity.

Permagit Technology Corp.'s Enterprise Archive Data Center Series, introduced in early 2008, addresses those challenges with a combination of sub-file deduplication and compression, as well as a grid-based architecture called RAIN-EC. Dividing data among multiple hardware nodes in the grid yields several benefits: it boosts performance to up to 2 Gbps through parallel processing; keeps archive data available in the event of a node failure; allows for rolling migrations to new hardware over time; uses standard interfaces like CIFS, NFS and WebDAV that keep data accessible to multiple applications; and includes features such as data verification, replication, write once read many (WORM) and encryption for compliance.

The data integrity features and grid architecture also mean users can build the archive using space-efficient 1 TB drives, while avoiding lengthy RAID rebuilds in the event of a drive failure. All of that adds up to scalability—up to 3 PB, to be exact. When asked about the product's appeal, one user of the Permagit archive simply repeated the IT pro's data-growth mantra: "Faster and more capacity."



BACKUP HARDWARE

## BRONZE QUANTUM DXi7500



IF 2008'S THEME was ever more data growth, some storage vendors responded to the challenge by offering greater efficiency and more flexibility in their products. The Quantum Corp. DXi7500 data deduplication device was the first to offer users a choice between "in-line" and "post-process" data deduplication approaches. While other vendors bicker over the relative merits of the two methods, Quantum lets users decide which one works best for them.

The DXi7500 also allows both approaches to be used simultaneously for different backup jobs. If users don't want to get so granular with policy settings, the product's "adaptive mode" can automatically adjust the data deduplication process based on the data ingest rate.

The product gives users a choice when it comes to integrating physical tape into backup schemes, neutralizing another frequent bone of contention in the market for deduplicating virtual tape libraries (VTLs). Backup software certified with the device can initiate, track and control all writes to tape, or the DXi7500 can manage copies to tape with shadow tape creation. It can also write copies of backup files to a directly connected tape library, minimizing the overhead on the rest of the environment when creating tape copies.

The DXi7500 scales from 9 TB raw to 180 TB raw, and offers up to 4 TB per hour compressed throughput, according to Quantum. It can be used with the smaller models in the DXi product line to transmit backup data from remote sites to a central location. Replication is asynchronous, automated, encrypted and operates as a background process.

The influence of the DXi line extended well beyond Quantum in 2008. With vendors like Data Domain and Riverbed Technology Inc. already paying royalties for data deduplication to Quantum (based on a patent portfolio it bought with ADIC subsidiary Rocksoft in 2006), DXi was also picked up by EMC Corp. as the basis of a new data deduplication product line in May.



## DISK AND DISK SUBSYSTEMS

**GOLD BLUEARC TITAN 3200 NETWORK STORAGE SYSTEM**

**THE STRONG POINT** of BlueArc Corp.'s Titan unified NAS and iSCSI system can be summed up in one word: performance. Our judges consistently gave the Titan 3200 high marks for performance, and Titan customers often cite performance as the top reason for buying the BlueArc systems.

Bumping up performance was the main emphasis of the Titan 3200 that rolled out in March. The new system offers up to 200,000

IOPS and 20 GBps throughput. Both those numbers were doubled from the Titan 2000. BlueArc's new hardware also supports file systems of 256 TB, 64 virtual storage servers and eight cluster nodes for a maximum capacity of 4 PB, up from 2 PB with the 2000 series.

But BlueArc did more than bump up performance with the upgrade. It added a new set of open application programming interfaces (APIs) that lets BlueArc partners and other developers write applications for Titan. These applications include information lifecycle management (ILM), data virtualization, and data retention and reduction.

The Titan 3200 supports solid state, Fibre Channel, SATA and WORM-protected storage. Like previous Titan systems, it's best suited for applications requiring high bandwidth and IOPS, such as high-performance computing and Web-based services. Hitachi Data Systems sells the Titan 3200 as its high-performance NAS 3000 platform.

Like previous Titans, the 3200 implements its file system in silicon on a field-programmable gate array (FPGA) rather than in software.

One of our judges describes the Titan 3200 as the "fastest NAS in the market with its innovative FPGAs. The value is good because of the ability to consolidate multiple NAS systems, and scalability is impressive."

Starting price for the Titan 3200 is \$100,000.



## DISK AND DISK SUBSYSTEMS

# SILVER INTEL X25-E EXTREME SATA SOLID-STATE DRIVE



**INTEL CORP. PLANTED** its flag in the enterprise solid-state drive (SSD) market with the X25-E (for Extreme) device.

Intel's newest SSDs plug into 2.5-inch SATA drive sockets, and deliver up to 250 MBps sustained read, 170 MBps sustained write, 35,000 IOPS read and 3,300 IOPS write performance. The X25-E is available in 32 GB and 64 GB models. The 32 GB X25-E is capable of writing up to 4 PB (petabytes) of data over a three-year period (3.7 TB/day), and the 64 GB version can write up to 8 PB over that period.

What sets the X25-E apart from other early enterprise SSDs is the way it speeds the write process. Intel uses single-level cell (SLC) flash memory for its extreme drives, storing one bit per memory cell vs. two bits with the more common multi-level cell (MLC) flash memory. MLC drives have greater capacity, but SLC drives perform much faster writes. In the case of Intel's SSDs, its MLC versions write at up to 70 MBps vs. 170 MBps for the SLC drives. SLC SSDs are more expensive, as SSDs remain a premium play.

Verari Systems Inc. is shipping Intel X25-E SSDs in its HyDrive enterprise storage blade and BladeRack 2 X-Series blade servers. Sun Microsystems Inc. has also pledged support for the drives.

"Intel shook things up with this product," says one of our judges, a storage manager. He calls Intel "the first major player in the flash drive market that can really push the performance ceiling."

Flash will need strong performance to make a market impact; X25-E pricing begins at \$695 for 32 GB.



## DISK AND DISK SUBSYSTEMS

# BRONZE

## HITACHI DATA SYSTEMS ADAPTABLE MODULAR STORAGE (AMS) 2000 FAMILY



**BY OFFERING** a SAS backplane and active-active controllers in a midrange storage system, Hitachi Data Systems is at least slightly ahead of its time.

Industry experts agree it's only a matter of time until SAS replaces Fibre Channel as the performance driver in storage arrays, but Hitachi got the ball rolling among major vendors with the AMS 2000 family. The systems are available with SAS or SATA drives.

Hitachi also rolled out a new Dynamic Load Balancing Controller for the platform. The symmetric active-active controller doesn't require logical unit number (LUN) assignments to match preferred paths from server to controller, and can send I/Os to any host port without a performance penalty. The system also monitors controller utilization rates and rebalances the load between controllers. The AMS 2000 series supports disk spin down when no I/O activity is taking place to save power.

The 2000 family comes in three versions: the AMS 2100, AMS 2300 and AMS 2500. Hitachi claims an IOPS cache burst rate of 400,000 for the AMS 2100 and 2300, and 900,000 for the 2500. The sustained throughput ranges from 1,250 MBps for the 2100 to 2,400 MBps for the 2500. Capacity scales to 118 TB on the AMS 2100, and to 472 TB on the AMS 2500.

"This is proof that midtier storage is growing up, with true active-active controllers," says one of our judges, a SAN architect. "It's a real big gain in stability, performance and management for real-world applications."

Hitachi AMS 2100 pricing starts at approximately \$31,500.



## NETWORKING EQUIPMENT

# GOLD RIVERBED OPTIMIZATION SYSTEM (RiOS) 5.0



**RIVERBED TECHNOLOGY INC.'S** RiOS operating system impressed our judging panel with its ability to consolidate and simplify remote-location infrastructure.

RiOS reduces the hardware footprint in remote locations by virtualizing edge services, accelerating application performance and simplifying remote-location administration. The RiOS OS powers the company's Steelhead WAN acceleration appliances, mobile client and Riverbed Services Platform (RSP).

Among the new features RiOS 5.0 introduced were the RSP data services platform that enables virtualized edge services without deploying additional physical servers. Application acceleration features now eliminate the burdensome storage requirement to support widely used applications, including Web-based apps. And an application-level protocol optimization for Microsoft Exchange 2007 adds to the OS's support for Exchange 2000 and 2003.

One judge remarks that Version 5.0 offers "major improvements for Steelhead customers."


The company claims its Steelhead appliances will improve simple file-sharing operations by a hundredfold, cut disaster recovery windows tenfold or more, and reduce WAN bandwidth utilization by 60% to 95%. The judges give RiOS 5.0 high marks for innovation, performance and functionality.

The Steelhead appliance models range in price from \$3,495 to \$129,995. Ten mobile client licenses cost \$3,499, and can be purchased with the appliances or later. Riverbed estimates that organizations will require a license for every three to five mobile users, with the cost per user approximately \$87.



## NETWORKING EQUIPMENT

# SILVER

**BROCADE DCX BACKBONE**

**THE BROCADE DCX BACKBONE** network switching platform allows administrators to consolidate servers and SAN switches, while increasing application service levels and energy efficiency. One networking category judge calls it an “evolutionary upgrade” for the networking market.

Introduced in January 2008, the DCX Backbone provides 8 Gbps throughput, which allows broader server virtualization, data center consolidation, and energy and overhead cost savings.

Each DCX Backbone operates up to 384 ports. The eight-blade slot features 256 Gbps of throughput, which means the Backbone can deliver up to 3 Tbps of chassis bandwidth. Inter-chassis link ports can connect two DCX Backbones to deliver 6 Tbps of dual-chassis bandwidth. One judge calls it “crazy fast.”

It received the highest performance marks in the networking category, with the judges commending

it for functionality and ease of use.

Blade slots can be used to provide 10 Gbps FC connectivity, FCIP SAN extension and fabric-based applications; they can also support 10 Gigabit Ethernet, Fibre Channel over Ethernet (FCoE) and Converged Enhanced Ethernet (CEE) protocols, which simplifies LAN and SAN server connections.

The DCX Backbone uses less than .5 watts per Gbps, which the company says makes it 10 times more efficient than competitive products.

The network infrastructure environment can be administered with the new Data Center Fabric Manager application, and it’s compatible with Brocade’s B-Series and M-Series components. Brocade says DCX Backbone pricing begins in the low six figures and varies by configuration, software licensing and support options.



## NETWORKING EQUIPMENT

**BRONZE** NETEX HYPERIP 5.5

**NETWORK EXECUTIVE SOFTWARE (NETEX) INC.'S** HyperIP 5.5 software enables faster WAN data replication and recovery by accelerating TCP applications using

standard Internet connections. HyperIP provides an alternative to dedicated point-to-point data lines, and is available in a 1U-sized appliance using standard, off-the-shelf components.

The judges give it high marks for ease of integration and value. One judge says HyperIP is “slick data transport optimization that works, is very reasonably priced and does what it says.”

HyperIP technology compresses block-level data up to a 15:1 ratio, which enables it to increase WAN performance to up to 800 Mbps throughput for a single TCP/IP connection, the highest performance of any optimization appliance on the market, according to the company, and 25% to 100% faster than competing products.

As an appliance, it enables data transport at wire speed and leverages the proprietary HyperIP TCP application acceleration technology to mitigate latency and packet loss, and protect against variable circuit quality conditions.

The current HyperIP 5.5 release extends support for multisite configurations, which allows for large-scale, remote data replication deployments. It has more than 8,000 application TCP connections, thereby enabling concurrent replication, migration and recovery processes. A command line interface (CLI) provides single response commands, and multiple-system images let users install new software image upgrades while the software is running.

HyperIP is easy to implement by simply attaching it to a distance-separated Ethernet network segment, so users can leverage existing Ethernet/IP WAN infrastructure. Traffic is pointed to HyperIP via the network, server, NAS device or storage array routing statements. HyperIP parameters and feature options can be modified through the Web-accessible user interface.

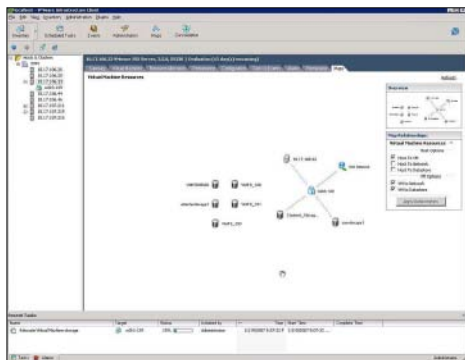
Pricing for HyperIP 5.5 starts at \$6,000.





## STORAGE MANAGEMENT TOOLS

# GOLD VMWARE STORAGE VMOTION



**VIRTUALIZATION WAS ONE** of the hottest topics in storage last year, fueled in no small part by the popularity of virtual server technology. Now VMware Inc., the leading server virtualization vendor, has extended its reach to storage and captured top honors in the management category with its Storage VMotion software.

Users had been asking for a storage version of VMware's VMotion tool, which moves running virtual machines (VMs) from one physical server to another. With Storage VMotion, VMware adds the ability to

migrate running virtual machine disk files within and across data storage systems with no downtime, ensuring continuous service availability and transaction integrity.

Storage VMotion "set the standard for virtual guest availability," according to one judge. "It increases application uptime and IO flexibility—the reason most [people] buy VMware today."

Storage VMotion works by moving a virtual machine's home directory (containing information on configuration, log files, etc.) to the new location before moving the virtual machine disk file. It creates a "child disk" for each virtual machine being migrated, and all disk writes are directed to that child disk once the migration starts. The parent, or original virtual disk, is then copied from the old storage device to the new one. The child disk is reunited and consolidated with the new parent disk, and the ESX host server is redirected to the new parent disk location.

The product is most useful for VMware customers who want to minimize service disruption while performing maintenance or migrating to new storage systems or different classes of storage. It's also helpful for balancing or optimizing storage workloads and addressing performance bottlenecks. Storage VMotion supports both FC and iSCSI SANs, and works with any operating system and application that runs on the VMware Infrastructure.

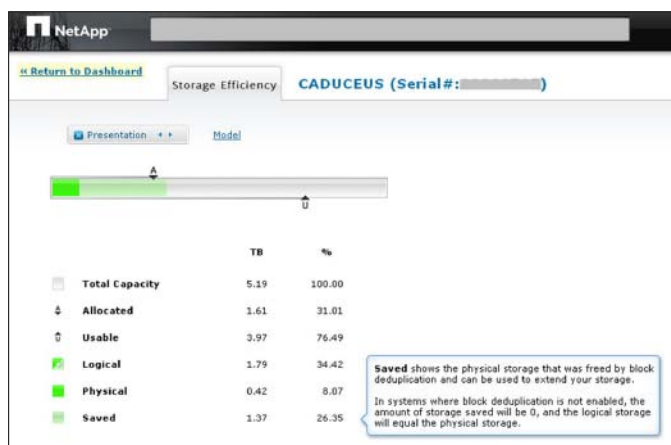
Customers who buy the Enterprise Edition of VMware Infrastructure 3—the server virtualization platform that includes VMware ESX Server 3.5—get Storage VMotion as part of the package. List price is \$5,750 per two processors. Both VMotion and Storage VMotion require VMware vCenter Server (licensed separately) to enable central management of the entire VMware virtual environment from a single console.

VMotion and Storage VMotion are available as part of a standalone package for purchasers of the less expensive Foundation and Standard versions of VMware Infrastructure. The standalone price is \$3,495 per dual processor.



## STORAGE MANAGEMENT TOOLS

# SILVER NETAPP DEDUPLICATION



**DATA DEDUPLICATION TOOK** center stage in 2008, as major storage vendors rolled out backup products, and corporate IT managers either tested the waters or clamored for more information about the red-hot technology.

NetApp Deduplication claimed runner-up honors in the storage management competition, differentiating itself from other large vendors with its ability to work with primary storage data in addition to backup and archival data.

NetApp announced in July 2008 that customers of its V-Series family of

storage virtualization products could also use the deduplication technology to reduce redundant copies of data on disk arrays from other major vendors, including EMC Corp., Hewlett-Packard Co. and Hitachi Data Systems.

Judges awarded NetApp Deduplication technology the highest marks of any storage management entrant in the areas of “ease of integration into environment” and “ease of use and manageability.” NetApp Deduplication is a simple command-based feature that’s free to users of the company’s Data Ontap operating system, which is part of all NetApp FAS Series and V-Series storage systems.

“Porting dedupe into the core OS is a winner,” comments one judge.

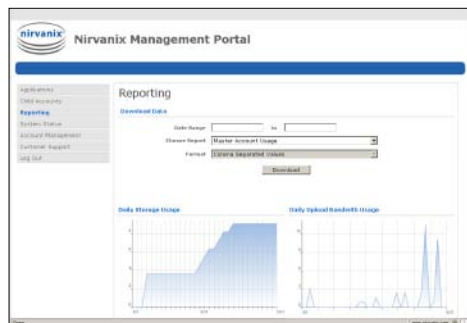
NetApp Deduplication operates as a background process on FAS Series and V-Series systems, and the impact on read/write operations is minimal. Users can schedule deduplication for off-peak times, which can be especially important in minimizing performance hits with heavily used primary storage apps. They can also select the data sets that will produce the greatest benefit, leaving out sets such as encrypted data that won’t deduplicate efficiently.

NetApp’s Web site has a “deduplication calculator” for customers to compute the space savings they can expect to see.



## STORAGE MANAGEMENT TOOLS

## BRONZE NIRVANIX CLOUDNAS



**CLOUD STORAGE GENERATED** a flurry of activity last year as a steady stream of products and services hit the market, so it's fitting that one of the promising new offerings, Nirvanix Inc.'s CloudNAS, rounds out the top products in storage management for 2008.

CloudNAS can transform any Linux or Windows server at a customer site into a virtual NAS gateway to the Nirvanix Storage Delivery Network's (SDN) encrypted offsite storage. The main distinguishing characteristic of the Nirvanix approach is the use of

standard storage protocols to access the SDN online storage service.

Nirvanix CloudNAS mounts the SDN as a virtual drive that can be accessed via NFS, CIFS, FTP or as a virtual tape library (VTL) target, through supported archiving and backup applications. Other storage services, such as Amazon's Simple Storage Service (S3) and Rackspace Hosting Inc.'s Mosso, require development to an API.

Released in October 2008, Nirvanix CloudNAS hasn't had much chance to make an impact yet, but the judges rate the technology as highly innovative. One calls the Nirvanix approach a "great alternative for second-tier unstructured data" with "excellent multisite [disaster recovery] capability."

Once CloudNAS is installed on a server, an administrator can point existing applications and storage processes to it and set file, directory or access permissions. Users access the Nirvanix-mapped drive from their existing applications.

Behind the Nirvanix SDN is a clustered file system that includes all of Nirvanix's globally distributed storage nodes under a single namespace. CloudNAS claims to offer built-in data disaster recovery and automated policy-based data replication on up to three of the geographically dispersed storage nodes.

Data is transferred from CloudNAS via the Internet or Nirvanix cross connect services, which provides a direct connection from the customer site to one of Nirvanix's storage nodes.

CloudNAS is available as a free download for customers that have a contract with Nirvanix for 1 TB or greater. Optional round-the-clock support is available at \$200 per month per server. ☉

---

Andrew Burton is the site editor for SearchDataBackup.com, SearchSMBStorage.com and SearchDisasterRecovery.com; Rich Castagna is Editorial Director of the Storage Media Group; Todd Erickson is the features writer for SearchDataBackup.com, SearchSMBStorage.com and SearchDisasterRecovery.com; Beth Pariseau is senior news writer for SearchStorage.com; Dave Raffo is senior news director for SearchStorage.com; and Carol Sliwa is the features writer for SearchStorage.com.



## 2008 Storage Products of the Year **FINALISTS**

The following products were finalists in this year's competition:

### Backup and DR Software and Services

Acronis Inc. Recovery for Microsoft Exchange

Asigra Inc. TeleVaulting Version 8.0

BakBone Software Inc. NetVault: Backup v8.1 and NetVault: Backup VMware Plugin v1.0

Continuity Software RecoverGuard v3.0

FalconStor Software Inc. Network Storage Server (NSS) 6.0

InMage Systems Inc. DR-Scout 4.2

Paragon Software Group Drive Backup 9 Server Edition

Spare Backup Inc. Spare Switch Data Transfer Version 4 (& Spare Backup Version 5)

Symantec Corp. Backup Exec 12.5 for Windows Servers

Symantec Backup Exec System Recovery 8.5

VMware Inc. Site Recovery Manager 1.0

### Backup Hardware

Cleversafe Inc. Dispersed Storage Network

Data Domain Inc. DD690 Deduplication Storage System

Hewlett-Packard Co. StorageWorks D2D2500 Backup System

Hifn Inc. Express DR 250/255

IBM Corp. System Storage TS1130 Tape Drive

Permabit Technology Corp. Enterprise Archive Data Center Series

Quantum Corp. DXi7500

### Disk and Disk Subsystems

3PAR InServ T-Class Storage Servers T400 and T800

Atrato Inc. Velocity1000 (V1000)

BlueArc Corp. Titan 3200 Network Storage System

EMC Corp. Clariion CX4 Series

Hitachi Data Systems Adaptable Modular Storage (AMS) 2000 Family

IBM Corp. XIV Storage System

Intel Corp. X25-E Extreme SATA Solid-State Drive

Isilon Systems Inc. X-Series Clustered Storage with OneFS 5.0 OS

NetApp FAS3100/V3100 series

Pivot3 Inc. Serverless Computing

Seagate Technology LLC Savvio 10K.3

Xiotech Corp. Emprise 7000

### Networking Equipment

Brocade DCX Backbone

Cisco Systems Inc. MDS 9500

Cisco WAAS

IBM Corp. SAN Volume Controller (SVC)

Network Executive Software (NetEx) Inc. HyperIP 5.5

Riverbed Technology Inc. RiOS 5.0

### Storage Management Tools

EMC Corp. ControlCenter 6.1

Mimosa Systems Inc. NearPoint 3.5

NetApp Deduplication

Nirvanix Inc. CloudNAS

Ocarina Networks ECOsystem

SANPulse Technologies Inc. SANlogics 2.0

Seanodes Inc. Exanodes Virtual Machine edition

StoredIQ Inc. Intelligent Information Management 4.5

StorWize Inc. STN-6000p Series

Symantec Corp. Veritas Virtual Infrastructure 1.0

VMware Inc. Storage VMotion

What's it like to feel 100% confident that  
your backup and restores will actually

# work?

Visit the SearchStorage.com "Advanced Guide to Backup" today:

[www.SearchStorage.com/backup\\_guide](http://www.SearchStorage.com/backup_guide)



*The Web's best storage-specific information resource for enterprise IT professionals*

TechTarget  
Storage Media



**STORAGE**

**Storage Decisions**




Timetable for

# 10 Gigabit Ethernet

*The next generation of Ethernet is likely to have a profound effect on storage; pumped-up iSCSI performance may challenge Fibre Channel's tier 1 dominance.*

*By Christine Cignoli*



**T'S BARELY EVEN HERE YET**, but 10 Gb Ethernet (10 GbE or 10 GigE) is going to have a hard time living up to its hype. Hailed as a “game changing” technology by some, it carries the burden of being a cure-all for storage (and network) managers’ problems. But when you look beyond the hyperbole surrounding 10 GigE, you’ll see the technology is, in many ways, still just emerging.

That’s not to suggest that 10 GigE won’t deliver on its promise—perhaps not a panacea, but certainly destined to give iSCSI storage and Fibre Channel over Ethernet (FCoE) topologies a big boost. Still, real products are few and far between at this time, and per-port prices are still at a very un-Ethernet premium. While there have been some early adopters, its first few inroads into the market have been in higher-end implementations such as super-high-performance computing.

Switch vendors are starting down the 10 GbE road with recent product releases and certifications. But it may still take some time for 10 GbE to hit the mainstream. Brad Booth, chairman of the board of the Ethernet Alliance, thinks 2012 is when costs will be low enough for the market to see widespread adoption. “Generally when a standard

is written, it takes about 10 years before it really hits what we consider the big volume adoption,” he says. Booth calls those 2012 adopters “tier 3,” and says that quicker adopting tier 2 data centers are probably looking at 10 Gb now and considering which cabling they’ll choose. He cites Google as a tier 1 data center; they’re now running a brand-new data center with 10 GbE throughout, according to Booth.

One research institute has already moved to 10 GbE. Bruce Allen, director of the Hannover, Germany-based Max Planck Institute for Gravitational Physics, chose a Woven Systems Inc. switch as part of his early 2008 systems upgrade. “I think 10 Gb is now a proven technology,” he says. And being an early adopter is standard for Max Planck. “We’re a research institute,” he says. “The one guaranteed way to fail is to wait. The wrong decision is waiting.” The institute stores and analyzes data, with approximately 1.5 petabytes (PB) currently managed. Allen says storage is secondary for them, with compute cycles the primary goal. “I wanted wirespeed nonblocking to all nodes, and that meant Ethernet,” he says.

Woven is one of the few vendors offering 10 Gb switches, along with

## 10 GIG RIDES ON COPPER OR FIBER

**OPTICAL FIBER:** The earliest 10 Gb Ethernet components were built around optical fiber, but copper cable-based products soon followed. In most data centers, copper is the standard transport for data and storage networks because it’s relatively cheap and easy to install. But copper cabling typically can’t carry a signal over long distances, so it’s used primarily within buildings, while optical cabling is usually used to link facilities over longer distances.

There are a number of IEEE standards for optical cabling with 10 Gb Ethernet, each designated by a unique identifier. For example, 10GBASE-SR (short distance) is the standard for the cheapest optical implementation that uses standard optical cabling; it can cover distances of approximately 30 meters to 90 meters. The next step up in price and distance is 10GBASE-LR (long range), which can carry signals up to 10km. There are several additional standards for 10Gig over optics that offer greater ranges.

**COPPER CONNECTIONS:** For copper cabling, 10GBASE-CX4 is an IEEE standard approved approximately five years ago that uses familiar twin-axial cable. Also referred to as 802.3ak, this standard is based on cabling and connectors used for InfiniBand, although some reengineering was done so they’re not identical. 10GBASE-CX4 is effective at distances up to 15 meters.

10GBASE-T, approved in 2006, is likely to gain widespread popularity. It uses unshielded (or shielded) twisted-pair cables and will work at up to 100 meters; companies may opt to use already installed Cat 6 cabling, but the distance will be effectively halved. Connectors for 10GBASE-T are familiar, too, RJ-45-style connectors rated at 650 MHz.

Cisco Systems Inc., Force10 Networks Inc., Foundry Networks Inc. (now part of Brocade) and Mellanox Technologies Inc. Joseph Ammirato, vice president of marketing at Woven, thinks 10 GbE will be more affordable by 2010. And, for more widespread adoption, “we need 10 GbE as a default feature on server motherboards,” he says. “Users still have to buy it as a separate add-on.” He targets mid-2010 for 10 Gb to be built-in.

Ammirato names several customer requirements he thinks will drive 10 Gb adoption: consolidation, server virtualization and what he calls desktop virtualization—a need for organizations to better control mission-critical applications, running them from the data center and not user desktops.

Mellanox is trying to move into the 10 GbE market from its base of InfiniBand products. T.A. Ramanujam, Mellanox’s senior product marketing manager, says that in the past few years, Ethernet has come into the spotlight as a consolidation or unifying data center fabric. “To accomplish that,” he says, “they’ve taken features from technologies like Fibre Channel and InfiniBand and added that to [a] new Ethernet specification, what’s going to be called data center Ethernet or converged enhanced Ethernet.”

Ramanujam says per-port prices are approximately \$300 to \$400 for a 10 GbE switch, and predicts that the market will start adopting 10 GbE when its price per-port cost drops closer to Gigabit Ethernet’s cost. “It’s not very far-fetched,” he says, estimating that the second half of 2010 will bring 10 GbE on the motherboard along with more affordable switches. “By then, the price of 10 gig switches will be in the region of \$100 to \$150 per port.”

It also remains to be seen how FCoE, a potential 10 Gb competitor, will affect 10 GbE’s adoption and pricing. Allen at the Max Planck Institute thinks 10 Gb will triumph. “I tend to think that Fibre Channel is one of those technologies that doesn’t really have a good place in the

## A SAMPLER OF 10GBASE-T NICs

EARLY ON, networking vendors were challenged by issues such as power consumption and overheating with their designs for copper-based interface cards for 10 Gig Ethernet. Many of these issues have been resolved. Today, a number of vendors offer 10 Gig Ethernet network

interface cards (NICs) that use the 10GBASE-T standard that supports twisted-pair cables and RJ-45 jacks, including the ones at right:

<i>Manufacturer</i>	<i>Product</i>
Chelsio Communications Inc.	S310E-BT Storage Accelerator
Intel Corp.	Intel 10 Gigabit AT Server Adapter
Mellanox Technologies Inc.	ConnectX EN MNTH18-XTC
SMC Networks Inc.	SMC10GPCIe-10BT TigerCard
Tehuti Networks Ltd.	TN7588-S and TN7588-D (dual port)



future,” he says. “You’re better off sticking to commodity stuff.” But storage users managing the deeply ingrained FC infrastructures of many businesses might disagree.

In the meantime, the Ethernet Alliance’s Booth says the group is busy putting on interoperability demonstrations and educating possible 10 GbE users on what they’ll need to do to be ready for the transition—namely, consider their current server utilization and decide on cabling. “Looking at a server running 1 gig, do they want to upgrade that server by putting in a new [network interface card] and upgrade it to run 10 gig?” says Booth. “Or do they wish to put a better server in there and put this 10 gig card in it?”

Users must also think about whether they can reuse their current optical data center cable, or whether they’ll need or want to replace copper cabling. Allen says the one mistake the Max Planck Institute made in setting up 10 GbE was its cabling choice.

They’re using copper 10 Gb cables, which Allen calls “bulky and inflexible.” For \$50 more apiece, he says, they could have had optical cables.

10 Gb also signals a shift in thinking toward consolidation and centralization. “Unlike previous Ethernet technologies, where primarily volume went to desktop,” says Booth, “now most of the volume of 10 gig is in the data center, into the core of network.” ☉

**"I tend to think that Fibre Channel is one of those technologies that doesn't really have a good place in the future. You're better off sticking to commodity stuff."**

—Bruce Allen, director of the Max Planck Institute for Gravitational Physics

---

Christine Cignoli is the Associate Site Editor for SearchStorage.com and SearchStorage.co.UK.

Special Report: Online Backup

Have you looked at

# Online Backup lately?

Visit the SearchDataBackup.com "Online Backup Special Report"  
to find out how online backup has evolved:

[www.SearchDataBackup.com/online\\_backup](http://www.SearchDataBackup.com/online_backup)

*Sponsored by:*



IRON MOUNTAIN



**SearchStorage.com**

*The Web's best storage-specific information resource for enterprise IT professionals*

TechTarget  
Storage Media



SearchStorage.com



SearchStorage.co.uk

STORAGE

Storage Decisions



SearchDataBackup.com



SearchDisasterRecovery.com



SearchSMBStorage.com



# The under-over on DR

*One size doesn't fit all, especially when it comes to disaster recovery planning.*

**Learn how to build a multitiered DR services capability.**

*By James Damoulakis*

# W

**HEN HURRICANE IKE** struck last September, many feared a repeat of the 2005 Katrina catastrophe. While Ike turned out to be one of the most destructive hurricanes on record, its impact was nowhere near as devastating as Katrina's. This was due, in part, to better preparedness across the board.

From an IT perspective, Katrina raised the level of consciousness regarding disaster recovery (DR) and spurred more organizations to invest in "recoverability."

However, building a DR capability, particularly in the current economic climate, can still be a tough sell. The business justification for DR is based primarily on risk avoidance—a so-called soft metric—rather than on hard cost savings. In addition, DR implementations often involve an investment in infrastructure that mostly sits idle waiting for that fateful day. As a result, organizations may be tempted to shelve DR initiatives or re-prioritize them further down the IT project stack.

Unfortunately, disasters don't distinguish between good and bad economic times. But you could easily argue that preparedness in tough times is even more critical as institutions are less able to tolerate instability.

## EFFICIENT DR

Disaster recovery is an insurance policy and, as with any other type of insurance, we dislike the possibility of paying for something we may never use. Therefore, when formulating a DR strategy, it's important to minimize unnecessary purchases and maximize utilization of DR assets wherever possible. The first step to attaining these goals is a solid understanding of business requirements with respect to DR. Over-designing and over-delivering beyond actual needs meets service-level requirements, but it's inefficient and won't help make your case for future DR dollars. Likewise, short-changing the process to save money is often a good way to get caught off guard when disaster does strike.

When it comes to storage-related disaster recovery services, many organizations fall into one of two categories:

- **ONE SIZE FITS ALL.** There's a single service option, which is often tape-based recovery. For mid- to large-sized organizations, it's highly unlikely that all application needs can be adequately met with a single solution; it's too much for some and not enough for others.

- **TWO SIZES FIT ALL.** In addition to tape, a data replication option exists. While this may be adequate in some cases, we've encountered situations where the gap between the two services levels—tape-based recovery vs. synchronous replication, for example—is simply too large. This forces a choice between a very high cost option that more than meets requirements or a much less-expensive one that falls short.

It's important for an organization to develop a catalog of DR services that can address the required range of recovery time objective (RTO) and recovery point objective (RPO) requirements at appropriate cost differentials to avoid over-delivery.

This top-down approach of gathering business requirements to formulate a strategy must then be tempered with a bottom-up understanding of the "step-function" cost implications of various protection options. The iterative process of testing assumptions and validating business needs against the likely cost to meet those needs results in a more realistic stratification of requirements from which we can develop a DR service catalog that aligns efficiently with the business.

## THE MANY FORMS OF REPLICATION

The key to quick data recovery is replication, and while the gold standard continues to be synchronous replication, the available alternatives have become quite extensive. This allows for a wide range of options at a variety of price points.

Enterprise arrays tend to offer the broadest range of options, including synchronous and asynchronous replication, as well as potentially critical features like consistency groups and various types of multi-hop replication. As these systems now also support a growing range of disk types (from solid state to Fibre Channel to SATA), it's conceivable that a multitier DR

solution can be configured within a single storage platform.

More commonly, organizations look to midrange storage platforms for lower cost replication options. The challenge here is that the replication capabilities of these systems can vary significantly, from vendors that offer essentially the same or nearly the same functionality as on their tier 1 systems to those that offer only basic replication.

In environments with more limited needs—a handful of key applications requiring replication, for example—host- or application-based replication can often be leveraged at an even lower price. These approaches, including database log shipping, volume- or file-based replication, and application-specific replication tools, can provide a very cost-effective way to meet requirements as long as the number of applications and servers remains manageable. But configuring and monitoring a large number of these systems can increase complexity, and a range of software solutions may be required to support different applications. At that point, people typically look to a broader option, typically at the storage level.

In environments with more limited needs—a handful of key applications requiring replication, for example—host- or application-based replication can often be leveraged at an even lower price.

### REDUCING DR COSTS

Server virtualization promises DR many benefits and is causing organizations to rethink their overall DR approach. One of the most significant cost-related impacts relates to the problem of idle systems. Not only can virtualization dramatically reduce the number of physical servers required for a DR site but, due to the very nature of virtualization, these servers may actually be leveraged for daily operations. For example, DR systems may be deployed under normal circumstances for test and quality assurance, but take on the alternate (or additional) role of hosting production virtual machines failed over from the primary site in a DR scenario. When trying to justify DR costs, non-idle, multifunction assets can be the difference between a winning and losing business case.

Given our fascination with virtualization, it's only reasonable to look there for potentially more affordable storage options. The promise of virtualization at the SAN fabric level has yet to be realized on a large scale. One service capability offered by such technology is replication, and it can be done heterogeneously between different types of arrays. This opens the possibility of replicating to a less-expensive array.

Beyond replication, disk-based backup can also play an important role in a multitiered DR strategy. Virtual tape libraries (VTLs) with deduplication and replication capabilities can provide a level of service below

primary storage replication but higher than tape-based recovery. Because data is deduplicated, the bandwidth requirements are usually less; by replicating to a similar platform at the DR site, recovery through restore can be significantly quicker than tape.

A recent global DR study by Symantec Corp. of 1,000 IT managers and C-level executives found that 56% of applications are now classified as mission critical, which is up from 36% in 2007. This has serious implications for the IT infrastructure in that it indicates a continuing rise in business demands and expectations. Meeting these demands in a time of increased budget constraints requires the careful application of the appropriate technologies to satisfy requirements without over-delivering. Understanding not only the potential benefits but the operational implications for these additional options is, of course, essential before heading down a given path. But it's clear that meeting current and future demands will require such a multifaceted approach. ☉

---

Jim Damoulakis is CTO at GlassHouse Technologies, a leading independent provider of storage and infrastructure services. He can be reached at [jimd@glasshouse.com](mailto:jimd@glasshouse.com).



## Scale-out NAS poised for growth

*The compelling economic benefits of deploying scale-out NAS have helped the technology increase its footprint in the general storage space.*

**UNSTRUCTURED DATA WILL** make up 80% of the information in the corporate data center by 2012. Much of that will continue to be the type of file data IT managers have dealt with for years—home directories and file shares—but new file storage requirements are also growing in importance and impact. The massive amount of rich file data generated by audio, video and rich graphics, combined with Web 2.0 applications, is creating demand for innovative file storage products that can economically scale bandwidth and performance to heretofore unheard of capacities. Scale-out NAS refers to systems designed from the ground up for economically dynamic scaling and support of extremely high-bandwidth applications. Enterprise Strategy Group (ESG) defines scale-out NAS as systems that can be independently scaled in multiple directions—processor, bandwidth or capacity—and managed as a single system in a global namespace. Scale-out NAS has seen early adoption in media and entertainment, as well as in high-performance computing (HPC), but the combination of out-of-control file data growth and richer file formats is bringing the requirement for scale-out systems squarely into the data centers of corporate America.

Scale-out NAS architectures have a number of cost advantages over scale-up products.

- **Low entry cost.** The entry cost for scale-out systems varies depending on the minimum configurations supported. Most systems start as small as two nodes and scale out from there. With clustered scale-out systems, you can add resources as needed.
- **Just-in-time scalability.** Because of their modular nature, there's no need to buy (and power or cool) frames, power supplies and mostly empty cabinets in advance of storage capacity. Capacity and/or performance can be added as needed.

Scale-out NAS refers to systems designed from the ground up for economically dynamic scaling and support of extremely high-bandwidth applications.

- **Riding the commodity curve.** Because users don't need to buy frames, processors or disks far ahead of time, they can typically get better pricing as Intel processors and disk prices decline in cost over time.

- **Higher utilization rates.** Better utilization means deferred purchases of new capacity. Because all of the NAS heads in scale-out systems can address the entire pool of usable capacity in the cluster, there's no capacity locked away behind underutilized NAS heads.

- **Reduced change management planning cycles.** When one file is multiple terabytes in size, conventional three- or six-month change management planning cycles are no longer effective. Requirements are unpredictable and time to provision is more important than ever. The modular and easily scalable characteristics of scale-out NAS allow for extremely fast provisioning.

- **Non-disruptive technology refresh.** In a clustered NAS architecture, everything is redundant—data paths, NAS heads and the data itself. Several scale-out vendors provide both forward and backward compatibility with new versions of hardware, firmware and software, so new versions can co-exist in the same system as older versions. There are also several scale-out vendors that provide non-disruptive file migration between nodes. The combination provides users with the ability to do rolling upgrades, plugging new nodes into the system and unplugging nodes when they need to be retired.

- **Ability to scale capacity without scaling headcount.** Essentially, it should be just as easy to manage a clustered storage system with 100 nodes as it is to manage one with two nodes. Scale-out NAS systems enable this through a global namespace. In layman's terms, a global namespace is a virtual representation of a group of physical file systems. It sits between clients and the assorted file servers in a given environment, and adds a layer of abstraction that divorces what the client sees as mount points from the physical server mount points. It's the "secret sauce" that enables a single point of management and non-disruptive data migration.

- **Automated, policy-based management.** Removing the need for human intervention in low-level storage management functions is another way that scale-out NAS reduces management cost. Most systems are plug-and-play—add a storage/processor node, and the system self-discovers and expands the file system, and then incorporates it into load-balancing algorithms on the fly. There's typically no disruption of service and no requirement to plan data layouts, create LUNs or migrate data.

Essentially, it should be just as easy to manage a clustered storage system with 100 nodes as it is to manage one with two nodes.



Based on the compelling economic benefits of deploying scale-out NAS products, it's no surprise that recent ESG research indicates scale-out NAS is increasing its footprint in the general storage space. A recent survey of more than 500 IT managers in North America and Western Europe found that 11% of those surveyed indicate they use scale-out NAS systems today, and 38% plan to deploy it within the next 12 months. Another 37% are actively investigating the technology. Forty-two percent of the IT managers surveyed say they're looking at scale-out options because of the management efficiencies. The bottom line: Scale-out products offer economics that just plain work when it comes to managing vast amounts of file data. ☉

---


Terri McClure is a storage analyst at the Enterprise Strategy Group, Milford, Mass.

# Storage shops holding up under economic stress


**It's only February**, but it's apparent that the repercussions of a stressed economy are filtering down to IT operations. But while storage isn't immune to shrinking budgets, *Storage* readers indicate that their storage environments may dodge a few economic bullets in 2009; as one respondent noted: "[You] can't stop storing data just because the economy is weak." Approximately one-third of those surveyed say their 2009 storage budgets will be lower, while 56% expect increased or stable budgets vs. 2008. Whether budgets are up or down, most storage managers are tightening their belts, ready to forego some features or performance in favor of lower price tags. Thirty-eight percent say they'll put some projects on hold, 27% say they'll consider products that help to use already installed gear more efficiently, and the same number plan to delay some technology upgrades.

—Rich Castagna

**How will your 2009 storage budget compare to 2008's budget?**



**To save money, which of these actions are you likely to take when you make storage purchases in 2009?**



# 48%

Respondents who say power consumption will be a greater consideration when making storage equipment purchases.

**Why is your budget rising this year?**

- 38%** We have planned projects that can't be delayed
- 24%** We still need to add disk capacity
- 14%** Our storage arrays are aging and need to be replaced/augmented
- 14%** The money we spend this year will save us money later
- 5%** Management considers IT purchases strategic investments

**"Lower revenues will force us to buy the most efficient solution available regardless of name-brand recognition."**

—Survey respondent

Scale-out NAS

Build multitiered DR

10 GbE timetable

Best products

Where to start dedupe

## Check out the following resources from our sponsors:

### Bakbone Software Inc., page 8

[Bakbone Software](#)

[Using Business Continuity to Ensure Uninterrupted Operations](#)

### Fujifilm, page 3

[AGPS Surveillance of Data In-Transit](#)

### Maxell Corp., page 11

[Is Your Information at Risk?](#)

### Quantum Corp., page 6

[Quantum Deduplication Solutions](#)

[White Paper: Quantum DXi7500 Deduplicated Backup for the Enterprise](#)