Geospatial Multistate Archive and Preservation Partnership (GeoMAPP)

State Architecture and Storage Challenges

December 31, 2011
Version: 1.0

Storage was one of the major challenges faced by the GeoMAPP partnership. The partnership discussed at length how to predict how much storage each partner would need, scoping the project for data transfer, as well as multiple copies in the form of preservation versus access copies. However, none of the partners anticipated to the amount of time it would spend in talking about their storage needs and actually being able to purchase storage to meet those needs. This document details the current storage environment of each of the GeoMAPP full partners. In addition, each partner details the success and challenges they experienced purchasing additional storage and configuring it to meet their digital preservation and access needs. This document is intended to serve as a companion document to the GeoMAPP Storage Primer.

Kentucky

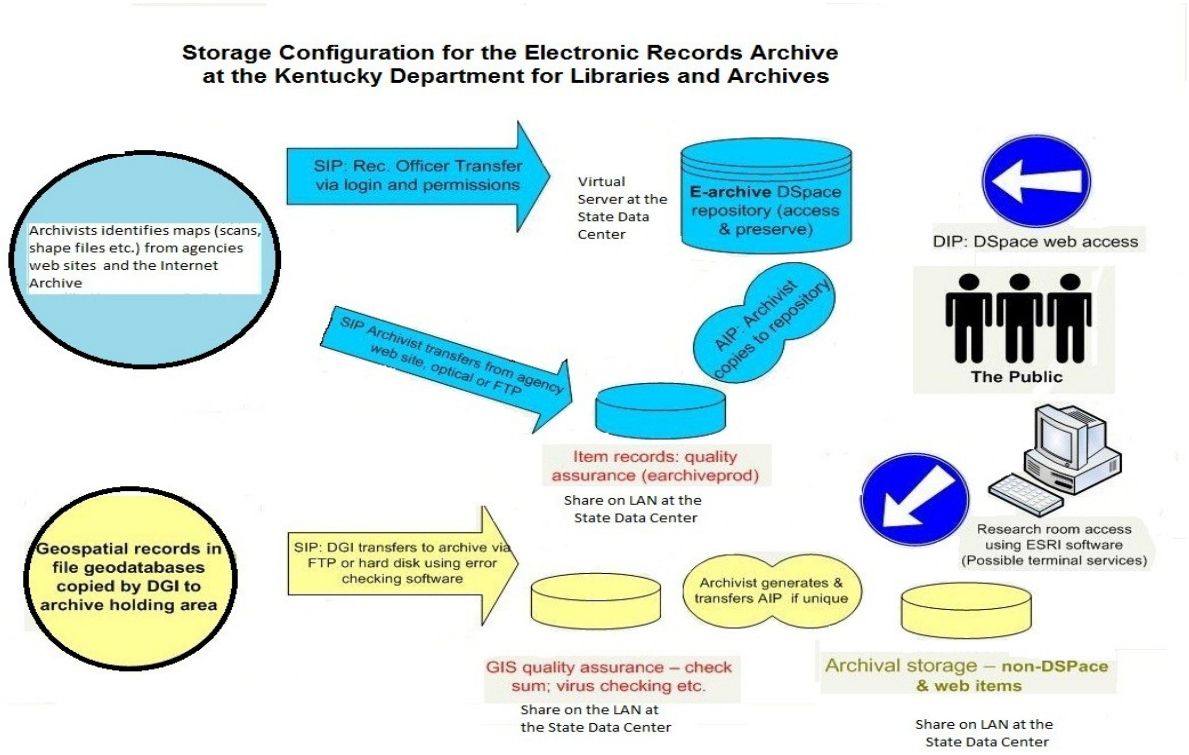

Document current State SAN approaches

The storage for archival records has evolved during the GeoMAPP project as new hardware has been acquired. The SAN storage which now totals 20 TB is accessible via shares on separate disks which have been designated for copies of geospatial records that are generated as the records are received, checked, replicated and backed up. With the final purchase of blade servers the archival SAN storage will be segmented from other KDLA processing with its own pool of server hardware.

All KY vector databases and selective KY raster files were copied to the e-archives by Division of Geographic Information (DGI) staff via direct network connection using robocopy and hashing to document and verify authenticity of transfer. Each snapshot is accompanied by a report in PDF format that lists layers in the database grouped by category. Additionally, archives staff downloaded various shape files and PDF files from agencies websites for import into DSpace, the Archives access tool for electronic information. Once the State Archives received the database snapshots, it is verified by rerunning the hashing software and an access entry is created in DSpace along with the snapshot report which provides documentation in a format that all users can open.

Before the project began DGI had previously transferred two copies of older imagery to the archive via portable hard disk and CD. Due to volume, the bulk of the tiled orthos that are converted to mosaic format and used by the KYGeonet were transferred via hard disk to network storage at KDLA. These periodic image snapshots, which date back to the 1990's were brought to the archives, but were not replicated, because they are also being maintained as snapshots in the KYGeonet. When they are no longer needed as active files in the KYGeonet, KDLA will replicate the years that are no longer being maintained in the KYGeonet.

Other than maintaining check sum files (using MD5 protocol) for each of the copies of the geospatial records, KDLA does not use a formal copying and auditing software like ACE, although it hopes to establish that capability in the near future. KDLA will test Vice Versa a program that maintains files properties such as date during the copying process

Document successes and challenges with procuring and maintaining SAN storage in state government

Purchase of storage and network equipment is governed by state price contract and adherence to the Enterprise Architecture and Standards (EAS). If an item is not on state price contract (select vendors), then it has to go through a bidding process. If it does not follow the Enterprise Architecture and Standards, then a written request for exception must be submitted to the Enterprise Architecture and Standards Committee for approval. Permission to purchase usually has to be approved at the cabinet level.

The servers and SAN storage owned by KDLA are maintained in the state data center by KDLA IT staff. The central IT organization (COT) charges KDLA for the floor space in the data center, thus providing physical security and risk mitigation in a proper environmental setting. Unlike most other agencies records that have been moved to the State Data Center, storage use is not charged to KDLA, because KDLA IT purchased and maintains the equipment (backups, operational support and rights control). Assignment of share quota capacity is controlled by KDLA IT in consultation with electronic archives staff. Permissions to edit, view, delete and copy electronic records are assigned by KDLA IT staff following request by the archives staff.

The storage challenge is to acquire and maintain affordable storage that meets the state architectural standards as well as archival needs. KDLA's current strategy has been cost effective, but it continues to face the challenge of satisfying the desire of the state IT unit to consolidate resources including hardware and staff.

Montana

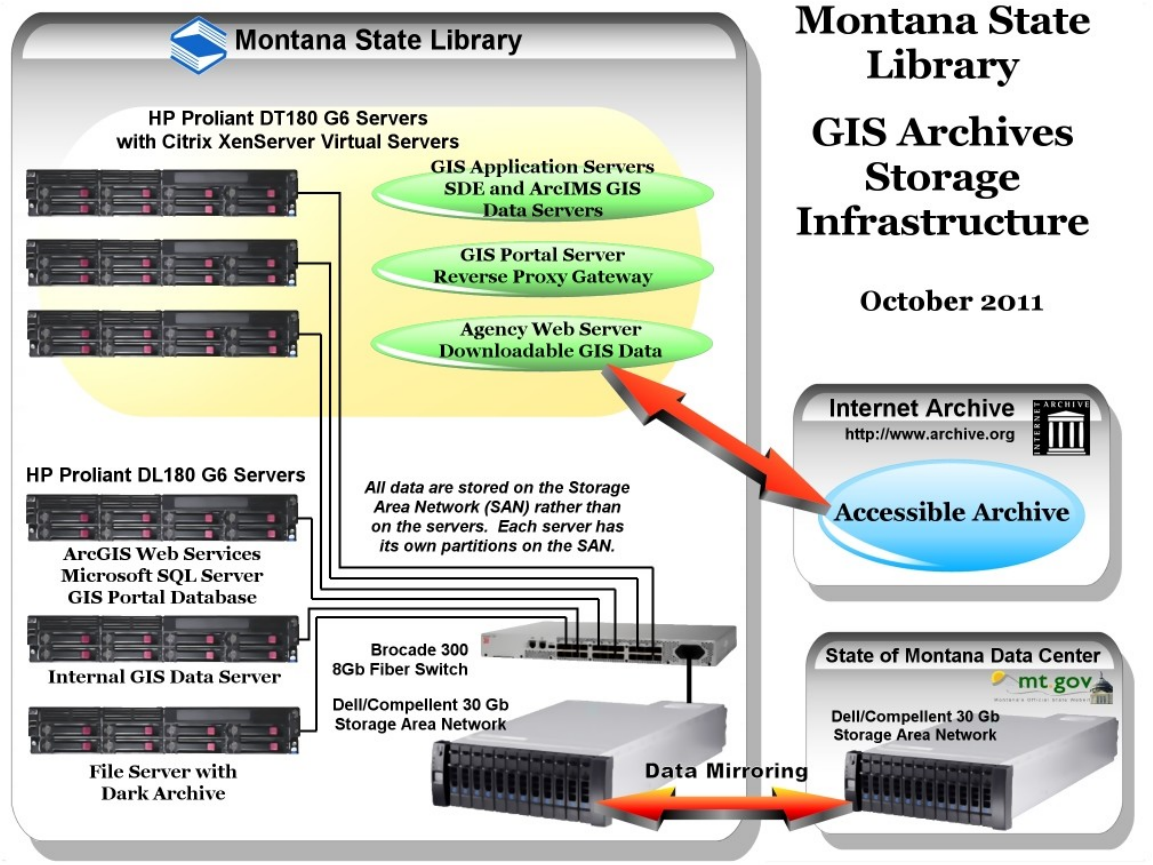

The Montana State Library has a production GIS environment that consists of virtual web servers, virtual SQL database servers, a virtual server that supports the GIS Portal Toolkit, two physical servers that support ArcGIS Server web mapping services and a physical server that supports an internal SQL database with SDE for versioned editing. Most of these servers are connected to a production storage area network (SAN). This SAN is one of two thirty-terabyte Compellent SANs owned by MSL, one of which was purchased, in part, with funds from GeoMAPP. The production SAN is housed in a data center located at MSL. The second SAN which will serve as a mirrored, offsite backup of our production SAN will be located in the new State of Montana Data Center during the fall of 2011. Both SANs are fully configured and managed by the MSL Network Administrator who is a direct report to the Digital Library

Director and who works closely with NRIS staff. These two SANS are connected via a redundant State-managed one gigabit fiber network. Although MSL will rely on the State Data Center for power and network connectively, configuration and management of the mirrored SAN is fully under the auspices of MSL.

Master files of archived GIS data will reside in a dedicated volume on the production SAN. Data will be managed in a folder structure that is organized by date within folder structures for series data such as cadastral data, collections data such as the Yellowstone River collection and general GIS data. Current capacity for this volume will be set at ten terabytes but can be increased on the fly. The MSL Network Administrator is currently developing a replication strategy that will replicate data on the production SAN to the mirror SAN based on the frequency of data change. Archived GIS data will likely be replicated on a quarterly basis.

MSL is currently developing an inventory tool that will ingest GIS data, place a copy of the data in the MSL dark archive, generate administrative metadata and will manage the integrity of the data. This inventory will rely on Bag-it to generate a checksum that will be used for on-going integrity checking. Corrupted files will be compared to access copies and backup copies for verification and replacement if necessary. Initial system requirements for the MSL inventory system do not include scripts that will automate data management including integrity checking however a variety of automation is planned for subsequent releases.

MSL's primary access repository is the Internet Archive. MSL is still researching the use of the Internet Archive as an access repository for archived GIS data and will rely on the Internet Archive's hosted solutions for storage of accessible archived GIS files. The Internet Archive maintains metadata records for archival content that are derived from metadata records that MSL maintains and that are accessible through the Montana GIS Portal.

The two SANs were purchased over two fiscal years. The total cost, which included vendor support for installation, was $135,000. Ongoing internal costs to support the MSL data center are difficult to isolate as they are included in agency fixed costs. The cost to host the mirrored SAN at the State of Montana Data Center is approximately $5,000 annually paid to the Montana State Information Technology Services Division (SITSD).

Document successes and challenges with procuring and maintaining SAN storage in state government

The MSL Network Administrator researched, reviewed and presented alternatives for MSL storage solutions. This research was conducted over approximately four months. The state maintains a list of preferred vendors which eliminates the need for a request for proposal. Once a solution was decided upon and agreed to by the Digital Library Director (a quick process

given MSL's small staff) SITSD required MSL to complete an information technology procurement request. This request was reviewed and approved by the SITSD executive management team including the state Chief Technology Officer. Total time for review and approval was no more than two weeks. Once approved, a purchase order was issued and an order placed. The equipment was purchased, is owned and is managed by MSL. Management includes complete control over the hardware, configuration of the server software including user access and storage allocation as well as performing backups of the data.

North Carolina

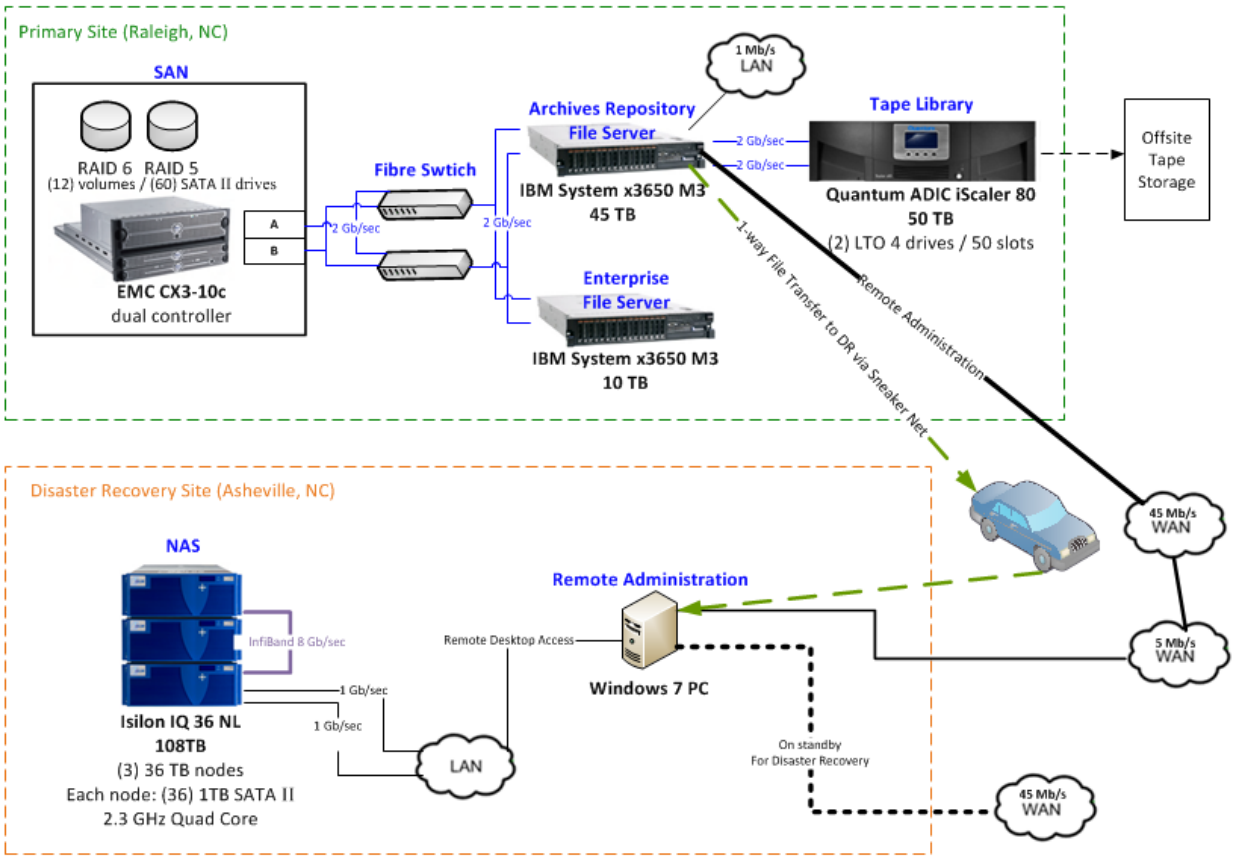

Document current State SAN approaches

The North Carolina State Archives began building its digital repository storage capacity in 2005. Using funds from the operations budget of the NC State Archives and the Division of Information Technology of the NC Department of Cultural Resources, the Department purchased an EMC CLARiiON CX3-10c SAN from EMC. Initially, the SAN had 15 TB of usable storage. Of that 15 TB, approximately 4.8 TB was allocated to store archival files of the North Carolina State Archives and the State Library of North Carolina.

In 2009, through GeoMAPP grant funds, the NC State Archives & Records Section purchased 12 TB of usable storage that were used to fill unused hard drive bays in the existing SAN. This storage was dedicated to the GeoMAPP project and used for testing and file storage.

In 2009, the NC State Archives increased the storage capacity of the digital repository. In conjunction with DCR –IT, the Archives replaced 40 existing drives with sizes ranging from 300 GB to 750 GB with 2000 GB drives, thus increasing the usable storage to above 70 TB, of which 45 TB were allocated for archival storage. We have an additional 13 TB of storage that has not been allocated that may one day be used for archival storage.

During the in-place upgrade, we converted all of the RAID volumes to use RAID level 6 to decrease the risk of data loss due to hard drive failure. Initially, the NC State Archives created two 22.5 TB drives, one for GIS data for the GeoMapp initiative, and the other to be shared by the NC State Archives and the NC State Library for other archival file types. Several months later we combined these two logical units (LUNs) to a single logical unit so that either network shares could grow beyond the original 22.5 TB. (The total of the two of course would never grow beyond 45 TB without adding additional storage.)

This product will reach the end of its supportable life in 2014, and we expect to make no more expansion to this SAN unit. The North Carolina State Archives plans to replace or upgrade this unit at the end of the supportable life.

Backups are made with a ADIC Scalari80 tape library. This tape library has the capacity to hold 50 LTO-4 tapes at one time (upgradable to 80). It has two drives and is capable of backing up 8000 MB / minute. We have separated all archival data into files that are completed (preservation and access) and those that are still in progress (working and staging). Those that are completed are backed up weekly using a differential type backup and backed up fully every three months. The full backups are kept for 6 months, while the differentials are kept for 1 month. Those files that are still in progress are backed up daily using a incremental style backup and backed up fully every 3 months. The full backups are kept for 6 months, while the incremental are kept for 3 months.

In 2011 The NC State Archives & Records Section purchased another storage unit from EMC, however instead of purchasing the successor to the CLARiiON, we decided to purchase a 3-node ISION Networked Attached Storage (NAS) device with 70 TB of usable storage.

The ISILON NAS is housed in at the Western Office of the Office of Archives & History, in Asheville, NC and will be used for disaster recovery purposes.

Instead of purchasing sufficient bandwidth to enable automatic replication between these two sites we plan to leverage the Audited Control Environment software (ACE) along with some

custom software to help us determine what changes need to be "shipped" out to the western office. By mailing external e-SATA drive back and forth and relying on staff at the Western Office to un-box and connect the drives to the dedicated front-end server we will be able to keep the two storage units roughly synchronized. This together with our off-site backups should allow us to recover from a disaster only affecting Raleigh or Asheville.

Currently we only have a 5MB/s DSL and that is being used by the staff at the Western Office. We are considering going ahead and requesting that a 45 MB/s fiber connection be installed at the western office and then disconnected. In this way when we need to have it reconnected we can do so in a matter of days instead of weeks.

When it is time to replace the CX3-10c in Raleigh we will replace it with a ISILON NAS as well so that we will have a matched pair.

Document successes and challenges with procuring and maintaining SAN storage in state government

When the GeoMAPP project began, the Department of Cultural Resources was in the middle of being "consolidated" with the statewide Information Technology Services service pack. Because we had a SAN unit on premises from 2005, the North Carolina State Archives was able to purchase the additional storage for GeoMAPP. Somewhat confusing during this time was the actual time period in which the Archives could purchase storage. The Chief Information Officer instructed staff that it needed to be ordered during one of ITS' mass storage orders. This meant that the NC GeoMAPP team faced a delay in transferring records. Since that time, it has become easier to purchase storage as storage vendors are now on contract.

The State Archives is in part supported by the Archives and Records Management fee, a $5 fee assessed to all recorded deeds with the exception of Deeds of Trusts. In the latter part of the GeoMAPP project, the General Assembly began to more closely monitor expeditures and the Archives sought an exemption to using the ITS Western Data Center. By purchasing our own storage, the North Carolina State Archives maintains more control over the data itself, a core critical component of a trusted digital repository. With that being said, the North Carolina State Archives is allowed to purchase storage but will have to generate additional documentation and seek more levels of approval to purchase additional storage.

Utah

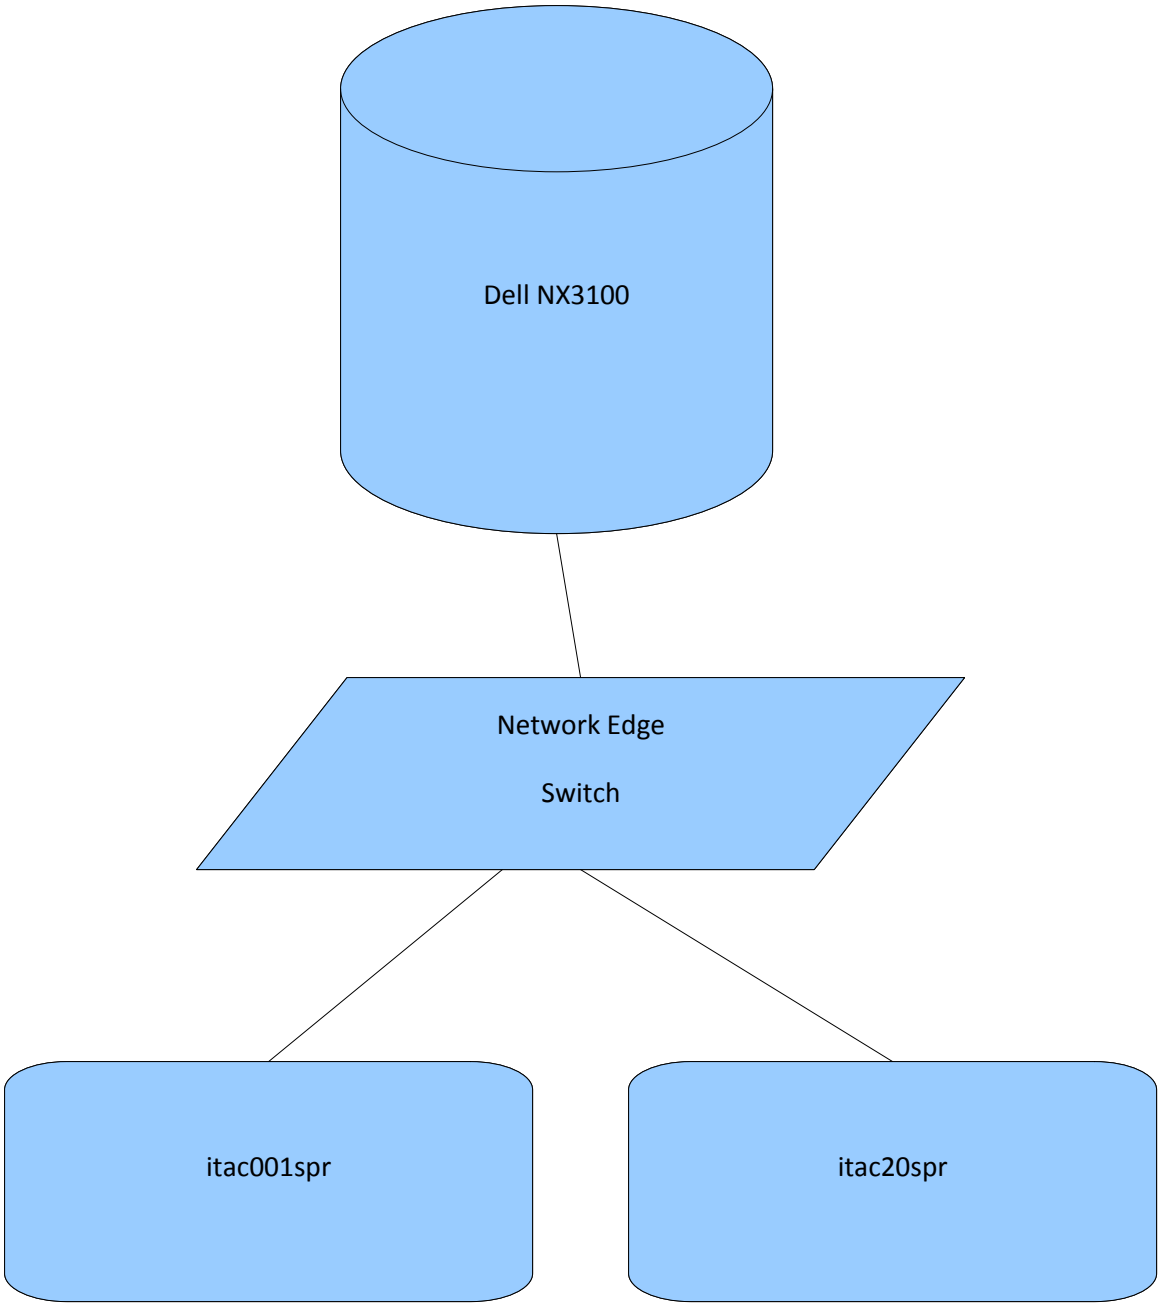

Document current State SAN approaches

The enterprise storage space that the Utah State Archives uses is provided by the Department of Technology Services (DTS). While SAN storage is available, the Archives has negotiated a lower-cost solution. A couple of years ago, DTS obtained a NAS device, which allowed them to offer storage at a lower rate than its SAN offerings. Since the NAS device is not directly owned by the Archives, the device itself, how it is networked, and capacity are unknown. The Archives has had 19 TB of space installed on the device during the past year when DTS offered a flat fee for storage. In the next fiscal year, the rates will be much higher. One estimate from DTS for how much 50 TB would cost using the NAS device was in the neighborhood of $250,000/year. This amount is not budgeted, and is unlikely to be obtained through normal funding requests.

To order more storage, the Archives uses the form provided on DTS' website. Generally, all that is required is to provide them with the billing code for charging the Archives, and they do the ordering. The vendors they use to procure this storage from are unknown. Purchases of hardware and software made through the DTS process have several levels of approval: Archives budget director, Administrative Services (the department over Archives) Finance director, and the DTS IT director responsible for the Department of Administrative Services. Service level agreements are in place between Archives and DTS. They are reviewed and updated each year.

Document successes and challenges with procuring and maintaining SAN storage in state government

Storage is administered by DTS when it comes to the online, enterprise-class version. The Archives has no ability to set user access or power the storage device on or off. For new users to gain access to this storage, a ticket is placed with DTS. This storage is fully backed up by DTS on an automated schedule. One full backup is kept, plus five versions of files which are edited frequently. The Archives can request that individual files be restored from backup if it notices file corruption.

Temporary storage space used by the Archives when processing records consists of portable hard drives, stored in the Archives' permanent records room. These are not backed up, but the Archives is in full control of them. They are low-cost, high-risk devices, whereas the enterprise storage is high-cost, low risk. Since money spent on enterprise storage is very real, but loss from hard drives is only potential, the bulk of Archives' electronic records are currently on hard drives. Many of those records will be ingested into the Archives' AXAEM system, which utilizes the NAS storage space. AXAEM will provide integrity checking for stored records.

The Archives' network connection was recently upgraded to 1 GB. Data ingest can take place either from the desktop, using network connections, or via the server, which makes ingest much faster. When a user recently ingested over 3,000 files over the network, the process (which included metadata extraction, lengthening the time requirements) was run overnight. Since the server is located off-site at the state capitol and managed by DTS, it is a little bit of a hassle to transport the files on media to the server, load the records, and then go back to the office, login to the application, and run the ingest process. So running it from the desktop is often easier even if it takes a bit longer to finish. The threshold that would make the run to Capitol Hill "worth it" would likely involve several thousand more files that need ingesting at the same time.

While the storage space impacts the Archives' ingest process, it also impacts data replication and creation of preservation masters. One form of media that is currently being investigated for dark archives purposes are Millenniata disks, http://millenniata.com/, a form of DVD that isn't manufactured the same as other DVDs. Data is etched onto the media and does not make use of dyes, which appears to make the product highly durable. The DVD is created using a special drive but can be read by any standard DVD player. While the storage capacity (just like other DVDs) is 4.7 GB, its durability is viewed as very attractive. For any media choice, the twin risks are media failure and technological obsolescence. If technological obsolescence becomes the only remaining risk (not counting format migration/emulation issues associated with all electronic records), then the Archives should be able to save a great deal of money by keeping records on the M-disks until the next generation of media becomes available. The greatest advantages are these:

- Cost per disk is about $3, and each writable drive is about $100. Fifty terabytes could be stored for about $32,000 if each disk was filled to capacity.
- DTS would not charge the Archives a monthly fee, regardless of how many terabytes are accumulated.
- The media can be purchased outright without needing permission from DTS, so Archives would "own" the media as well as the records.
- It could be stored locally with other archival records, giving the Archives and visiting patrons easy access.
- It could be replicated and distributed off-site, even out of state/country
- Automated checksum audits may still potentially be available if the disks are stored in some kind of jukebox setup that the AXAEM application has access to.

The disadvantages are these:

- Some records won't fit on a 4.7 GB disk
- The storage space won't help users access the records online, unless by some miracle the jukebox system could work with a web server.
- Constantly adding disks to a jukebox could become annoying or labor-intensive

- The disks only write at 4X speed currently, which means that it could take considerable time to write terabytes of data to disk, and then replicate the disks

If the M-disks are used, this would change processing procedures somewhat. The data would first be ingested to AXAEM as a SIP, then likely be manipulated in-house on portable hard drives, then written to M-disk from those drives. The final AIP ingest would happen from the desktop using the M-disk version, so that the checksum that is captured during ingest should reflect what is on that disk. Then the disk would be stored. The copy sent to the server would either stay as the DIP, provided that the Archives has funding for online storage, or would be removed after ingest and leave all copies on M-disk. Disks would then be bar-coded and their locations tracked. This process does present a problem of not storing a copy of the metadata with the record, if the M-disk is written prior to metadata being added. The sequence of procedures might change due to this factor. If the M-disk is written as the last step, then the network capacity will be stretched to its limits as data moves back and forth during the ingest/write-to-disk process. This will likely annoy users.

One piece of the storage equation that would be nice to modify is the use of portable hard drives as working space during processing. The Archives is very interested in obtaining either a NAS device that could be housed at the Archives and accessible to their LAN, or perhaps have their own server for all ingest/storage purposes. Currently, DTS has a policy of virtualizing all servers, which is why the hardware has been consolidated at the state capitol. To purchase their own hardware, the Archives would be in direct conflict with DTS' policies. In the past, these purchasing decisions have required DTS' approval. Due to this conflict, Archives management has developed a strong desire to be independent of DTS, at least where storage is concerned.

Use of cloud storage may be explored, but that would bring its own set of advantages/disadvantages. The AXAEM application could be hosted on the cloud, along with any records ingested into it. However, this application is also integrated with other software at the database level, and the other software will not run on the cloud. It is unknown whether the shared Oracle database could be accessible to both applications. Another negative is that large ingests would likely take much longer and would necessitate being done via desktop. Records would be stored outside of the Archives, which makes some archivists nervous about whether or not the records would be truly preserved. If the connection to the cloud is broken, the Archives would be unable to access its data. Cloud security is likely still an issue, since some archived records are very sensitive.

Conclusion

The GeoMAPP partners will continue to explore storage options and opportunities in the future. The partnership in an informal manner will continue to exist. Certainly, technology will continue to evolve. The partners are committed to continuing in the work of preservation and will continue to explore access avenues. This requires us all to work with our Information Technology partners. As we do so, we will continue to evolve our storage solutions in order to meet our responsibilities and ensure that the records are trustworthy, preserved, and accessible.