

Get peace of mind with the right technology.

Tripp Lite SMART1500SLT Tower UPS
\$349.99 **TRIPP LITE**



Replay



Shop CDW-G

Subscribe to the FedTech E-Newsletter **GO**

FEDTECH

Search...

GO

FedTech

Google

CDW-G

Subscribe to Fedtech magazine. **ORDER NOW!**

HOME MAGAZINE RESOURCES TECH TRENDS BEST PRACTICES HOT TOPICS E-NEWSLETTER

FedTech » Magazine » August 2010 » Feature » Rite of Passage

[Feature]

Rite of Passage

As data needs and requirements grow, agencies dynamically adapt their storage infrastructures to meet demand.

By Karen D. Schwartz

7/29/2010



By combining SAN and virtualization technologies, the Choctaw Nation Health Services Authority gained the upper hand on managing quick growth of electronic health records, say the authority's Shane Miller (left) and Dwane Sorrells.
Photo: Shane Bevel

Like most healthcare facilities, the Choctaw Nation Health Services Authority relies on a data center, complete with the latest technologies. But unlike many facilities under the purview of the Indian Health Service, this healthcare system is run by the tribe itself — a fact that the IT team never takes for granted.

To support the facility, the Choctaw Nation has built and maintains a state-of-the-art data center, because, as IT Director Dwane Sorrells says, "we were able to run it on our own, so we did."

The tribe took over management of its healthcare services in the 1980s. Since then, it has funded and built a 45,000-square-foot health facility, which provides medical care to the Choctaw Nation and other Native American tribes of Oklahoma from its base in Tahihina.

The data center, built with the hospital, has grown from a handful of green-screen PCs, says Sorrells, to a fully functioning data center built on about a dozen HP servers hosting 100 virtual machines, along with a comprehensive and growing storage infrastructure.

The organization's storage needs grew so quickly that the IT department had to take serious steps to keep up, he says.

Just a decade ago, the only storage requirements were for the patient database and a registration system. Today, the data center stores and backs up a Microsoft Exchange e-mail system and a host of in-house applications and databases, and also has accommodated the shift from paper to digital forms and the organization's web presence. In addition, the move to virtualization, while beneficial in shrinking the data center footprint and improving service to users, has resulted in a new need: more storage for virtual desktops.

The explosive storage demands of the Choctaw Nation Health Services Authority are typical of government institutions, which tend to be extremely storage-intensive. According to a recent report by the Enterprise Research Group, only the communications and media industry has as much storage capacity as the federal government, on average. In both cases, 36 percent of organizations maintain 250 terabytes or more of storage, compared with 19 percent of educational organizations and 14 percent of retail businesses.

But how do agencies address quick data growth smartly? A look at three facilities — the Choctaw Nation's health data center, the Air Force's personnel records center and the FBI criminal records center — illustrates that agencies have set a path for consolidation and networked storage to ensure quick and easy access to data even as federal data stores grow.

SAN with a Side of Virtualization

For the Choctaw Nation facility, perhaps the biggest data spike has come with the shift to electronic health records, which must be

- » comment del.icio.us
- » print digg this
- » email reddit
- » rss feeds

RELATED MOST POPULAR

When Disaster Strikes

As WiMAX emerges, does it offer a COOP option for secure communications?

The Many Faces of Cloud

DISA, NASA and other early adopters are cutting costs and improving IT services through the cloud.

On a Cloudy Day

As agencies face across-the-board budget cuts in the coming fiscal year, IT shops must define exactly how and where cloud computing will figure in their agencies' infrastructures.

Peeling the Onion

Are agencies ready to make the most of data center consolidations?

Beat the Heat

These tips can help effectively and efficiently cool high-density computing environments.

Brief Byte

Where do agencies expect to reap cost savings by reducing footprint?

Data on Tap

In federal IT shops, when it comes to unlimited access, the idea really keys in on data: how feds can tap it, use it, store it — from any device and from anywhere in the world.

The Many Faces of Cloud

DISA, NASA and other early adopters are cutting costs and improving IT services through the cloud.

Rite of Passage

A look at three facilities illustrates that agencies have set a path for consolidation and networked storage to ensure quick and easy access to data even as federal data stores grow.

SANs: Taking Advantage of Storage Pools

USDA service center finds SANs offer highly available data and cost-efficient storage.

The 21st Century Government™



Everywhere Our Government Works, CDW-G Is There

[Visit the 21st Century Government](#)

SUBSCRIBE



Get what you need to know about information technology solutions to improve your agency.

[subscribe now »](#)

hosted and securely stored. That addition alone has raised the patient database from less than 20 gigabytes a year and a half ago to about 80GB today, and the growth shows no signs of slowing. To address the health center's rapid data expansion, the data center sports three HP StorageWorks EVA4400 storage area networks, along with a smaller HP StorageWorks 1500cs Modular Smart Array as a file server.

"Before, if we needed storage for anything, we would overbuy," explains Systems Engineer Shane Miller. "For example, when we moved to digital dental X-rays, we bought a siloed pizza-box-style server and loaded it with hard drives. Even though we may have had only 20GB worth of digital X-rays per year, we bought enough storage space to cover them for the next 10 years. That's a lot of unused space that's not being properly utilized."

The new SAN infrastructure avoids this problem, he says, because all storage is in the same place physically, and the IT staff can provision more as needed. It's also easy to scale out storage for a specific server, he says.

The group also is looking into EMC's Avamar backup and recovery solution to add deduplication to the mix.



Suite of Storage Options

Few groups demonstrate the government's data explosion better than the Air Force's HQ Personnel Center. Located at Randolph Air Force Base, Texas, the center collects, manages and stores personnel data from active duty, reserve, civilian and retired Air Force staff around the world. That's a big change from a decade ago, when each of the service's bases managed its own personnel center.

Today, the Air Force Personnel Center stores records for about 780,000 active and retired Air Force personnel, and it must keep that information for 67 and a half years after an employee has died, according to Office of Personnel Management rules.

But that's only half of the story. No longer does the center simply enter pertinent personnel data into databases. Now, it stores entire scanned

documents as PDF or TIFF forms. Although the system requires much more storage, the documents are fully searchable and therefore much more valuable to the Air Force, says Mark Stewart, the center's enterprise storage and SAN manager.

Leader of the Pack: Although federal government adoption of deduplication lags some other industries today, **50%** of federal government respondents say they plan to implement dedupe technology over the next two years — more than any other industry.

SOURCE: 2010 Data Protection Trends, Enterprise Strategy Group, April 2010

"In 2004, we were estimating a 10 to 13 percent annual data growth rate, which equated to about 285 terabytes," Stewart says. "But that quickly jumped to 25 percent the following year and 35 percent the year afterwards. Right now, we're at about 790 terabytes."

The exponential storage growth has forced Stewart's team to change the way it uses storage technology. To keep up with demand, the storage infrastructure has evolved from a collection of Fibre Channel disk arrays connected via direct-attached storage to a scalable, flexible storage system.

Today, data for civilian personnel is housed on one HP StorageWorks XP12000 disk array with 40TB of storage and a StorageWorks XP24000 disk array with more than 250TB of storage. Information for military personnel is stored on a pair of HP StorageWorks XP12000 disk arrays with a total of 35TB of storage. About 500 Windows-based systems that run applications and web servers are hosted on two HP StorageWorks Enterprise Virtual Arrays, an 8100 and an 8400.

To complete the storage upgrade, the organization last year installed three Data Domain DD690 gateways for data deduplication, which Stewart says were needed to replace a maxed-out

tape-based backup system.

"We went from having to get data off of the backup system every 10 days to never having to delete it," he says. "We've had all of our data on the Data Domain system for more than 15 months and haven't had to delete a single backup. We're only using 40 percent of our 100TB array, and we've saved about 1.2 petabytes of data in that 40 percent."

Tiered Techniques

The FBI's Criminal Justice Information Services Division also is dealing with spiraling storage requirements. Located in Clarksburg, W.Va., the division manages criminal data from state, local and tribal law enforcement organizations. These records cover a wide spectrum of types, including fingerprints, warrants and stolen property reports.

Digital images consume a large portion of the division's storage capacity; changing requirements several years ago to increase fingerprint resolution caused a staggering fourfold increase in storage needs. That might not sound like much, but the storage growth rate went from 10 percent a year to 40 percent a year, says Jean Archambault, the division's chief of technical planning and control.

Fingerprint images account for the bulk of the growth: The division stores about 65 million sets at 45 megabytes per fingerprint card uncompressed, and it must keep cards for 99 years. The division currently has about 4 petabytes of storage, an amount that continues to climb.



That growth prompted the division to switch from an aging direct-attached storage architecture to a tiered infrastructure that now consists of several EMC Symmetrix VMAX disk arrays, along with an EMC SAN, Fibre Channel and serial ATA drives, and solid-state storage.

“Some of the imagery that we don’t access all the time is stored on slower SATA drives, and we also have some solid-state storage for a lower tier,” Archambault says.

Eventually, the division would like to move from traditional Fibre Channel to Fibre Channel over Ethernet. FCoE would let it combine its communications and storage networks by virtualizing a single cable into both, thus reducing the cost and management load.

Further down the road, the bureau is contemplating object storage, which segments information by data objects, including data attributes and metadata.

The idea is to move away from relational databases and find needed information by relating data together, Archambault explains.



Photo: Gary S Chapman/Getty Images

Optimizing Storage

The dramatic growth in storage needs among federal organizations doesn’t just add up to increased capacity — it also puts the impetus on agency IT shops to optimize data stores that already exist. By doing so, agencies can postpone costly upgrades and meet cost and energy reduction guidelines.

Archiving is a solid way to optimize storage. After all, if an organization can archive 20 percent of its data per year, it will have room for 20 percent growth.

“It may not be exciting, but it’s a time-tested, proven route for stretching your storage dollar and optimizing your storage,” says Greg Schulz, founder of the StorageIO Group.

Tiered storage, which uses less expensive storage to retain data with less value and more accessible, flexible storage for current needs, also is a storage optimization winner. To make the most of it, Schulz recommends using a discovery tool (examples are the Hitachi Data Discovery Suite, EMC Ionix and CommVault Simpana) to monitor data in storage, its location and when it was last accessed.

Compression is another underused technique. Although it’s commonly thought of as a tape-only method, it’s also applicable to databases and file systems. Schulz even considers deduplication as a sort of “intelligent compression” that increases reduction ratios.

Finally, thin provisioning is an excellent technique. Much like an airline overbooks its flights, thin provisioning prevents wasted space by allowing administrators to provision enough storage for users without over-allocating storage. With this method, when more storage is needed, it is automatically allocated to the user.

That’s the approach the Choctaw Nation Health Services Authority recently took with its new Virtual Desktop Infrastructure system.

“If an operating system disk is 20 gigabytes, we don’t actually provision all 20GB at once — only what’s needed,” says IT Director Dwane Sorrells.

Another method Sorrells’ staff is using to optimize storage is auditing and monitoring data storage usage, performance and capacity. By using Vizioncore’s vFoglight Pro, the IT staff can set a threshold for data stores and receive an alarm if that threshold is exceeded, prompting them to reprovision the storage.

[Related Articles]

- [When Disaster Strikes](#)
- [The Many Faces of Cloud](#)
- [On a Cloudy Day](#)
- [Peeling the Onion](#)
- [Beat the Heat](#)
- [Brief Byte](#)
- [Data on Tap](#)
- [The Many Faces of Cloud](#)
- [Rite of Passage](#)
- [SANs: Taking Advantage of Storage Pools](#)

» comment | » print | » email | del.icio.us | digg this | reddit | rss feeds