

The Storage Management Imperative

Customers Need Strategy Now More Than Ever

By Graeme Thickins
March 2003



TidalWire

A networkengines company

The Storage Connection.SM

www.tidalwire.com

With IT budgets stagnant over the past two years and data growth continuing to advance at a 40-100% compound annual growth rate (depending on whose number you believe), it doesn't take a genius to see something big is looming.

The term "train wreck" comes to mind.

Overly dramatic? Perhaps so – but industry experts say one fact is clear: the enterprise without a strategy for getting its arms around the storage management problem is living very dangerously indeed.

Welcome to the developing new world of awareness for storage management, and the long talked about software category of "storage resource management" (SRM). Though the term is far from new, some are saying 2003 just may be its year.

"The big problem in storage shifted from capacity to management about 24 months ago," says Richard Scannell, VP of strategy for Glasshouse Technologies. "What was interesting then is pretty much passe now." He notes the focus today must be on how an enterprise can deal with the *complexity* of its storage. "There are lots of issues that come into play now besides the technical piece."

Another complication, Scannell points out, is that storage-industry leaders (vendors included) are well ahead of the market – they're now looking forward to utility-class products. While he notes those new technologies hold much benefit, the rest of the market has a "competency defect" and is running behind. The gap is glaring, he thinks – maybe 18 to 24 months wide. "The complexity is such that the customer simply doesn't know where to begin," Scannell says.

MANY-HEADED MONSTER

Actually, a confluence of factors is driving the need for storage management, even beyond the obvious dichotomy of rapid data growth and reduced IT budgets.

"The lack of standards is causing interoperability challenges and cross-training nightmares," says Paul Mayer, product manager for storage resource management at Datalink – another vendor-neutral storage consulting firm, based in suburban Minneapolis. "And the cost to administer storage has risen to between \$3.50 and \$13.00 per gigabyte per year. SNIA says the former, Yankee Group the latter – you pick." Then do the math for your own storage, and it likely won't be pretty.

Another stat one often hears cited is how many dollars are being spent today on storage management compared with each dollar spent on storage infrastructure. Various research pegs that factor at 3x to 8x.

Of the three components making up the enterprise storage budget – hardware, software, and labor – Gartner says labor is now the only one increasing as a percentage of the total budget.



And that rate of increase continues to accelerate, its research has found, even as hardware and software costs are dropping.

A MATTER OF DEFINITION

It seems the categorization of storage management products is less than a perfect science – often being in the eye of the beholder. And that can add to the confusion. Gartner defines SRM products as those "providing data collection and automation agents to consolidate and operate on information from multiple platforms supporting storage management tools on multiple OS, storage, and SAN devices." Got that?

Datalink takes a more palatable approach to describing the functions of SRM:

- Global storage resource discovery at the file level, reaching beyond SAN and into DAS and NAS
- Focus is more on *the data* than the storage device
- Trend analysis
- Projection capabilities
- Alert/notification (email or SNMP)
- Extensive reporting
- Provisioning
- Charge-back capability

Gartner also allows that key functions of an SRM product include "capacity reporting/analysis, performance reporting/analysis, capacity/performance management automation, storage management product integration, application and database integration, and hardware integration."

Another key requirement for an SRM product, says Gartner, is that it integrate with device resource management products and media management products to include launch of hardware configuration utilities from the SRM console, collection/reporting of agent information, and integration of logical level data. It now considers products that provide for discovery, topology mapping, and monitoring of SAN components (such as the products we mention below) to be part of the "SRM subsegment," because many are being included with SRM software, or are evolving to include SRM functions.

The whole question of whose definition is best may be moot, however – at least at this point. Glasshouse Technologies' Scannell cites a survey his firm did recently of storage professionals. "About 60% of them had never ever *heard* of SRM," he says.

As we referred to earlier, it's hard to deny that the vendors are ahead of the market – way ahead.



WHAT'S REAL, RIGHT NOW?

SANavigator 3.5 from McDATA is one example of a mature, well-featured software product for more efficient storage management. It's a network management application that promises better management and control of heterogeneous, multi-vendor storage networks. Think of it as a "SAN fabric management" product.

"We're really about the network - that's where we're strong. And SANavigator reflects that," says Matt Vawter, senior product manager at McDATA. "SANavigator gathers information from the network and allows users to gain greater control over their fabrics so they can decrease the complexity of the storage networks."

SANavigator provides capacity planning, device discovery, performance monitoring, and asset management tools to allow for better management of storage networking assets and optimization of resources. It also features policy-based automation tools, which McDATA says makes for greater operational efficiencies, increased ROI, and lowered TCO for storage networking.

"It's a single interface, a single pane of glass, to manage all the layers from the HBA down through the network into the storage," says Vawter. "This allows our users to simplify their storage networking, optimize their resources and minimize their storage networking risks."

CLIMBING UP THE PYRAMID

Datalink, which defines itself as a storage architect, tells its customers to think of storage management as a series of layers in a pyramid, going from basic foundation steps at the bottom to advanced (and still evolving) technologies as you go up the pyramid.

Customers often start simply with storage redundancy, Datalink's Mayer says. "That's okay – it buys more time. But it creates more problems down the line, more things to manage."

Datalink says it's helpful to think of the process or evolution of storage management as follows, from lowest level to highest:

- 1) Storage consolidation and device-level management
- 2) SAN fabric management
- 3) Virtualization
- 4) Storage resource management (SRM)
- 5) Automated resource management (ARM)

Vendors' products don't necessarily fit neatly into each level, Mayer says, but each step of the pyramid nonetheless represents distinctly different concepts and technologies, which do interrelate with one another.

Datalink is seeing the categories of "SAN fabric management" and "storage resource management (SRM)" come together, which would appear to jibe with Gartner's move last year to begin combining them for its category-tracking purposes. Nonetheless, Datalink's steps-up-the-



pyramid view of storage management is a helpful way to understand the levels of sophistication, and provide a path for those just getting storage management religion.

Of McDATA's SANavigator product, Datalink's Mayer says this: "It's very simple to install, configure, and use, and it offers real simplicity for the common features of discovery, zoning, monitoring, and reporting. And McDATA has a vision for an agentless, feature-rich version of their product, which will be interesting to follow." (McDATA is not commenting further on that vision at this point. But more on the issue of "agent vs. agentless discovery" in SAN management later.)

How does an experienced, hands-on SAN troubleshooter view the matter? "In theory, all these products should be able to manage all devices on the SAN. But, practically speaking, without the latest device-specific APIs, managing from a SAN management tool will always be one step behind the device vendor's management tool for that device," says Shravan Pargal, director of storage consulting at Compellent Technologies, a storage technology startup in Eden Prairie, MN. "Why is this important? Because, when the chips are down and I need to look at diagnostic or error logs of a device to debug a SAN problem, I'm forced to go to the device vendor's interface to debug that device. In other words, most of the time I can't debug the SAN using a high-level SAN management tool - rather, I have to use the device vendor's tool."

VENDORS APLENTY

End users looking for storage management solutions – whether SAN fabric management, virtualization, SRM, or even the new, whiz-bang "ARM" – have no shortage of companies to consult. A great online resource to start the process, and even compare features, is <http://www.storagemountain.com/software-SAN.html>. It lists some 28 vendors in the "storage management software" category alone. This web site was first launched by storage consultant W. Curtis Preston as BackupCentral.com. Preston has also authored a useful new book called *"Using SANs and NAS: Help for Storage Administrators,"* published by O'Reilly & Associates (2002).

Datalink regularly evaluates many storage management products. Who are some of the other leaders? "Veritas SANPoint Control is one to watch," says product manager Mayer. "It's a SAN management product that, in my opinion, has crossed over into the SRM world. Veritas has a solid vision for SANPoint, and their recent acquisition of Precise would appear to put them in the driver's seat to execute on that."

Just what is that vision of Veritas' for SANpoint Control? "My interpretation of it," says Mayer, "is that they want to deliver complete management and control at every layer of the storage management hierarchy – including device management, fabric management, storage resource management, and automated resource management. Veritas has taken big strides in the SRM space, and has an architecture to deliver the self-healing and automation provisioning functionality that people expect from ARM. Of course, they aren't delivering true ARM yet, but neither is anyone else."

Another vendor to watch? "Computer Associates (CA) also has good story with its BrightStor SRM product," notes Mayer. "But they came to the game a little late, and seem to be in catch-up mode. They're devoting a lot of resources to the product, and could be a contender soon."



InterSAN is also a firm whose name keeps cropping up on the radar of the vigilant storage-market observer. As this white paper went to deadline, the company was about to release version 2.5 of its Pathline storage management software, which seems to move it more into the SRM space. The new version is said to contain improved planning, monitoring, reporting, chargeback, and automation features. The company says many users are having to reconfigure their SANs often because of host bus adapter upgrades and other adjustments. Pathline now enables such companies to automatically re-map their SAN storage in these upgrade or modification situations. InterSAN has also added automated path reconfiguration in Pathline v2.5, automating the re-mapping of paths from applications to data storage paths, a requirement after SAN reconfiguration. In addition, the Pathline upgrade includes new storage resource management capabilities, policy compliance auditing, path proofing, subscription-based event management, and chargeback reports.

"SANs can start looking like spaghetti," says Anne Skamarock, senior analyst at Enterprise Management Associates. "They're great till they grow, then they become complex and difficult to manage. No one task is difficult in and of itself – but managing the spaghetti is no longer possible without computer assistance to automate the auditing and provisioning tasks."

STRATEGIES FOR SAN FABRIC MANAGEMENT

Getting back to a more basic level, relating to choosing a product to help you manage your SAN fabric right now, there are several things to be aware of in the way of what product strategy is right for you, including:

- Agent vs. agentless discovery
- Multiple levels of access and control
- In-band vs. out-of-band management
- API-based device management

Each of these is a discussion in itself, but let's just take the first one for now, a topic every end user should be aware of.

In technical circles, there's much debate over the use of agents as a means of reporting in a storage management product. "The debate mainly focuses on the balance between how much intrusiveness and performance impact an administrator can tolerate," says Mayer, "versus how much information they need in order to optimally manage their storage environment.

"We're still trying to get our arms around this issue ourselves," Mayer notes. "Even the vendors don't seem to have a solid grasp on it, so us consultants and integrators – not to speak of the end users – are left to our own devices to navigate this maze."

First, understand that storage management products generally gather information through two primary paths. The first is from within SAN, in which the software probes through fibre channel ports to receive its information; the other is via a LAN connection to a storage host (server), where it probes through the Ethernet card to obtain its information. Given that, Mayer says, there are some variables when choosing whether an agent-based or agentless approach makes more sense.



NOT-SO-SECRET AGENT, MAN

"In agent-based reporting, the software product includes what is, hopefully, a *small* piece of software that runs on each server that's under management of the product," Mayer says. The agent runs in the background on the server, collecting information about the storage that is managed by that server. It also can collect information about applications running on the server, and report on storage statistics pertaining to the application. "For instance," says Mayer, "this agent could track where Oracle's redo logs are physically located, or track trends on storage growth for particular Oracle tables."

Agent-based reporting can report back to the storage management application either via SAN or LAN. In general, this approach has the potential to yield the most information of all the storage monitoring approaches, according to Datalink's Mayer, depending upon the extent to which the software vendor exploits this capability. "The impact on the server depends on how much memory is consumed by the agent while it is running, as well as what processes are involved in gathering the information," he says. Another variable is how frequently the agent is configured to poll the host's resources for changes, and how frequently it reports back to the storage management server. If this is configurable and is set to run too frequently, Mayer says "it can bring a server to its knees."

Agentless reporting in a SAN is an approach that captures information on a data path from the fibre channel card (HBA) to the port on the switch, to the port on a storage device. "This can give you information such as firmware levels on all of these devices, traffic data, and can signal problems with any of these devices," says Mayer. "It does not, however, allow much information to be gathered about data allocation in the environment. It's a quick way to get some graphical visualization and a resource audit of the SAN fabric, but does not offer much in the way of useful, storage-related data. On the other hand, it's very non-intrusive and not resource-intensive."

Okay, got that about agent vs. agentless discovery? Good, but you're just getting started. "We suggest customers take calculated steps, based on sound ROI analysis, as they move up the storage management pyramid," says Mayer. "There is no end-to-end solution – only 'Legos' that must be used under adult supervision. "

AVOIDING YOUR OWN TRAIN WRECK

So, if you're an end user, where are you on this continuum of storage management? Baby-faced rookie? Hardened storage veteran? Or somewhere in between?

Wherever you are, you have an interesting future in front of you, as data growth advances perilously, and technology marches on to control and manage it. "Being a storage professional today – and especially in the future – is definitely *not* a part-time job," says Glasshouse's Richard Scannell.

The challenge, quite simply, begins with keeping up, then expands in no small way to sorting through all the vendor claims and hype.



One storage analyst puts an interesting spin on things. "The customers aren't the problem. It's the vendors who haven't done a very good job, " says Enterprise Management Associates' Anne Skamarock. "They need to (1) integrate broader capabilities to provide real value, as opposed to selling point-products or solutions today; (2) educate the market; and (3) sell management software with their storage, to eliminate the second decision to purchase management. Most customers have used, and will continue to use, the management provided with the storage itself." Maybe this line of thinking is part of the rationale behind the latest spate of acquisitions by the big players, to broaden their offerings.

"The 'train wreck' may not be so much for the customers as it is for the vendors," adds Skamarock. "As storage becomes less expensive, the opportunity presents itself for vendors to retain prices by adding management functionality. The more integrated the management functionality across the infrastructure, the higher the value to the customer, at all levels of storage. And customers will go with the vendor that continues to provide the most value, all the way up the chain."

ABOUT THE AUTHOR

Graeme Thickers, who has been reporting for Tidalwire since early 2002, is a freelance technology writer and marketing consultant based in the Twin Cities and Southern California. He began working in the storage industry more than 25 years ago and today, among other things, provides news and perspective on storage networking, including coverage of industry events. His reporting and commentary have also appeared in other online journals and newsletters, including Conferenza, Business 2.0, Office.com, The Journal of the Hyperlinked Organization (JOHO), NanoTechNews.com, and Darwin Magazine. He may be reached at graeme@thickers.com.

