Managing the information that drives the enterprise

STORAGE

Vol. 8 No. 1 March 2009

Sales Competence
Initial Quality
Features
Reliability
Tech Support

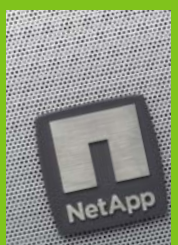

Best of the Midrange Arrays

A surprise winner emerged as readers rated their midrange arrays for service and reliability in our FOURTH ANNUAL QUALITY AWARDS

p. 12

ALSO INSIDE

4 Data on the brink
9 Think data, not equipment
20 Take tape out of remote offices
27 Dedupe directions for 2009
31 Where does dedupe belong in backup?
34 DR testing

Data on the Brink

4 EDITORIAL You might think your company's data is secured and safely backed up, but there's likely to be a lot of crucial data that's out of the reach of your backup systems.
by RICH CASTAGNA

Think Data, Not Equipment

9 STORAGE BIN 2.0 If IT can start treating every decision from the perspective of the data itself rather than all the equipment that's been acquired to hold it, our data centers could cost less and be more effective at the same time.
by STEVE DUPLESSIE

Quality Awards IV
Dell Plus EqualLogic:A Winning Combination

12 In our fourth annual service and reliability survey for midrange arrays, Dell's midrange systems shake up the field by soaring from a last-place finish in 2008 to first this year.
by RICH CASTAGNA

Remote-Office Backups Made Easy

20 With plenty of economical disk-based backup products and cloud-based services available, remote offices can be brought back into the data center fold. by RICK COOK

Dedupe Update: What's Coming in 2009

27 Everybody knows that the hottest thing in storage in 2008 was data deduplication. Don't expect it to cool off in 2009, as even more storage vendors are getting into the dedupe race.
by DAVE RAFFO

Where Does Deduplication Belong in Backup?

31 HOT SPOTS The benefits of data deduplication in disk-based backup are evident, but with backup applications now incorporating the technology you'll have to decide if software- or hardware-based dedupe is best for your shop.
by LAUREN WHITEHOUSE

More Testing, More Confidence in DR Plans

34 SNAPSHOT According to the latest *Storage* magazine survey, more companies are testing their disaster recovery (DR) plans regularly, and storage managers are more confident that their DR plans can avert a significant business disruption.
by RICH CASTAGNA

Vendor Resources

35 Useful links from our advertisers.

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

Quantum
BACKUP. RECOVERY. ARCHIVE.

Data on the brink

You might think your company's data is secured and safely backed up, but there's probably still a lot of crucial data that's out of the reach of your backup systems.

OUR DATA CENTER storage systems are buttoned up and battened down, and everything has been deduplicated and replicated. But what about all of those laptop computers, smartphones, PDAs, thumb drives—and even MP3 players—out there? They might be out of sight, but they shouldn't be out of mind if you're truly serious about safeguarding your company's data.

A lot of companies try to keep all user files on network storage where they can be centrally managed, and properly backed up and archived. But remote workers and offices, which often use portable devices, can be disconnected from the corporate local-area network (LAN) for long periods of time. The data those devices create is usually called "edge" data, data that resides on the far reaches of an enterprise's network. But it's really data on the brink.

We're talking about an awful lot of stuff that may be falling through the data protection cracks. Some industry experts say that in many companies, the volume of data created (and stored) on the fringes may amount to even more than is created on data center storage. That might overstate the case a bit, but there's little doubt there are enough Word files, Excel spreadsheets, PowerPoint presentations and other business documents floating around to make any storage manager just a little less confident about data protection.

Some industry experts say the volume of data created on the fringes may amount to more than is created on data center storage.

There are actually two separate issues. The first is backup or just making sure there are copies of all those files so they can be recovered if needed. Then there's the compliance and security angle. With all that stuff living on devices that tend to get left in taxi cabs, pinched by airport thieves or simply misplaced, backing up the data isn't enough.

There are plenty of solutions to address backup, although none of them could be considered perfect. There are backup applications that specifically address portable data protection, like Atempo Inc.'s LiveBackup, IBM Corp.'s Tivoli Continuous Data Protection for Files and Yosemite Technologies' (now a Barracuda Networks company) FileKeeper, which

are just a few examples in a very crowded field of products. These can continually check for changes and ship them back to a central repository. They keep working when the portable PC isn't connected to the company network and update the repository upon reconnection.

The gap between connection times could be a little risky because even though the backup data has been collected and isolated, it's still on the same machine until it's hooked into the network again. A relatively "gapless" alternative is to use online—or "cloud"—backup services exclusively to protect mobile data. Generally, these services work the same as the backup apps, but don't require connecting to the corporate LAN, so any Internet connection will do. You'll have to be comfortable with having a third party hosting your backup data, although the services that align with enterprise needs offer central management and may even integrate with corporate backup systems to some degree.

So, if someone swipes a laptop from one of your users, but one of the backup methods mentioned here was being used, it shouldn't be such a big deal to recover the data. But backup won't keep the crook from looking at the data and, if it's sensitive company information or a customer's personal data, that is a big deal. That's the compliance and security side of the mobile data equation.

The Trusted Computing Group, an industry consortium, is helping to create standards for hardware-based computer security.

The Trusted Computing Group (TCG), an industry consortium, is helping to create standards for hardware-based computer security. The TCG's Storage Work Group, as its name implies, focuses on security for storage devices, and it recently released a few sets of specs for hard disk drive encryption. But you don't have to wait for disk vendors to get around to implementing the specs as most drive manufacturers sell self-encrypting disks. These drives were developed specifically for laptop computers where there's the greatest risk of losing data, but they're slowly making their way into data center-class storage systems.

Seagate LLC, a TCG member, has been shipping encrypting drives for more than three years. Other TCG members also offer full-disk encryption drives: Hitachi Data Systems shipped its first encrypting drive nearly two years ago; Fujitsu (which sold its disk operation to Toshiba) rolled out theirs nearly a year ago; and Samsung, Toshiba and Western Digital Corp. all offer self-encrypting drives.

Data on the edge doesn't have to be data on the brink of disaster. With so many options available today, you should be able to find one that's right for your environment. Many of you are probably well along in this effort, so drop me a line and tell me what you're doing to protect mobile data.

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

** Click here for a sneak peek at what's coming up in the April issue.*

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™

COMING IN APRIL

Data Replication

Replication has become a key component of most organization's disaster recovery and data protection environments. And with its growing popularity, the number of replication options has also grown. We examine the pros and cons of different data replication methods and products, and offer a look at the state of the art for replication technologies.

How Storage Jobs are Changing

Virtualization has a grip on the server side of IT shops, and is making inroads into storage. Add network virtualization into the mix, and traditional barriers between storage, systems and networking groups begin to fall. Who will manage what in the data center?

Power-Smart Disk Systems

Disk drive systems use more power than just about any other data center gear, but storage vendors are addressing this problem with a variety of technologies. This update describes how vendors are enhancing their disk systems and offers some tips on reducing storage power consumption.

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

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

Subscriptions:
www.SearchStorage.com

IS YOUR INFORMATION AT RISK?

LTO Cartridge Memory Analyzer knows!

See the LTO Cartridge Memory Analyzer at work.

maxell
www.maxell.com

Think data, not equipment

If IT can put itself back into a position of treating every decision from the perspective of the data itself, our effectiveness could be optimized.

RECENTLY, I HAD A DISCUSSION with a senior operations manager at a major European telecommunications company. We were discussing his experience with a backup consolidation effort (hundreds of remote data sites) that leveraged virtual tape library (VTL) systems with data deduplication technologies. What I learned was entirely unexpected.

It seems that not only was the company able to realize the assumed and obvious benefits one would expect in this situation, but because the new process was so much more efficient, it altered the way the company fundamentally viewed all of the processes and services IT delivered to the business. There were substantial gains in far-reaching areas—from an ability to significantly improve their IT delivery capabilities—and across regulatory bodies and borders to security, privacy and beyond. Because of the success of what began as a fairly simple consolidation exercise in the area of backup and recovery, the company is now aggressively investigating other opportunities to improve efficiency and gain even greater benefits. The firm has enjoyed so much success, so quickly, that even in this economy they're accelerating the consolidation of several European data centers. If all of this were happening because of something as seemingly simple and benign as backup consolidation, I was interested in what other value could be derived by pushing overall efficiency improvement through other mainstream consolidation efforts.

It's interesting when you hear of people spending money in times like these, rather than talking about doing more with less.

It's interesting when you hear of people spending money in times like these, rather than talking about doing more with less. Either they're insane or they're able to find real value and return on their efforts. No one spends money today just because they can.

My European friends found such value because by solving their original problem, the solution created so many obvious downstream benefits that it became impossible to ignore. Consolidation began for one reason, but the results crossed so many parts of IT that it was impossible not to see them.

About the same time this effort began, I went to China and spoke at a "Green IT" conference in Beijing. I talked about how "green" is simply a modern metaphor for "efficient," only instead of operating efficiency it focuses on power, cooling and space efficiency. What makes us bad green citizens is the same thing that makes us bad service providers: 50 years of having way too much stuff. You don't need to hold an advanced degree to realize that more stuff is way harder to deal with than less stuff.

After my speech, I spoke with a senior official responsible for China's global economic analysis who was curious as to why I spent so much of my presentation on process vs. technology. My answer was simple: History has left us to contend with the sins of our past. Process can be changed even when economics don't allow us to change technology as easily. "But why do you focus on the symptom [meaning infrastructure] before the cause [meaning data]?" he asked. It was a great question. Most commercial Chinese IT operations are relatively new and unburdened by many of the long-term historical issues Europe and North America face. China's relatively new venture into the world of IT means they haven't faced the infrastructure "sprawl" issue others have spent 50 years dealing with and can stay focused on the "cause," which is the data itself. Most of us aren't so lucky. We have to deal with a lot of history and the problems past decisions have placed on current realities.

Process can be changed even when economics don't allow us to change technology as easily.

This realization helped me see the current situation more clearly. If IT can put itself back into a position of treating every decision from the perspective of the data itself—and not simply the previously established infrastructure—our effectiveness could be optimized.

Global macro-economic conditions are spotlighting IT operations and capital spending in both a positive and a negative light. The downside is that harsh economic realities force business cost centers such as IT to support continued demands and data growth with flat/diminishing budgets and IT resources. The positive aspect is that it also forces IT to become creative and even challenges IT shops to re-evaluate everything they may have taken for granted during economic boom times. Investigating and implementing efficiency improvements across the organization is always good, but they're often overlooked or patently ignored in good economic times. Prosperity, unfortunately, has a tendency to relax our standards. ⊙

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

Sales Competence
Initial Quality
Features
Reliability
Tech Support

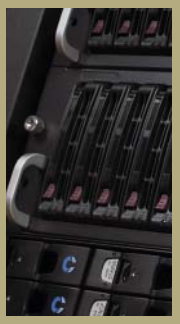

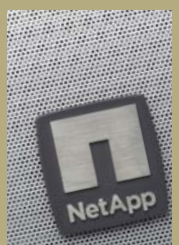

Dell plus EqualLogic: A winning combination

Dell's midrange arrays soar from last to first in our user service and reliability survey.

By Rich Castagna

IF THE PEOPLE AT DELL INC. are reflecting with more than a little satisfaction on their decision to acquire EqualLogic Inc., it's certainly understandable based on the most recent results from the *Storage* magazine Quality Awards (see "About the survey," p. 13). Until it acquired EqualLogic, Dell had relied mainly on its OEM versions of EMC Corp. Clariion arrays for its midrange storage line, and it didn't fare well in any of our previous surveys, ranking ninth out of nine in 2008. But adding EqualLogic to its roster of storage systems appears to have had a dramatic impact. In the 2009 edition of the Quality Awards for midrange arrays, Dell has fashioned

a dramatic reversal and leaped to the top of the field based on user surveys that for the first time combined Dell's and EqualLogic's products.

This was not the first time EqualLogic products made it to the top of the charts. EqualLogic won the award in 2007 and was runner-up to Compellent Technologies Inc. in 2008—missing out on the top spot by the slimmest of margins. However, Compellent wasn't able to repeat last year's performance, falling to sixth place in the current survey. IBM Corp. also made a significant move up the ranks, rising from a seventh-place finish last year to second this year. IBM was remarkably consistent across all categories of our survey, placing in the top four in every evaluation category—a feat matched only by Dell.

NINE PRODUCTS JOCKEY FOR POSITION

Of the 15 companies and their product lines included in the survey, nine garnered a sufficient number of responses to be considered as finalists, the same number as last year (see "Products included in the survey," p. 14). For the first time, Xiotech Corp. was included in the final results, but its debut was somewhat inauspicious as it finished last overall. With the same number of finalists, the overall ratings were much tighter this year. The spread between first and last in 2008 was 1.44 on a 1.00-8.00 scale; this year the spread was just 0.60. Respondents also seemed less euphoric about their products, as the winning score dropped from Compellent's 7.02 last year to Dell's 6.55 this time around.

Although Dell won the award decisively, it was by no means a runaway. Of the five categories in our survey, four different companies came out on top in at least one category. Dell won two categories, while Compellent, IBM and NetApp each took a turn in the spotlight. The scores bounced

ABOUT THE SURVEY

The *Storage* magazine Quality Awards are designed to identify and recognize products that have proven their quality and reliability in actual use. The results are derived from a survey of qualified *Storage* readers who assessed products in five main evaluation categories: sales-force competence, product features, initial product quality, product reliability and technical support. Our methodology incorporates statistically valid polling that eliminates market share as a factor. Our objective is to identify the most reliable product on the market regardless of vendor name, reputation or size. Products are rated on a 1.00-8.00 scale, where 8.00 is the most favorable score.

The respondent pool for this survey skewed toward the small- and medium-sized enterprise, with 70.9% of respondents from companies with less than $1 billion in revenue. In keeping with most of the prior surveys, financial services/banking was the most represented industry (15.8%) followed by healthcare/pharmaceuticals (13.3%) and IT services (12.3%). Most respondents had operationally oriented titles, with storage administrators being the largest group (29%). In addition, 8% were either CIOs or CTOs, which is a somewhat higher percentage than usual; this probably reflects the composition of smaller organizations where a senior individual is more likely to have direct interaction with storage issues.

around for Hewlett-Packard (HP) Co. and Hitachi Data Systems, but both companies always placed somewhere in the middle. And while EMC, Sun Microsystems Inc. and Xiotech notched respectable scores, they placed consistently toward the back of the pack.

Perhaps reflecting an increasing focus on midrange systems, this year's survey attracted the most responses of any of the four midrange array surveys we've fielded, with 631 respondents providing 1,056 system evaluations. EMC had the largest pool of evaluations with 262, followed by NetApp (155) and HP (151). Dell had 94 evaluations.

Of the individual respondents, 39.5% had arrays from two or more vendors. For those who provided evaluations of more than one vendor, we compared the head-to-head ratings. The results were quite surprising, turning the order of finish on its head. Xiotech led the finalists with a 62.5% win rate. In contrast, Dell won just 22.4% of its comparisons, besting only Compellent among the finalists (18.8%). There's no direct relationship between head-to-head results and order of finish, but it's unusual to have the lists nearly reversed.

PRODUCTS INCLUDED IN THE SURVEY

The following product lines were included in the Quality Awards IV survey for midrange arrays:

- 3PAR Inc. InServ E200*
- Atrato Inc. Velocity1000 (V1000)*
- BlueArc Corp. Titan 2000 Series (iSCSI)*
- Compellent Technologies Inc. Storage Center
- DataDirect Networks Inc. S2A6620*
- Dell Inc. CX Series or Dell EqualLogic PS Series
- EMC Corp. Clariion CX Series
- Hewlett-Packard (HP) Co. StorageWorks EVA Series
- Hitachi Data Systems Universal Storage Platform VM, Thunder 9500 V Series or AMS Series
- IBM Corp. FAStT or DS4000/DS6000
- LeftHand Networks Inc. SAN/iQ*
- NetApp FAS200/FAS900/FAS2000/FAS3000 Series
- Pillar Data Systems Axiom 500/600*
- Sun Microsystems Inc. StorageTek 6000 Series or StorageTek FlexLine Series
- Xiotech Corp. Magnitude 3D or Emprise

** Didn't receive a sufficient number of responses to be included among finalists.*

SALES-FORCE COMPETENCE

While Dell dominated the overall survey, defending champ Compellent earned its one bright spot in the sales-force competence category, albeit by a whisker. Compellent scored a 6.43 to lead this category, just 0.01 higher than Dell. This was the lowest high score of any category. The way the two companies achieved their scores was somewhat different as well. Compellent had its highest rating (a 6.77) for the statement "My sales rep is easy to negotiate with."

MIDRANGE ARRAYS

- Dell CX Series or Dell EqualLogic PS Series
- IBM FAStT or DS4000/DS6000
- Hewlett-Packard StorageWorks EVA Series
- NetApp FAS200/FAS900/FAS2000/FAS3000 Series
- Hitachi Data Systems Universal Storage Platform VM, Thunder 9500 V Series or AMS Series
- Compellent Technologies Storage Center
- EMC Clariion CX Series
- Sun StorageTek 6000 Series or StorageTek FlexLine Series
- Xiotech Magnitude 3D or Emprise

Based on a 1.00-8.00 scoring scale

SALES-FORCE COMPETENCE

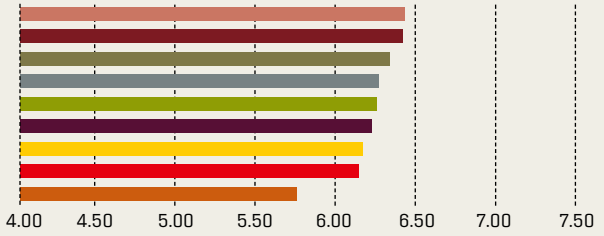

INITIAL PRODUCT QUALITY

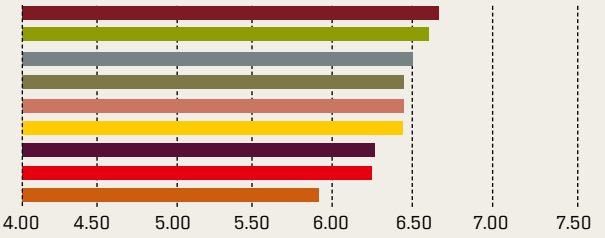

TECHNICAL SUPPORT

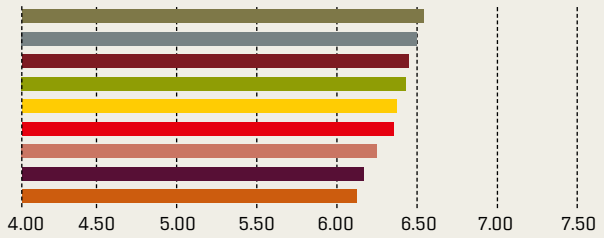

OVERALL RANKINGS

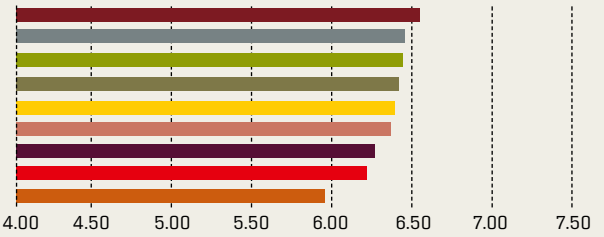

PRODUCT FEATURES

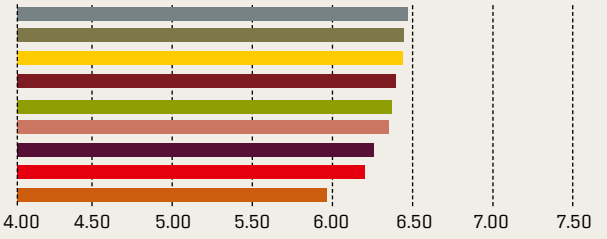

PRODUCT RELIABILITY

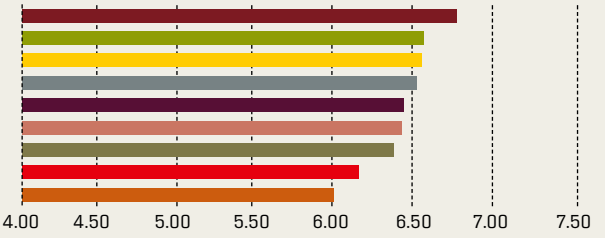

BUY AGAIN: All things considered, would you buy this array again?

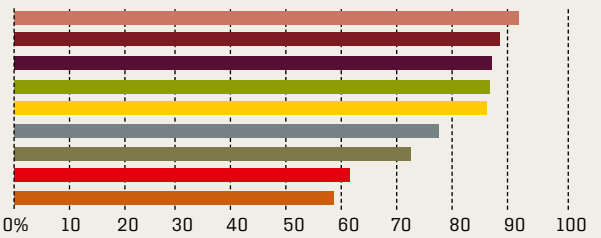

RelayHealth, a division of Atlanta-based McKesson Corp., has three Compellent arrays with another on its way. “Sales came in as a team—there was a sales rep and a storage architect,” says Whitney Gray, storage architect at the healthcare provider. “They worked with us on price, [and] they worked with us on technical configuration,” says Gray. “They were very good to work with.”

Dell’s best score in this category (a 6.59) was for “The vendor’s sales support team is knowledgeable.” Compellent received a high mark (6.68) for “My sales rep keeps my interest foremost,” while Dell fared well with a 6.58 rating for “My sales rep is easy to negotiate with.”

“They took time to learn what we did, to understand what we were doing,” says Donald Wilkins, IT director at Navicure Inc. in Duluth, Ga., of his experience in purchasing multiple EqualLogic arrays. And since Dell’s acquisition of EqualLogic, “for us it hasn’t changed that much,” he says. “The majority of their support staff is still in place, and they still provide some of the best support around.”

IBM apparently has the most knowledgeable sales support team, having received a 6.63 in this regard. NetApp’s sales reps apparently do their homework, as its 6.61 score for the statement “My sales rep is knowledgeable about my industry” outpaced the other eight finalists. Miles O’Neal, IT specialist at Austin, Texas-based engineering firm Intrinsity Inc., was pleased with his dealings with NetApp’s sales team. “They spent a lot of time with us answering our questions, talking about roadmaps,” he says. At the time Intrinsity was considering the NetApp equipment, they were also looking at other vendors’ offerings. “They knew they were up against some competition,” say O’Neal.

NetApp placed third overall in this category with a 6.33. The overall ratings in this category for all vendors topped 6.00, except for Xiotech’s 5.75.

> "Sales came in as a team—there was a sales rep and a storage architect. They worked with us on price, [and] they worked with us on technical configuration. They were very good to work with."
>
> —Whitney Gray, storage architect, RelayHealth

PRODUCT FEATURES

The product features category asks users about the usefulness of various product attributes, including management capabilities, mirroring and replication. This category illustrated the topsy-turvy nature of the top four positions in our survey. IBM bested a bunched-up field with a 6.48, just ahead of NetApp (6.46) and Hitachi (6.45), with Dell a not-so-distant fourth with a 6.40.

“Everything we’ve needed has been there,” says Mathew Colona, IT manager at Pite Duncan LLP in San Diego, referring to the Hitachi Adaptable Modular Storage 200 (AMS200) they installed about eight months

ago. “With the AMS, I actually think a layperson could use it,” he says. “I think that’s how easy it is; it’s really a point-and-click deal.”

IBM’s best score was for the statement “This product scales to meet my needs” (6.68). Users also rated IBM highly for its overall feature set, giving it a 6.60 for “Overall, this product’s features meet my needs.”

Despite finishing sixth in the product features category, Compellent had the highest score for the statement “Overall, the product’s features meet my needs” (6.75). Hitachi’s 6.66 was the second highest rating for that statement.

Having a comprehensive feature set built in is another advantage that midrange systems users consider. “You buy the product and you’ve got all the features,” says Navicure’s Wilkins of his Dell/EqualLogic arrays. “It was not à la carte pricing, it was all-inclusive.”

Having a comprehensive feature set built in is another advantage that midrange systems users consider.

A full slate of features is certainly a requisite for competitive midrange arrays, but performance counts, too. “When we run any type of disk I/O or disk benchmarking utility against it, it just blows us away with the I/O,” says RelayHealth’s Gray of their Compellent systems. “Performance is just screaming.”

Compellent’s category score was dragged down by a 6.09 for “This product is interoperable with other vendors’ products.” While not a bad score, it was enough to lower Compellent’s average score for the category.

NetApp’s highest score (a 6.75) was in response to the statement “This product’s snapshot capabilities meet my needs.”

INITIAL PRODUCT QUALITY

One of Dell’s category wins was in the initial product quality category. With a 6.67 rating, it nudged out HP (6.60), while IBM ranked third (with a score of 6.50). Very few scores higher than a 7.00 were given to any of the category statements, but Dell had a 7.17 for “This product was installed without defects.” Compellent also received a very high 7.00 for that statement but, once again, its category average was pulled down by a relatively weak 5.75 score for “This product requires very little vendor intervention.”

Apparently, the combination of Dell and EqualLogic has also been good for product quality. “Their quality assurance is still one of the best processes they have out there,” says Navicure’s Wilkins. “They just don’t ship out a defective product.”

HP’s second-place finish in the initial product quality category was helped by a 6.74 for the statement “This product was installed without defects,” plus consistency across the board in the category. Its lowest score, a still-solid 6.53, was for “This product is easy to use.” The highest-

rated product in regard to ease of use was Compellent (6.83).

The initial product quality category includes a key statement regarding the system's value: "This product offers good value for the money." Compellent was rated most highly in this regard with a 6.68. HP was second with a 6.59, followed closely by IBM with a 6.55. Dell was in the middle of the overall pack for this statement with a 6.44 rating.

While Hitachi ranked only sixth in the category, its score was still just 0.24 points less than leader Dell. "Smooth, smooth" is how Pite Duncan's Colona describes the installation of his Hitachi Data Systems' AMS200. "Hitachi includes installation—they won't sell it to you if they don't install it. It's a warranty issue," he says.

PRODUCT RELIABILITY

Dell's other category victory was in product reliability, a notable achievement as it led both the initial product quality and reliability categories. In reliability, Dell's 6.79 was the highest single category score in the survey. HP followed right behind Dell in this category with a 6.57, and Hitachi Data Systems clinched third with a score of 6.55.

Dell's top category score was highlighted by a very high 6.99 for the statement "This product experiences very little downtime." It also garnered an impressive 6.87 for "This product requires very few unplanned patches." Dell didn't do quite as well for "Patches can be applied non-disruptively" (where it scored a 6.56), but it was still the high score of the group for that statement.

Sometimes, direct vendor support is so personalized that it seems more like it's a local reseller providing help than a big, distant vendor.

Wilkins says Navicure's experiences with their Dell/EqualLogic systems reflect that solid score. "Their firmware process has pretty much been a non-event for us," he says. "We haven't had any problems over the last few years."

Compellent had the second-highest score for the "patches" statement with a 6.50, but the statement was NetApp's undoing in this category, as it scored a 5.43, well below its 6.39 category average.

Although he concedes there were some minor issues with a firmware upgrade to his NetApp array, Intrinsity's O'Neal is more than satisfied with the system's reliability. "In terms of day-to-day operations, it's been rock solid," he says.

TECHNICAL SUPPORT

In the technical support category, NetApp took its turn at the top with a 6.54 that beat out IBM's 6.50 and Dell's 6.44. Intrinsity gets support for its NetApp machines from their reseller, a situation O'Neal describes as "pretty good" although he adds, "it's not quite as good as when we're dealing

direct with NetApp." But with both the reseller's and NetApp's services, he says, "Overall they work hard, and things get fixed generally pretty quick."

Sometimes, direct vendor support is so personalized that it seems more like it's a local reseller providing help than a big, distant vendor. "They gave us an 800-number," says Pite Duncan's Colona regarding Hitachi Data Systems' support for his company's AMS200, "but the engineer who's on site actually gave us his card and cell phone number and said just call him anytime."

This category rarely offers strong differentiators among vendors and such was the case again. The range between NetApp and last-place Xiotech's 6.11 was the narrowest of any category—a mere 0.43 points. This was also Xiotech's highest category score.

NetApp's top score in this category was a 6.97 for the statement "Vendor supplies support as contractually specified." Among NetApp's survey respondents, 45% purchased their NetApp gear through a reseller. Thus, its high 6.63 score for the statement "Issues rarely require escalation" speaks well of both its own first-level support as well as any first-level support that might be provided by its resellers. Its lowest category score was a 6.08 for "Vendor provides adequate training." It should be noted, however, that most vendors received lower ratings in this area.

BUY IT AGAIN?

When our respondents were asked whether they would make the same purchase decisions today, the results were somewhat surprising. Despite its sixth-place finish overall, Compellent had the highest "Yes" responses to our repurchase question (91.7%). It's no wonder Compellent received a high buy-again rating with satisfied customers like Relay-Health's Gray who says of his firm's Compellent array, "It's the greatest thing since RAID."

Dell was second with 88.3% of those surveyed saying they would repeat the same purchase—a huge increase over its 2008 rating of 60.7%. EMC also jumped up significantly in this regard, garnering an 87.0% positive response vs. a 73.6% for last year. HP had a rating of 86.8%.

This year's survey indicates a dynamic, competitive storage system market segment. With 82% of respondents indicating they would make the same purchase decision again, it's clear that users are generally quite satisfied with their choice of vendors. In this economic climate, that's something to hang on to. ⊙

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

Remote-office backups made easy

With plenty of economical disk-based backup products and cloud-based services available, remote offices can be brought back into the fold.

By Rick Cook

IF BACKUPS ARE A HEADACHE for storage managers, backing up data at remote offices is a migraine. The traditional method of handling remote backups is to have a tape system on site so that local staff can run regular backups. This has a number of obvious drawbacks, not the least of which is reliability because of the need for human intervention.

To address the issue of the manual nature of tape-based backups, more and more companies are replacing tape with other solutions in their remote offices. The trend is nascent, but catching on fast.

Remote offices are generally characterized by relatively small amounts of data that have to be backed up, as well as a lack of technically adept staff. In an office with just a couple of gigabytes to back up, tape can be relatively expensive. While tape still has the lowest cost per gigabyte, that's only true of the system as a whole above a certain threshold. Tape drives aren't cheap, and tape loaders and libraries are even more expensive. If the backup volume isn't big enough, the cost of the tape system, especially hardware, dominates the economics. If you use a system that costs several thousand dollars to back up a few gigabytes' worth of files, the economics are seriously skewed.

Illustration by Neal Aspinosa

THE PROBLEM WITH TAPE

But the biggest hitch in a remote tape backup scenario, and the one most likely to jeopardize effective data protection, is that most remote offices don't have technically knowledgeable people on staff who can maintain a tape system. Backups are often handled by an administrative assistant or clerical staffer with little or no training in the tape technology. Even a grandfather-father-son rotation is hard to handle consistently and more complex rotation schemes like the Tower of Hanoi are just about impossible. A multi-tape loader or library mitigates the problem, but those devices are more expensive than a single drive and don't completely eliminate the difficulty. Even a grandfather-father-son scenario requires rotating the oldest tape out of the loader or library for offsite storage on a regular schedule.

"Tape management is a problem; we felt a lot of time and energy was being spent daily just managing tape."

—Edward Ruffolo, IT director, Miron Construction Co. Inc.

"We often had problems," says Edward Ruffolo, IT director at Miron Construction Co. Inc., a Neenah, Wisc.-based firm that's in the process of converting its tape backup to an ExaGrid Systems Inc. disk-to-disk storage device. "You'd think something was on this tape, go to restore it, and there was something wrong with the tape or the drive was dirty," says Ruffolo. "Tape management is a problem; we felt a lot of time and energy was being spent daily just managing tape.

"A receptionist or someone has to change tapes daily," adds Ruffolo. "They mean well, they're trying hard, but they're busy." There have also been cases where someone has left the same tape in the drive all week and, in one case, an employee locked the tapes in a safe and lost the key.

Miron Construction has four remote offices spread across the upper Midwest, some of which are four or five hours from its Neenah headquarters. That makes it hard for the IT staff in Neenah to handle problems.

It could be worse. Grey Healthcare Group Inc., a New York City-based healthcare advertising firm, was very nearly a worst case for tape in remote offices. The company was using a combination of LTO-1, LTO-2 and LTO-3 to back up its remote offices, and a DLT library to back up 8 TB to 13 TB of data at its central location. Not only was its tape capacity maxed out, but the tape systems were old and growing unreliable, requiring constant maintenance efforts by the company's small IT staff to keep them running. Even when the tape systems were running, they weren't running very well. The worst part, says Chris Watkis, Grey Healthcare's director of information technology, was that the backups were unreliable and couldn't be counted on to recover lost data.

Grey Healthcare replaced its tape system with two FalconStor Software Inc. Enterprise Storage Appliance virtual tape libraries (VTLs) and now backs up its other remote offices over the WAN to the libraries.

Grey Healthcare Group has four remote sites; one of the FalconStor VTL appliances is in its headquarters and one is in a remote office.

RUNNING "DARK"

Switching to another form of backup won't magically produce someone with the training to run remote-office backups. But most of the alternative technologies require much less operator intervention at the remote site, and in some cases none at all. Vendors and their customers are aiming for remote backup that runs "dark," without any human interven-

STILL A PLACE FOR TAPE

DESPITE SOME OF THE DRAWBACKS OF TAPE, it's still the most common method of backing up data at remote offices. The big reason, says Subodh Kulkarni, vice president of global commercial business at Imation Corp., a maker of tape as well as tape-, optical- and disk-based backup systems, is cost. "Tape continues to have the lowest cost per terabyte," he says. "A terabyte cartridge costs $50 to $70, much less than disk."

Although we tend to assume that remote-office staffs lack the skill or time to manage tape, that's not always true. "I may have someone who knows how to deal with tape," says David Hill, principal at the Mesabi Group, a consultancy that deals with backup issues.

And the assumption that remote offices have modest data storage needs doesn't always hold water. Remote sites such as engineering field offices or laboratories can produce very large data sets that have to be archived, something tape is often ideal for. Even if a site has modest storage needs, the data may need to be archived. "Dentists want to keep their X-rays for a long time, but they don't want to keep their business data for 20 years," says Hill. That tends to lead to a hybrid solution with tape as part of the mix.

Hill also suggests that backing up to a central repository may not be ideal. "I may not want to depend on a central site to restore," he says. "If something happens and I need to restore, I may not feel I can restore from the central site because of bandwidth reasons."

Some businesses may want to move their remote sites away from tape, but economic constraints come into play. "It also comes down to the timing of decisions," says Stephanie Balaouras, principal analyst at Forrester Research, a Cambridge, Mass.-based market research firm. "Corporate policy often says the useful life of IT equipment is five years. They have to use what they have for five years."

This is especially true in today's economy. Not everyone is satisfied with tape, but those who are don't see much point in changing. And tape can be extremely reliable. Drives can work flawlessly for years, so many companies see no reason to replace them. Kulkarni notes that Imation is still making cartridges for obsolete tape formats like Travan because there's still demand for them.

Kulkarni agrees that tape has some disadvantages. "On the negative side, tapes do need some knowledge, especially of the tape identification scheme," he says. "It's not quite an off-the-shelf product."

tion at all. Some of the replacement technologies, especially remote backups over the WAN or into the cloud, are dark technologies as far as the remote office is concerned.

Many backup vendors provide the ability to manage their products remotely, often from a Web-based console. This allows IT staff to monitor and control backups without intervention at the remote office.

10 TIPS FOR USING CLOUD BACKUP SERVICES

1. CHECK YOUR BANDWIDTH. You need to know how much data you expect to back up to the cloud service and if your current bandwidth is adequate not just to handle backups in a reasonable time, but for restores.

2. ENSURE RELIABILITY. A cloud backup service, like any online service, can experience outages. Check on the service's record, noting how many outages they've had and how long they've lasted.

3. TALLY THE COSTS. Because services have different fee structures, it's important to know how much data you'll ship to their site, how frequently you'll run backups and how often you expect to restore data. With that information in hand, you'll be able to make accurate cost comparisons.

4. EVALUATE ACCESS CONTROLS. You may want your users to be able to do their own restores, but access to backup data should be controllable to limit unnecessary backups/restores and to protect the data, especially if access from anywhere is allowed.

5. MAKE SURE YOUR DATA IS SAFE. Ask what measures the service provider takes to safeguard your data. They should have backup data centers and offer encryption for data in flight and at rest. If encryption is an option, get it and make sure it's turned on.

6. STOP AND RESUME. A cloud backup service should allow you to stop a backup in progress and then restart it from the point it was interrupted. Having to rerun an entire backup is costly and time consuming.

7. BIG RESTORES. If a disaster strikes and you have to restore your entire backup data set or a large part of it, online transmission will likely be impractical. Find out how the service handles these requirements.

8. PROTECT DESKTOP AND LAPTOP DATA. If all of your company's user data is stored on servers, you don't have to worry about desktop or laptop PCs. But if you have a mobile workforce or allow local storage, ask if the service provider can also protect the data on those systems.

9. AGENTS AND OTHER SOFTWARE. Many services require an agent to run on the servers you're backing up. Find out if the agent will affect the servers' performance or interfere with other applications, and if they can be managed centrally.

10. CONTINUOUS OR SCHEDULED BACKUPS. Some services can back up your servers and other systems continuously (or nearly), while others do backups on a regular schedule. Make sure your provider offers the types of services that best fit your company's environment.

"This is just a box hanging on your rack," says Ruffolo of Miron Construction's ExaGrid system. "You don't have to worry about it."

Restoring from backup is usually not quite as simple, but can still be managed from the central office in many cases.

Watkis says Grey Healthcare gained a number of advantages with its new backup system. "It has reduced the cost, the amount of system administration time, storage media requirements, the recovery timeline, our service warranty guidelines and improved security because the tapes are encrypted," he says. Deduplication wasn't part of the original specification, but it came with the FalconStor products and has proved to be a major advantage, adds Watkis.

The cost savings realized by switching to offsite VTLs can be substantial. "We were spending over $60,000 or $70,000 on tape media itself," says Watkis. "That's now down to $12,000 and maybe lower." But the company's savings went well beyond the cost of tape media. Staffing costs were cut, and they were even able to avoid some storage system purchases. "We were spending $5,000 to $7,000 a month for people to stay late; that's gone away," he says. "We're no longer purchasing any external storage for our SAN. Deduplication freed up more server space. We budgeted $60,000 for additional storage and that wasn't done."

"We were spending over $60,000 or $70,000 on tape media itself. That's now down to $12,000 and maybe lower."

—Chris Watkis, director of IT, Grey Healthcare Group Inc.

Miron Construction is in the process of switching to disk-based backup. It installed an ExaGrid system in its data center a year ago and intends to switch the remote offices to disk backup this year.

"The first thing you need to be certain about is to make sure you have a reliable WAN connection," says Watkis. "You need a dedicated switch connection into the internal network for the VTL. Don't try to share a connection."

Watkis says that most central offices already have these kinds of connectivity if they have remote sites. However, having enough bandwidth is vitally important for doing remote backup.

"The VTL method is new and I took a risk in implementing it," says Watkis. "We were really putting business at risk and the solution had to work with an immediate return on investment."

SPEED MATTERS

Another important remote backup issue is speed. Specifically, any remote backup planning has to consider how long it takes to do the actual data backup and how long it will take to restore data.

"We found out the hard way that the amount of data being stored matters," says Victor Liu, president at Link High Technologies Inc., a

Denville, N.J., reseller specializing in backup technologies. "On a pure Internet solution, more than 50 gigabytes is just not practical." Link High Technologies' Liu says one customer with a large amount of data stored over the Internet took more than four days to do a complete restore.

There are several ways to deal with this. One way is to keep several of the most recent backups on local disk and transmit them to the remote site. Another, which is available from some vendors such as Link High Technologies, is to have the data transferred to USB drives and overnight expressed to the customer needing the restore. With large quantities of data, the difference between reloading from a local disk and reloading over a network can be considerable, even if you take into account the time it takes to ship the disks.

ONLINE BACKUP ALTERNATIVES

One of the more intriguing backup alternatives for remote sites and smaller offices is cloud-based, or online, backup.

"I wouldn't say it's becoming popular yet," says Eric Burgener, a senior analyst and consultant at Taneja Group, Hopkinton, Mass., "but I think it will over time." Burgener notes that unlike the storage service providers (SSPs) of a few years ago who tried to provide main storage for corporate data, Storage-as-a-Service (SaaS) companies using the Internet "cloud" are concentrating on backing up and archiving data where there are fewer performance constraints.

If a remote office has a relatively modest amount of data to back up, a cloud service may be a good fit.

Many cloud providers, such as Carbonite Inc., are currently targeting individuals and very small businesses. But cloud backup is well suited for remote offices because it can handle dark backups—automatic backups that don't require manual intervention. If a remote office has a relatively modest amount of data to back up, a cloud service may be a good fit; however, bandwidth may still be an issue. At the least, a broadband connection is required, and consideration should be given to bandwidth requirements for large restores.

Reliability could also be an issue. There have been some well-publicized outages at large cloud providers, but a connection that's performing poorly or not working at all is more likely to be a problem.

Costs vary widely among cloud backup providers, as do the methods of calculating them. Some offer "unlimited" amounts of storage for a monthly fee. Others, such as Amazon's Simple Storage Service (S3), charge a few cents per gigabyte per month to store data, with additional charges to upload or download data. Other vendors have sign-up fees or minimum monthly charges.

A major advantage of backing up to the cloud is that, in most cases, you can recover to any computer. This restore-anywhere feature may sound a

bit risky, so most cloud services offer encryption as part of their packages.

Because most cloud backup services rely on the Internet to transfer data, there are practical size limits on how much data you can effectively back up. By its nature, cloud backups are slower than most LAN backups, although the speed depends almost entirely on the bandwidth of your connection. The other determining factor is the size of the backup window. With a fairly typical 6 Mbs connection speed and a weekend backup window, 1 TB or 2 TB of actual data transmitted may be pushing the limits. But most transmissions are far smaller, as differential or incremental backups only ship the changed data. The first backup, which is a full backup, takes far longer than subsequent backups.

And just because someone else is handling your backups doesn't mean you can forget about them. Analysts say it's important to trust your data to vendors who have the capability to protect it.

"Consider the class availability and disaster recovery capability," says Stephanie Balaouras, principal analyst at Forrester Research Inc., Cambridge, Mass. "Are they backing up to another site? If I were an enterprise-class company, I'd want to make sure my data was further protected."

Finally, consider the legal implications of where the data is being stored. "For European operations, you need to ask where the data center is located," says Balaouras. "Are you in violation of any national privacy laws by transmitting the data offsite?"

There are still plenty of things to like about tape, so it's unlikely it will disappear from remote offices completely. But it will undoubtedly become much less common as vendors develop replacement technologies at lower costs that enterprises can easily integrate into their existing backup infrastructure and operations. ⊙

Rick Cook specializes in writing about issues related to storage and storage management.

DEDUPE UPDATE: What's coming in 2009

Everybody knows that the hottest thing in storage in 2008 was data deduplication. Don't expect it to cool off in 2009.

By Dave Raffo

EVEN WITH ALL of the data deduplication product rollouts in 2008, we can expect plenty more throughout this year.

CommVault got the ball rolling when it launched Simpana 8 with block-level deduplication and dedupe for data on tape and disk in January. Quantum Corp. brushed up the integration and replication capabilities of its DXi7500 dedupe backup box around the same time. While there was plenty of news and new products in 2008, we don't expect the onslaught to let up—even more data deduplication products, including some for primary storage, are slated to roll out soon. The following is a roundup of dedupe products currently in the pipeline:

DELL INC. In November, Dell said it would bring out dedupe products in 2009 that will integrate with Quantum's dedupe products. Quantum also licenses its target deduplication software to Dell's storage partner EMC Corp., so it's expected that Dell's products will be compatible with EMC's Disk Library, if not an outright OEM deal. Dell said it would also sell Quantum's replication software to copy deduped data for disaster recovery.

EMC'S CELERRA. EMC will offer primary storage dedupe for file systems with the next upgrade of its Celerra NAS platform, due early this year. EMC will use single instancing from its Avamar dedupe product and compression from its RecoverPoint appliance. The Celerra dedupe will compress files with low usage activity, and single-instance files to remove duplicates. EMC's NAS rival NetApp began offering dedupe for primary storage a few years back, but no other major storage vendors have followed suit until now.

HEWLETT-PACKARD (HP) CO. There have been no public announcements, but David Rogers, manager of product marketing for HP StorageWorks data protection, tells SearchStorage.com that the company will add replication on the firmware of its dedupe products, the StorageWorks D2D Backup System and StorageWorks Virtual Library System (VLS). Rogers says the replication was developed specifically for the dedupe products and will be a licensed feature.

In addition to these new products, expect enhancements to the dedupe products already out. For instance, FalconStor Software Inc. added a NAS interface to its virtual tape library (VTL) dedupe product in early December.

NEXT DEDUPE FRONTIER: ARCHIVING

Stephen Foskett, director of data practice at storage consultancy Contoural Inc., says dedupe will become a necessary technology for backup and take big steps into archiving in 2009. But dedupe isn't yet ready for many types of primary storage, to the chagrin of some of the storage customers he talks to.

IN BRIEF: DEDUPE IN 2009

The following deduplication products have been unveiled or are expected to be rolled out this year.

Vendor info	What's notable	Availability
CommVault Simpana 8	Adds block dedupe and tape/disk dedupe	Available now
Dell Inc.	Dedupe systems based on Quantum technology	Not announced yet
EMC Corp. Celerra	File dedupe using single-instance storage and compression	Not announced yet
FalconStor Software Inc.	Added NAS interface to virtual tape library (VTL) dedupe product	Available now
Hewlett-Packard Co.	Plans to add replication to current dedupe products	Not announced yet
Quantum Corp. DXi7500	Enhances replication and integration of DXi7500 dedupe system	Available now
Riverbed Technology Inc. Atlas	Applies some Steelhead technology to primary storage dedupe	Expected mid-2009

"There's some general disappointment at this point that people can't use dedupe on primary storage, but they're excited about the potential it has on archiving," says Foskett. "I'll be shocked if every product in the archiving space doesn't have advanced deduplication pretty soon."

Frank Slootman, CEO at Data Domain Inc., agrees dedupe will become more of an archiving play this year, but says dedupe products will change more in size and scope than in capability. He says they'll get bigger and faster on the high end, and smaller and cheaper on the low end.

"The technology is still developing, and will always be developing," says Slootman. "We're riding a relentless wave of microprocessor improvements, mostly on the Intel side. And that kind of stuff is manna from heaven for us."

CommVault took the lead on one new development. Besides adding block-level deduplication, Simpana 8 became the first product to allow writes to physical tape libraries without requiring re-inflation of deduplicated data.

But vendors are working on primary dedupe, too. Riverbed Technology Inc. is preparing a primary dedupe product, although it has been pushed out until 2010. Riverbed began alpha testing its Atlas device in September, with the expectation that it would ship around the middle of this year. But testing showed the product needs more work to make it easier to install and manage, so Riverbed will wait until next year. Atlas will use the deduplication technology that Riverbed employs in its Steelhead WAN optimization products to shrink primary data. Atlas' closest competitor is NetApp's deduplication software for primary data.

Eric Burgener, a senior analyst and consultant at Hopkinton, Mass.-based Taneja Group, says Riverbed will also raise the bar with Atlas. "The scalability in a distributed environment is better than anything out there," he says. ⊙

Dave Raffo is the Senior News Director at SearchStorage.com.

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

Sponsored by:

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

Where does deduplication belong in backup?

Should you go with a software-based approach that allows for policy-based deduplication or a hardware-based approach because it can be implemented quickly and easily?

THE IMPACT THAT data growth is having on backup windows is driving more organizations to implement disk-to-disk backup. This has created tremendous interest in data deduplication because the capacity optimization resulting from deduplication means that data can be retained longer on disk, which increases the likelihood of a disk-based recovery vs. a slower, manual, tape-based recovery.

While deduplication has been a feature of several backup offerings for years, the technology has been most widely adopted in backup hardware, such as virtual tape libraries (VTLs) and network-attached storage (NAS)-based disk targets. Meanwhile, deduplication implementations in backup software require organizations to switch out legacy solutions, which the hardware-based deduplication vendors have made sure to point out isn't always a desirable path. Now that mainstream backup software vendors such as CommVault, EMC Corp., IBM Corp. and Symantec Corp. are incorporating data deduplication into their backup products (reducing the amount of disruption caused by implementing deduplication), the question is being asked again: Where does deduplication belong in backup?

> Organizations are just as likely to purchase and implement data deduplication technology from backup software vendors as they are from disk/appliance hardware vendors.

SOFTWARE-BASED DEDUPLICATION

Software-based approaches are differentiated in a few ways. First, they have knowledge about the data in the backup stream; they can look at patterns in the data stream (the bytes that make up a file) and determine the optimal segment boundaries, which maximizes the likelihood

of identifying duplicates. In short, backup software understands the content, whereas target-side deduplication solutions typically don't. Targets simply receive a "blob" of data from the backup application. Those target-side deduplication devices that are content-aware typically have to extract the meta data associated with the backup and "reverse engineer" the backup stream to understand its contents. Second, integration with the backup software allows for policy-based deduplication. Deduplication can be disabled for selected data sets where it doesn't make sense to turn it on (such as an MRI image) or for other data types (like databases) where you don't want to interfere with performance.

One of the drawbacks of a software-based approach is that adopting a deduplication feature could require an upgrade in backup application and/or client agents. Another factor is that deduplication may be processor-intensive and, when performed at the source application server, it may compete with and slow down apps. The scalability and performance of the media server performing deduplication could also be limiting factors. It will be important to investigate the upper limits of deduplication "pools" and performance capabilities for large volumes of data.

One of the drawbacks of a software-based approach is that adopting a deduplication feature could require an upgrade in backup application and/or client agents.

HARDWARE-BASED DEDUPLICATION

Hardware-based deduplication is less disruptive; that is, it's seamless to deploy because it's compatible with any backup software and can be implemented quickly and easily. It typically leverages powerful, purpose-built storage appliances to accommodate processing of the entire (non-deduplicated) backup load either pre- or post-ingestion. Hardware-based solutions also have the advantage of processing data streams from multiple backup applications.

There are a few trade-offs to consider. More data than may be necessary traverses the network between the source system and target device (creating unnecessary congestion), as deduplication happens at the end of the data path. Depending on the solution, scalability could be another drawback. Some vendors are limited to single-node systems, which can result in multiple islands of deduplication and points of management, as well as underutilization in capacity per silo. Data streamed to a single-node system is only compared with other data directed to the node.

The goal of many target-side deduplication vendors is to deduplicate across clustered nodes. Global dedupe allows backup data to be dedu-

plicated against all other backup data, regardless of which head actually receives the data. This capability is seen more often in software-based and grid architecture approaches, but may also be supported for target deduplication systems that replicate in a hub-and-spoke fashion (with global deduplication occurring at the hub). Global deduplication can result in higher deduplication ratios—as data is deduplicated within and across backup sources—and greater economies of scale with respect to operational overhead and capital costs.

COST IS A FACTOR

Enterprise Strategy Group research has found that organizations are just as likely to purchase and implement data deduplication technology from backup software vendors as they are from disk/appliance hardware vendors. The top considerations when evaluating and selecting a data deduplication provider are cost, ease of integration, performance, ease of use and scalability, with cost clearly outranking the others. Now that deduplication is becoming a mainstream feature integrated in backup software, it will be interesting to see if "bolt on" deduplication systems can maintain their premium price.

As with any new technology, it will be important for IT organizations to evaluate software- and hardware-based approaches vs. the requirements of the environment. Having a clear understanding of how deduplication works, especially in conjunction with other requirements such as performance, ease of use and offsite copy creation, should go a long way toward selecting and designing a solution that delivers maximum business, operational and financial benefits. ⊙

Lauren Whitehouse is an analyst focusing on backup and recovery software and replication solutions at Enterprise Strategy Group, Milford, Mass.

More testing, more confidence for DR plans

Disaster recovery (DR) planning is a work in progress and this month's Snapshot survey results bear that out. More respondents are testing their DR plans regularly, but overall it's still only 59%. Non-testers cite lack of a DR site, inadequate staffing or lack of funds for not testing. But a tough economy may be contributing, as staff and money issues rose by 7 points and 5 points, respectively. Staffing issues are also affecting DR site management—last year, 48% of respondents said their own staff ran their DR sites vs. 27% this year. Still, 38% are very confident that their DR plan can avert significant business impact vs. 32% in 2008. That's bolstered by the 49% who have met their testing RTOs/RPOs. But for the sheer number of applications successfully recovered, the numbers are less encouraging, with 35% claiming to recover all applications vs. 57% last year. *—Rich Castagna*

How often do you test your DR plan?

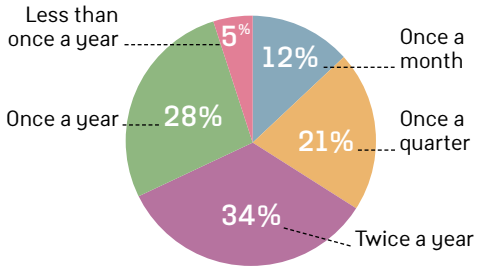

Did you meet your RPOs and RTOs during your last test?

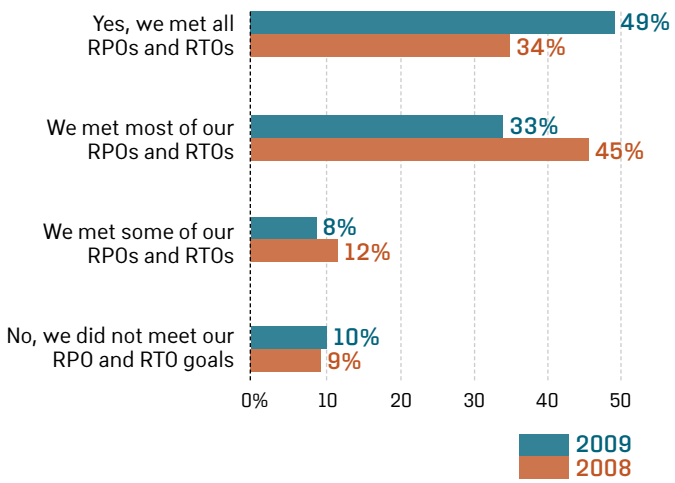

41%

Respondents who don't perform regular disaster recovery plan testing.

How far is your DR site from your main data center?

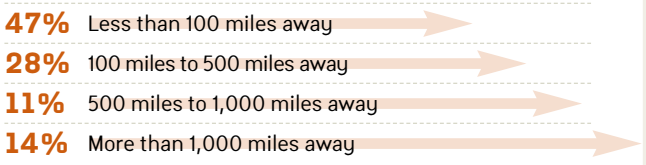

> "We do a DR test of each of the individual applications before they go into production, but doing the full suite of apps for all business units simultaneously has always been cost prohibitive." *—Survey respondent*

Check out the following resources from our sponsors:

BakBone Software Inc., page 6

BakBone Software

Reducing Risks Through Flexible and Automatic Oracle Data Protection

Maxell Corp., page 8

Is Your Information at Risk?

Quantum Corp., page 3

Quantum Deduplication Solutions

White Paper: Quantum DXi7500 Deduplicated Backup for the Enterprise