

White Paper

April 1, 2002



User Survival Guide to Network Storage Management



SNIA
Storage Networking Industry Association

User Survival Guide

to

Network Storage Management

Forget all of the marketing hype about networked storage and the silver bullet it provides in terms of storage cost reduction. The simple fact is that without effective management, networked storage topologies offer no more cost savings than the traditional server-attached topologies they seek to replace.

Without management, networked storage technologies like network attached storage (NAS) and storage area networks (SANs) may actually contribute to the total cost of ownership (TCO) for a company's storage infrastructure. While these technologies

may indeed enable storage capacity to scale non-disruptively in the face of burgeoning data, more often than not, increased capacity translates simply

to a need to hire more staff to administer and "groom" the storage platforms. Moreover, in an unmanaged

networked storage environment, the increased accessibility of data afforded by the topology may well translate to increased opportunity for data corruption and increased need for administrator oversight and intervention.

When presented with projections from Framingham, Mass.-based market research firm International Data Corp. of an anticipated 67% to 68% increase in networked storage technology adoption through 2004, longtime industry watcher and Disk/Trend President James Porter observes that NAS and SAN would not displace, but would more likely co-exist with server-attached storage and internal server disk for at least a

Abstract

As users strive to control their exponentially expanding data, the need for heterogeneous storage management is more important than ever for the storage networking industry. This situation has created both an opportunity and a challenge for vendors that must work together on the creation of standards if they are to ultimately provide the full value of their products. This Computerworld White Paper, written by noted writer and industry consultant Jon William Toigo, takes a look at current data management solutions, suggests interim user strategies, and lays the groundwork for long-term implementations. It also features practical case studies.

decade. "It isn't as though companies fielding networked storage solutions are ripping out their existing investment in storage technology," Porter says.

By **Jon William Toigo,**
Independent
Consultant and
Author

Porter's observation alluded to what has become a nagging industry truth today: networked storage acquisitions have not simplified storage management, but rather have added greater complexity – and cost – by multiplying the number and type of “targets” that the typical enterprise IT organization must manage.

Storage is the last bastion of enterprise information technology that lacks either open or de facto management standards. In the absence of standards that can be used to manage the heterogeneous components and processes that comprise a typical storage infrastructure, companies continue to struggle with the costly challenge of managing their mission-critical data armed only with the 21st Century's equivalent of stone knives and throwing sticks.

Is a solution to this problem waiting in the wings? The answer is a highly qualified “Maybe.”

The current state of storage management tools

Storage administrators often use the metaphor of a quiver of arrows to describe the tools they use to perform their work. They are referring to an assortment of “point management” software products, each of which is designed to perform a certain discrete task. These point solutions are often limited to a single vendor's storage component.

Point management tools have proliferated over the past several years to address a host of specific tasks – from device discovery to disk tuning to backup-and-restore functions. For the most part, storage administrators

have kludged together what they feel are best-of-breed products in each functional area; they then carry these tools with them as they travel from one server or storage platform to the next.

According to administrators, there are problems with the point-management-tool approach. For one thing, there are so many products that it's difficult to maintain a current knowledge of the capabilities and limita-

Forget the marketing hype about networked storage and the silver bullet it provides in terms of cost reduction. Without effective management, networked storage topologies offer no cost savings over the traditional server-attached topologies they seek to replace.

tions of each, or to ensure that the best tools for the job are identified and used. For another, using a multiplicity of products carries a potentially huge cost in terms of licensing, version management, and knowledge transfer.

A glimmer of hope was seen two to three years ago, when some of the larger storage management software

vendors began gobbling up some of the smaller fry and repackaging the point products as storage resource management (SRM) software suites. However, the results of this acquisitions frenzy, according to many storage managers, have been somewhat disappointing. According to corporate IT managers, many of the SRM suites introduced to market, despite the claims of their marketing brochures, are poorly integrated and include tools that are either not best-in-class or are simply irrelevant within the administrator's particular storage product mix.

Ideally, users say, storage management software would enable the policy-based automation of management tasks. That way, fewer administrators would be required to manage a growing storage pool. Automated management has been the goal of another class of storage management product vendors, those offering “management frameworks.”

Framework products use agent-based software to monitor storage infrastructure components, and also to return status or event messages to a centralized monitoring console. Ultimately, the goal of most framework vendors is to have their products do more than simply monitor for events and alarms. These vendors want to enable their products to perform most storage management tasks as well – with or without the participation of the console operator. At a minimum, these vendors want users to be able to launch point tools or SRM suites from the console, effectively centralizing the management task. Ideally, they want to instrument

storage so that a given action will be initiated automatically in software when a particular “out of threshold” event or constellation of events is detected.

When such a centralized and intelligent management approach will come to market is anyone’s guess. Currently, the bulk of the programming resources within most framework vendor development shops are dedicated full-time to writing device drivers, according to industry insiders. They must ensure that their product maintains its capability to monitor the myriad new releases of host-bus adapters, array controllers, and SAN switches and hubs that are introduced by storage vendors each year.

Moreover, they must spend substantial money, time, and effort in order to cultivate good relations with storage hardware vendors so that they’ll continue to enjoy access to application programming interfaces (APIs), command line interfaces, and self-articulated Web pages with which the vendors instrument their platforms for management. Under these circumstances, finding additional programming resources to write automation routines is a major challenge.

Finally in this taxonomy of storage management products, it’s worth mentioning that a new category of management has been introduced very recently – in the past several months. This new class of storage management software seeks to manage storage as a service, complete with service-level agreement (SLA) monitoring and reporting, as well as

back-end links to customer relationship management (CRM) and accounting systems.

Vendors in this space tend to be the early storage service providers (SSPs) that appeared in the late 1990s, but found little market for the concept of delivering storage as a service. Prospective clients were less interested in these vendors’ services than they were in the software the vendors had developed to manage, measure,



and bill clients. The SSPs have mutated into software vendors and are bringing their storage SLA reporting tools to market, creating yet another tier of storage management software.

The crux of the matter

Though it is rarely perceived, the core problem confronting storage

administrators in their quest for effective storage management is the industry’s resistance to the development of a common management standard. Some support the Simple Network Management Protocol (SNMP), which is used to monitor network devices such as LAN switches and routers. But this is by no means a universal characteristic.

For the most part, the products of a particular storage vendor are instrumented for management using proprietary software tools. Most vendors leverage their proprietary management schemes to sell more of their own products. The message from each vendor is that tremendous TCO advantages can be realized by the IT shop that buys that vendor’s products alone.

In organizations that decide to standardize on the platform of a single vendor and to field a homogeneous storage infrastructure, storage TCO improvements based on reduced labor costs are, in fact, possible. However, in most enterprises, the heterogeneity of the storage infrastructure can be discerned by a quick glance at the multi-colored racks and cabinets in the data center and in business unit equipment rooms. Often, homogeneous storage translates to a vendor lock-in – an uncomfortable proposition to many IT managers, CIOs, and CTOs, especially given the lack of a clear industry leader in the storage market. The old dictum “no one ever gets fired for buying IBM” has found no equivalent in the storage space.

In many shops, heterogeneous storage is guaranteed by the “tactical”



– a polite way of saying “knee-jerk” – manner in which storage products are acquired. Unaided by effective storage management tools, data is often poorly groomed, capacity utilization is often poorly tracked, and proactive or strategic product acquisition decision-making falls prey to reactive seat-of-the-pants purchasing in response to “Disk Full” error messages.

The resulting infrastructure is hard to manage without a common management standard. However, vendors lack motivation to commoditize their storage offerings by cooperating meaningfully with any standards-making effort, especially given a lack of a compelling business reason for doing so. IT organizations just aren't clamoring for a common storage management solution, according to one

vendor spokesperson. Until they do, the Balkanized world of storage management will remain much as it is today.

Coping with disarray

There have indeed been multiple industry initiatives aimed at introducing storage management standards. One that has received lip service (but little action by vendors) is the Common Information Model (CIM) effort spearheaded by the Storage Networking Industry Association (SNIA). Championed by Mark Carlson, an SNIA director who works on storage management tools at Sun Microsystems Inc. for a living, CIM seeks to describe all storage components as object. CIM proponents say this is a first step in instrumenting heterogeneous storage for common management.

CIM has reared its head before, in the desktop management arena. There it encountered some of the same results as current storage CIM efforts. Vendors agreed that management was needed, but rebelled against the notion of any one-size-fits-all standard that could capture the unique features and functions that differentiated their products from those of competitors.

Operating system vendors, including Microsoft Corp. and Sun, have instrumented their products for CIM management, so storage that is still “owned” by servers can be monitored and managed to some extent by CIM. However, the thrust of the networked storage effort is to divorce storage from direct attachment with servers, and this means storage product ven-

dors must implement CIM on their own platforms. To date, no vendors have done so.

The challenge to CIM comes in the form of EMC Corp.'s AutoIS initiative. Last fall, the storage vendor began to introduce a set of technologies, built on its proprietary products, that are intended to enable cross-platform storage management. Many people who are involved in the CIM effort perceived the AutoIS announcement as a shot across their bow – an effort by EMC to impose a de facto standard (a la Microsoft in the desktop world) on the management of storage.

EMC proposes an API middleware layer, called WideSky, that vendors may elect to use to interface their products to EMC's proprietary management software console,

Point management tools have proliferated to address specific tasks. Most storage administrators have kludged together what they feel are best-of-breed products in each functional area; they then carry these tools with them as they travel from one server or storage platform to the next.

In most enterprises, the heterogeneity of the infrastructure can be discerned by a quick glance at the multi-colored racks in the data center. Too often, homogeneous storage translates to vendor lock-in, an uncomfortable proposition to many IT managers given the lack of a clear industry leader.

ControlCenter/Open Edition. What gives legs to this approach is EMC's prominent position in the market; smaller storage equipment or software vendors seeking to sell product into an EMC account may find it useful to be able to claim that their technology can be managed by the management software already used by the customer.

The idea has attracted the interest of many storage players, including Compaq Computer Corp., on the one hand. On the other, it has also created a furious cry of "Foul!" from many of EMC's leading storage competitors. These detractors, including Hitachi Data Systems and IBM, are now rallying around SNIA and CIM, which they view as open alternatives to EMC's "Open Edition."

In the final analysis, the EMC ini-

tiative may bring about a common management standard – whether a de facto standard imposed by an industry leader, or an open standard reluctantly agreed to by storage vendors unwilling to give EMC bragging rights for having solved the storage dilemma.

Until a standard does arrive, enterprises must cope with burgeoning data growth and shrinking IT budgets using a kludge of point, SRM, and framework management products. Networked storage, SANs in particular, are not making things any easier.

Current-generation SANs, which are based on Fibre Channel technology, do not provide IT organizations with the means to support in-band management, or "management-in-the-wire." Such functionality was deliberately excluded by the people who invented Fibre Channel at IBM about a decade ago in order to preserve the maximum amount of bandwidth for moving data and SCSI commands from server to storage platform. This is an issue the Fibre Channel Industry Association is

working to resolve by adding "Fabric Services" into the protocol.

At the same time, two camps of vendors – Cisco Systems Inc. and Adaptec Inc. on the one hand, and Hitachi Data Systems, Nishan Corp., and Alacritech Inc. on the other – announced in January that they will introduce a new IP-based version of SAN technology to market before year's end. These new IP SANs, which will use an early specification for iSCSI until the Internet Engineering Task Force completes work on a final version of the standard, will hit the streets about two years earlier than analysts predicted. Because these SANs are based on a protocol that uses IP networking, they will also be manageable via the same SNMP that has been used for years in the LAN space.

It remains to be seen which protocols and management schemes will come to the fore as industry standards. Until a leader is established, the cost of ownership of storage will continue to accelerate in all but the most homogeneous storage infrastructures.*

COMPUTERWORLD
CUSTOM PUBLISHING

This White Paper was created by Computerworld's Custom Publishing group. Comments on this supplement can be sent to Carolyn Medeiros at carolyn_medeiros@computerworld.com.

This White Paper, as well as other custom supplements, can be viewed online at www.computerworld.com.

Bank Goes Global, Finds 24/7 Solution

CA's BrightStor EB helps financial institution meet new demands

Founded in 1979, Hemisphere National Bank (HNB) began as a local financial institution dedicated to serving the needs of the South Florida community. Building steadily on its reputation for customer service and commitment, HNB was soon ready to expand beyond its original Miami branch; the bank opened additional Miami branches and expanded to Coral Gables and Aventura, Fla. HNB quickly diversified and began to offer a wide array of individual and business services, from Internet banking to commercial lending and credit card processing.

Managing data growth

In 1998, the management of HNB decided to broaden its services again to reach an international clientele. With the establishment of its International Division, HNB added a series of new offerings, including pre- and post-export trade financing, that doubled the company's growth. The breadth of these new products and the exponential increase in data that they generated made it necessary for HNB to dramatically increase its storage capabilities as well.

The IT management team at HNB decided to deploy a network-attached storage (NAS) system to meet this challenge, and evaluated numerous backup-and-restore solutions. The team determined that only Computer Associates International

Inc.'s BrightStor Enterprise Backup (EB) was robust enough to handle the storage management tasks for its global operations. Working with Aspedient Technologies, a CA-certified managed service provider, HNB implemented BrightStor EB in less than three weeks.

Addressing e-business demands

HNB viewed the ability to back up and restore growing volumes of data within a shrinking backup window as its most critical objective. "Providing high data availability for HNB's expanding network of international clients was essential," says Tony Silva, co-CEO at Aspedient Technologies. "The unsurpassed backup speeds of BrightStor EB have enabled HNB to offer its global customers unimpeded access to data

The bank viewed the ability to back up and restore growing volumes of data in a shrinking backup window as critical. Providing high availability for a network of international clients was deemed essential.

and e-banking services 24 hours a day, 365 days a year."

Maximum uptime has become increasingly critical as HNB continues to expand its e-banking services. In 2001, HNB acquired the assets of a competing bank with a global customer base, extending the demands on its storage infrastructure across multiple time zones. "Hemisphere National Bank has become an institution without borders, and this makes performance and data availability key to our long-term success," says Miguel Morera, HNB's CIO. "The ability of BrightStor EB to deliver a clean, central solution for high-speed backup has enabled us to take on these new 24/7 markets with confidence."

The speed and reliability of BrightStor EB aren't the only benefits HNB is leveraging to further its expansion into new markets. "Its centralized administration and reporting capabilities enable us to proactively manage the backup-and-restore process, reducing the burden on personnel," Silva says.

From the BrightStor console, members of the IT team can view reports about completed backup jobs, monitor the performance of HNB's storage servers, and evaluate the disk capacity of its DASDs. The IT group at the bank says BrightStor EB is a significant part of their successful growth.

Making News with NAS

USA Today chooses Network Appliance solution for redundancy, ease of use

Darrell Griffith is senior infrastructure architect within USA Today's Infrastructure Technology Group. He played a critical role in the newspaper's implementation of Network Appliance Inc.'s filer technology and continues to monitor the NetApp solution as it is extended and expanded. Griffith recently spoke with Bruce Hoard, executive editor of Storage Networking World Online (www.snwonline.com), about the challenges and benefits associated with network attached storage (NAS), data replication, and multi-protocol support.

Q: Describe the storage problems you had before implementing your Network Appliance solution.

Griffith: We had no comprehensive, large enterprise-grade storage solution on the landscape here when we began a serious technology investigation toward the end of 2000. What we did have was a lot of what we call single-channel storage solutions and vertical applications. So we had to find a comprehensive enterprise-grade solution that worked well with our existing infrastructure and complemented our core competencies from both technology and skill-set perspectives.

Q: Describe the existing infrastructure you had to work with.

Griffith: We've got a lot of Wintel type stuff – our Windows 2000 and XP server base. We also have our share of AIX and Solaris. We have no AS/400, midrange, or big iron. Most of what we have is Unix or a standard Microsoft-type environment. We have some pretty powerful RS/6000 and AIX-based implementations for a lot of our publishing systems.

Q: What process did you follow in selecting a storage vendor, and how long did it take?

Griffith: It was actually straightforward. Three or four of us were cobbled into an ad hoc task force during the middle to end of 2000, and we basically set about reviewing the industry. First and foremost, we assessed our current architectures and tried to understand not just the hosting platforms themselves, but also the vertical applications.

We could find certain fits that would work with RS/6000s, but those machines may have been running a third-party application that wouldn't accommodate certain types of storage solutions. So we spoke with a bunch of different vendors and did informal research.

We also attended a very beneficial Storage Networking World conference in Orlando that winter. That's really where, after we'd culled all our preliminary information, we got a much better feel for industry directions. One of the big issues I had with taking a particular SAN approach was that they kept wanting me to reinvent something that was pre-existing, that being the network.

Q: What happened after the Storage Networking World show?

Griffith: We basically came away looking at a couple of big players, including Network Appliance, in the NAS market. We compared them to a bunch of other NAS competitors. Our decision to go with Network Appliance was really more philosophical than it was based on cutting-edge technology. NetApp seemed to be the best entrenched, and they seemed to be the company with the most proven commitment through their R&D and their technology. They had a lot of strategic initiatives with Microsoft as well as a clear commitment to the Unix standards, which was good for us. It was a case of looking at a business partnership for the long haul in terms of robust technology and where it was going.

Q: What Network Appliance product or products did you buy?

Griffith: We ended up going with the filers, and in particular the cluster filers. Media and data availability was one of our strictest criteria, so we bought a pair of F840 clusters about a year and a half ago. We put one at our nationwide headquarters in McLean, Va., and the second over a very redundant WAN link at a disaster recovery location at our White Oak, Md., location in the event of a failure in McLean. Speed was important and feature sets were important – but the bottom line is that data has to be available.

Q: What was your vision for how the products would be implemented, and the resulting benefits?

Griffith: We tinkered with a lower end solution first. We were still fairly cautious and experimental with NetApp. We subsequently brought a couple of their F740 filers and tested them with critical concepts, including SnapCopy, SnapMirror, and SnapRestore. We were also concerned about the ability to take four Microsoft-standard IBM, Compaq, and Dell servers and consolidate them into the NetApp filers with zero impact on the desktop. But those four Microsoft Windows NT personalities transferred directly to the NetApp filer without any disruption.

Q: Why'd you go with an IP-based solution vs. Fibre Channel?

Griffith: What it came down to was the kind of access we wanted. For us and our particular implemen-

tation, Fibre Channel would have been overkill. It would have required us to redeploy our existing network infrastructure. We got our filers wired back via fiber-optic communications, and they really rock. They do everything they need to do as far as average file access for the average client or our average storage solution.

Q: Some companies use NAS for database-centric OLTP applications. Would you consider NAS in this environment?

Griffith: We are considering it currently. To be quite frank, one of the environments we have that is still fairly low-volume in terms of total storage, but which contains extremely critical information, is an application we call CCI, Color Composition International. CCI is huge in Europe, and now it is penetrating the States. It is basically a high-end database content management system that runs on an RS/6000 over [Serial Storage Architecture, an interface that was developed by IBM] directly to SSA storage.

The direction we'd like to move with that is to an Oracle database. We'd like to move that content – the Oracle database engine storage component – to NAS. Our issue with the database world is not speed, but latency. Local databases and other applications get finicky about timeouts and other situations. We haven't done this yet, but I know it's in line for exploration and pioneering sometime this year.

Q: Let's get back to your implementation. When did it begin, how

long did it take, what problems, if any, did you have along the way, and how did you deal with them?

Griffith: There were a couple of speed bumps relating to code issues that we encountered when we rolled out the NetApp F740s. They were sporadically crashing for no apparent reason. This was in the infancy of the rollout. And NetApp was responsive to our problem. We sent the debugs to them, and they did the debug analysis. They had a patch out for it in a matter of 48 hours.

The NetApp implementation itself went flawlessly. We had NetApp turned up in a half hour in terms of configuring volume, bringing up NetVirus aliases – bringing everything online and patching to the network. We had the typical migration pains and difficulty associated with cleaning up user data and NTFS permissions. But that was no different than if we had moved from one NT box on to another.

Q: Would you characterize the F740 filers as a pilot?

Griffith: Effectively. A pilot in terms of looking at the NetApp technology and concepts associated with it before we invested significantly in the pair of F840 clusters we now have that top out at 12 terabytes. I think we had F740s four to six months before we jumped directly to the F840s.

Q: What's the current status of the system?

Griffith: The F840s are up, running, and doing what we need them to do. In fact, we're getting ready to expand them. *USA Today's* Dot.com

online product is about to requisition a terabyte of storage for all their internal office automation, standard file, Microsoft-standard desktop applications storage. We're also looking to expand our core storage in our McLean headquarters to support a very high availability FTP solution.

Q: How easy is this system to monitor?

Griffith: Very easy. In fact, so easy, it's a bit freaky. We no longer have to reactively manage all our storage. When I say "kind of freaky," it's because now only two of us are required to monitor the NetApp solution, and it works so well that sometimes we forget about it. It just stays in the background and does what it needs to do. Each week, it posts a bunch of reports based on stats, usage, utilization, things like that.

Of course, there's a maintenance agreement which brings a NetApp engineer assigned to the *USA Today* account once a month to do a complete top-to-bottom health status, which includes checking on the filers. Basically, to be quite candid, all of our reactive storage management has pretty much been eliminated.

Q: Did you use an integrator or consultant in the course of this implementation, or did you do it all internally?

Griffith: We worked it. We did a lot of research and, quite honestly, NetApp worked very well for what we needed. We did not have to get a storage integrator. We understood enough about our landscape, our production environment, our editorial

environment, and our standard office automation environment that we were able to cobble it all together and architect this thing ourselves.

Q: What are the primary benefits you've realized as a result of this NetApp solution?

Griffith: First and foremost, redundancy of information for the disaster-recovery scenario. On September 11, we nearly had to implement the NetApp filer in terms of disaster-recovery functionality. Our old headquarters were right across the bridge from Washington, D.C., on the Potomac River and along the flight path of National Airport. The point where the Pentagon was hit by the plane was approximately in the flight trajectory of our headquarters. We could see the Pentagon burning from our building.

There was obviously a great deal of chaos that day. At one point, all of the technical staff had moved to the White Oak, Md., location and had begun bringing all our backup systems on line. We had the peace of mind that comes from knowing that the 3 terabytes we had online were replicated over to White Oak. With just a couple of command lines, we could resurrect the personalities of all the NFS and Microsoft servers on that network from McLean. As you might imagine, business continuity was very important. So the primary benefit we realized from the NetApp solution was data availability.

Q: What other benefits were there?

Griffith: I would say our ability to

extend and expand data with minimal, if any, impact to the production and operational status of our business. That's a major benefit. We can basically expand and extend the function of these NetApp files while having very little impact on the client side. It is done transparently in the background; I really have little to do with the desktop side of it, but it's very significant that we can go out, grab a file for a user, and restore it in a matter of minutes, as opposed to an hour. And multi-protocol support is a huge advantage for us, because we're as much a Unix shop as we are a Microsoft shop – we talk NFS, FTP, CIFS – you name it.

Q: Was one of your goals to realize a hard-dollar return on investment with this? And if so, have you done that?

Griffith: We have not done a formal calculation on ROI with this particular project, but I can tell you anecdotally that it's beginning to pay off in a big way. The initial investment is done, and we're scaling this existing technology. It used to be we had entire operations departments managing banks of servers for storage that had been spread across dozens of NT file servers, including at least four or five critical key servers. Now, that's all been consolidated and reduced into a NetApp cluster.

So there is no hard ROI value or assessment that we've done. However, the value we've gained has been very clear and tangible. We're definitely getting a return on our NetApp investment even as we continue to extend and expand it.*