

*The most leading-edge users—those working at genomic research and energy exploration organizations—are starting to build storage architectures that can be truly shared among multiple users in a distributed network*

Research **firms** are **approaching**

# pushing storage to a higher level

BY JOHANNA AMBROSIO

**t**HE INFORMATION-RICH AREN'T SO DIFFERENT FROM THE REST OF US. They basically use the same kinds of storage technology—they just use more of it and push it harder. “Everyone has access to the same tools,” says David Hill, research director of storage at the Aberdeen Group in Boston, MA.

Indeed, if there's anything that separates those who work in storage-intensive research and development (R&D) environments from less storage-hungry companies, it's more in the way R&D shops approach storage—as opposed to the kind of gear they use. These high-need R&D centers are looking at storage as an enterprise-level architecture, as opposed to seeing storage in the traditional application-by-application, piecemeal kind of way.

stora

Joan D  
enterp  
Anada  
data, d  
month  
flexibil  
of stor

A woman with dark hair, wearing a blue blazer over a light-colored collared shirt, stands on a modern building's exterior walkway. She is smiling and looking towards the camera. The walkway has a white metal railing. In the background, a tall, dark glass skyscraper rises against a clear blue sky. The overall scene is brightly lit, suggesting daytime.

storage as a non-traditional, **enterprise-level** architecture

Joan Dunn, manager of enterprise computing at Anadarko, has 100TB of data, doubling every six months: "That's why the flexibility and expandability of storage is important."

## SHARING IS A NICE THING TO DO

Enterprise-level means building storage architectures that can be shared among multiple users in a distributed network. For now, these shops are clearly on the vanguard and in the minority (see "Sharing is a nice thing to do," this page). And many of the largest shops—although certainly not all—use hierarchical storage management (HSM) to help with their backup and archival requirements, and to keep their disk capacity clear for the highest-priority data.

The R&D gang does something else, too. They segment their storage needs into what kinds of information they have, regardless of the type of application the information is used for, and then they select the most appropriate technologies to support those different needs. So it's not unusual to have, say, both a network-attached storage (NAS) and a storage area network (SAN) setup in these shops.

Michael Peterson, president of Strategic Research Corp., a consultancy in Santa Barbara, CA, says this kind of thinking represents a real shift in the industry. He calls this model application-intelligent storage, which is when storage systems become applications themselves by managing bandwidth, resource allocation, scalability, security and other things behind the scenes. Application-intelligent storage also calls for optimizing storage for specific types of applications.

Although Peterson likes to call this smart storage—which takes the next step beyond enterprise storage and is shared and virtualized more than ever before—he believes it's still essentially stupid because the storage software really can't understand what type of data it's storing, and make appropriate decisions based on this knowledge.

### Early pioneers

One company that has started to implement much of this application-intelligent model—although they're not calling it that—is Anadarko Petroleum Corp., Houston, TX. Anadarko is one of the world's largest independent oil and gas exploration and production companies. In the past two years, Anadarko has begun to install various types of storage technology for different kinds of applications, and now has a NAS for its seismic data and a SAN for its relational database.

"We found that the NAS and the SAN are optimized for different things," says Joan Dunn, manager of enterprise computing at Anadarko. The company has 110TB of primarily seismic data stored in its Network Appliance Filer servers.

Ken Nadolny, manager of exploration and production

Companies in the R&D space are—like the rest of corporate America—struggling to learn how to make all their storage resources available across the board. They're coming up with different ways to approach the problem.

CardioNow, based in Encinitas, CA, provides cardiologists with the ability to share cardiograms with colleagues across the globe. Each cardiogram ranges from 70MB to 130MB, depending on how the input device is set up and how detailed the cardiogram gets.

At the core of CardioNow's setup is a DataDirect Networks' Silicon Storage Appliance, a high-speed RAID system. The system, now storing 1.5TB of information, is set up at a data center managed by a hosting vendor in San Diego, CA. After performing the cardiogram the test goes through a TI-based virtual private network (VPN) into the DataDirect box. After logging on with a password, users and guests can access the cardiogram via the Internet, and download the compressed image for viewing in a browser.

At Anadarko Petroleum, Houston, TX, the information is so proprietary that their goal has been to show it to as few people as possible. Also, most of the people who need to view the data work out of the main office, in one location. Nevertheless, the company is starting to look into working with a data service provider to be able to easily share information around the globe. The project, says Ken Nadolny, manager of exploration and production systems at Anadarko, is in the very early stages.

John Reynders, vice president of information systems at Celera Genomics, Rockville, MD, says his goal is to share more information among business units via a global SAN based in the corporate data center. "We're discovering that different parts of the business pull data together in different fashions, so we might have different data from multiple projects exposed all at once." He says that he's accepted, as inevitable, the notion that consolidation brings with it some amount of replicated data. "But if it's consolidated, then the overhead is amortized over many more servers" than in the traditional storage model.

One of the chief goals of the Cambridge, MA-based Genome Center at the Whitehead Institute, "is to get rapid access to data, regardless of where it's stored," says Michael Zody, manager of sequencing informatics at the center. For some of the biggest research algorithms, the center is still working off of direct-attached storage.

The center would really like to get to an environment where one could attach hosts to a SAN and then do file-level sharing. But to do this would require them to use different protocols for file-level access in a SAN—and "that's what NAS does," says K.M. Peterson, manager of computer operations for the center. He calls this the holy grail of storage systems: to be able to plug something in and access files with the efficiency of a SAN. "But we're a long way from that," he says.

## It's not always the quantity that creates complexity —it's the number of objects that have to be managed

systems at Anadarko, says seismic data is akin to a “sonogram of a pregnancy, except it’s a sonogram of the earth. But the quality is about as good as the medical kind.” Essentially, seismologists use sound waves to transmit energy through the earth, analyze the way those sound waves look after they come back out and hopefully discover if there are any oil reserves hidden beneath. For each square mile explored, one can easily record 300GB of data, according to Nadolny. The company is currently doing exploration in the Gulf of Mexico—all 615,000 square miles of it.

Currently, Anadarko has approximately 15% of its data stored on standalone servers, 4% on an EMC-based SAN and approximately 81% of its information on the NAS, Dunn says, adding that their storage requirements generally double every six to 12 months. “That’s why the flexibility and expandability of storage is important,” she says. “We believe the combination we have today will serve our needs for the foreseeable future.”

It’s not unusual for the company to decide it needs to explore a huge amount of property in a hurry, she says. “We have to be prepared to support—on real short notice—the storage of new data.” Looking toward the future, Dunn says they would like to pursue HSM “to apply intelligence to manage files and usage, to minimize cost and improve performance.”

Although in an entirely different industry, the Celera Genomics Group, Rockville, MD, struggles with many of the same storage-related issues as Anadarko. Celera has two major businesses—one which provides information about the human genome to researchers around the globe, and the other that uses genetic information to help identify possible therapies for different cancer-related diseases.

John Reynders, vice president of information systems at Celera, says the company is in the process of moving from SAN islands to a global storage network that will facilitate sharing. Right now the firm is a straight Compaq

shop, using Compaq’s AlphaServer ES40 boxes and Storage Works arrays and SANs. All together, Celera has approximately 100TB of spinning disk and another 150TB of tertiary storage managed by UniTree HSM software.

Although Celera hasn’t selected which SAN it will use, the firm did some testing of NAS vs. SAN. “We found better consolidation, and better connection of all servers at a higher connection speed with the SAN approach vs. the NAS,” Reynders says. Celera’s fiscal year starts in July, and plans are to unfold the new SAN now.

Reynders uses two other techniques that assist him in forecasting the firm’s storage needs. One is activity-based costing (ABC), an accepted means within the financial community of figuring out moneymakers or losers, and where the break-even point may be. The only difference is Reynders does this for storage. “We can fairly accurately predict what each business requires, per project, for storage,” he says.

Reynders also monitors which databases are used and helps decide what information needs to be safely archived because it hasn’t been accessed in a while. “We don’t know much about what’s being asked,” he says, “but we know what classes of data are changing more rapidly than other classes. So we know what to leave in cache, and what to move.” One particular genome, for instance, may need to stay active, while another can safely be archived.



**K.M. Peterson:**  
“I would kill for  
an HSM system  
that works.”

## It is much more difficult to manage an elephant than it is a Chihuahua, but even more difficult to manage a herd of mongrels than one elephant

Another leading-edge research organization, the Center for Genome Research at the Whitehead Institute, Cambridge, MA, hasn't implemented HSM for its 21 TB of data. "But I would kill for an HSM system that works," says K.M. Peterson, manager of computer systems operations for the center. "We've talked to several vendors that claimed they could solve our problem, but the volume of data transferring [we must do] to tertiary storage is just too large."

For that reason, he says, "we're very interested in EMC's Centera, which we understand is not HSM but is storage-optimized for high reliability, low cost and permanent storage."

For now, the center uses DLT-based tapes for backup, and is just about to purchase their first SuperDLT library, Peterson says. Why just tape? "We've lost enough RAID 5 disk sets that we're very paranoid about backups. So all our data is backed up to tape, a full backup every two weeks and incremental every evening."

Peterson's biggest challenge is how to find out where bits and pieces of information are stored. "It's like Unix files," he says. "The pathname is almost as important as the actual file that gets pointed to. Here, if you pick out a trace file, it's of no use to you unless you know exactly where it was stored and which project it's associated with." Each project may have 30,000 or more different items associated with it. If someone's looking for a particular file, it can be daunting to figure out where the file is.

Ideally, he says, he'd like to "describe the files on the system almost more from a database perspective" rather than in a traditional storage file system format. He says his company's stored files have grown so big that "I wish we could have simplified things a few years ago when we had the chance."

Instead, he says, "We push the technology so hard that some scientists have thrown two or three million files in one directory. It's hell on the system to do that, because Unix file systems aren't architected to do that. Our search tools just aren't meant for that type of scale." The solution? "We modified them to make a database call rather than have them look for a file and open up a directory."

The center generates 20GB to 30GB of gene-sequencing data each day. Much of this needs to be on disk for fairly significant amounts of time, because of the way scientists work. Sometime scientists need to retrieve data that was generated at the beginning of a project, Peterson says, "and we support that."

The center has traditionally been a Compaq StorageWorks shop running TruCluster and a NAS/SAN hy-

brid. The system is configured to show one system image, Peterson says, adding there are four servers logically connected on the back-end via a Compaq Fibre Channel SAN. Sometimes with applications that need high-speed access to specific data sets, he says, they move disks off the SAN and onto different hosts.

Although Peterson says the Compaq gear works fine, the center bought a Network Appliance Filer NAS system. The primary motivation was to "decrease the amount of time it takes to deploy and reconfigure storage. We think the amount of time the systems administration team has to spend on managing a terabyte of data will drop significantly," he says, adding that the NetApp filer system will make it more efficient to support the center's Windows users.

When the center first built its server environment it placed a premium on flexibility, because the specific systems requirements of the then-new Human Genome Project were still unknown. Now, 20+ TB later, they understand their workload's scalability requirements much better. At this point, scalability takes a back seat to manageability, speed of deployment and integrated solutions.

"It's most efficient to support Windows over the Windows file-sharing protocols," including the common Internet file system [CIFS] or server message block [SMB], Peterson says. "To do that on the TruCluster system, you're faced with significant tasks. It's not difficult, just time-consuming, because you've got to implement Samba or Advanced Server, the two things that support Windows within the Unix environment."

In comparison, the NetApp filer already includes CIFS and SMB support. The company has bought a filer that can support up to 15.5 TB.

What it ultimately comes down to is how conservative one wants to be in choosing technology. "We're very risk-averse here, given what we do and our requirements for availability," Peterson says. In fact, the center resisted going with NAS two years back, because "We felt it was not sufficiently mature. Whether that was a good thing, I can't really say."

Lesson learned? The bottom line says Aberdeen's Hill, "is that it's not always quantity that creates complexity; it's the number of objects that have to be managed. It is much more difficult to manage an elephant than it is a Chihuahua. But it's more difficult to manage a herd of mongrels than it is one elephant." ☉

*JOHANNA AMBROSIO is a freelance writer in Marlborough, MA.*