The Perspectives of Digital Curators on Building Distributed Repositories

Caryn Wojcik
Government Records Archivist
Records Management Services
Michigan Historical Center
Department of History, Arts and Libraries
3405 N. Martin Luther King Jr. Blvd.
PO Box 30026
Lansing, Michigan 48909
wojcikc@michigan.gov

Richard Marciano
Lead Scientist
Sustainable Archives & Library Technologies (SALT) laboratory
Chien-Yi Hou
Digital Preservation Specialist
Sustainable Archives & Library Technologies (SALT) laboratory
Reagan Moore
Director of Data and Knowledge Systems
University of California San Diego/San Diego Supercomputer Center
9500 Gilman Drive, MC 0505
La Jolla, CA 92093-0505
marciano@sdsc.edu
chienyi@sdsc.edu
moore@sdsc.edu

Mark Conrad
Archives Specialist
Electronic Records Archives (ERA)
NHE, Room 1540
The National Archives and Records Administration
8601 Adelphi Road
College Park, MD 20740-6001
mark.conrad@nara.gov

Abstract
The Persistent Archives Testbed project (PAT: 2003-2006) brought together digital curators from libraries, archives, historical societies, scientific data environments, and museums, as well as IT researchers and staff. One of the principal goals of the project was to design a distributed community repository for electronic records management. Each site chose an archival collection for testing on the infrastructure, with the record types varying from web crawls, to e-mail, to images, to voting records, to reports. We implemented a distributed data management system that allowed each institution to separately manage their own records, while using a common software infrastructure. In this architecture, each site controls access and update permissions for their preservation environment independently of the other participants. Each site implemented a different preferred interface for interacting with their archival collections and linking the distributed data to holdings at their institution. Finally, each site tested one or more archival functions - appraisal, accessioning, arrangement, description, preservation, and access – with their archival collection. The ability to manage all of these types of records using common software infrastructure was one of the significant outcomes of the project. The ability to share a common infrastructure while implementing independent archives is a second major outcome.

This paper focuses primarily on the lessons learned and skills needed by the digital curators to automate ingestion, description, and validation of records. The approaches taken by each participating institution are described, as well as the benefits that were achieved by using common infrastructure.

Introduction

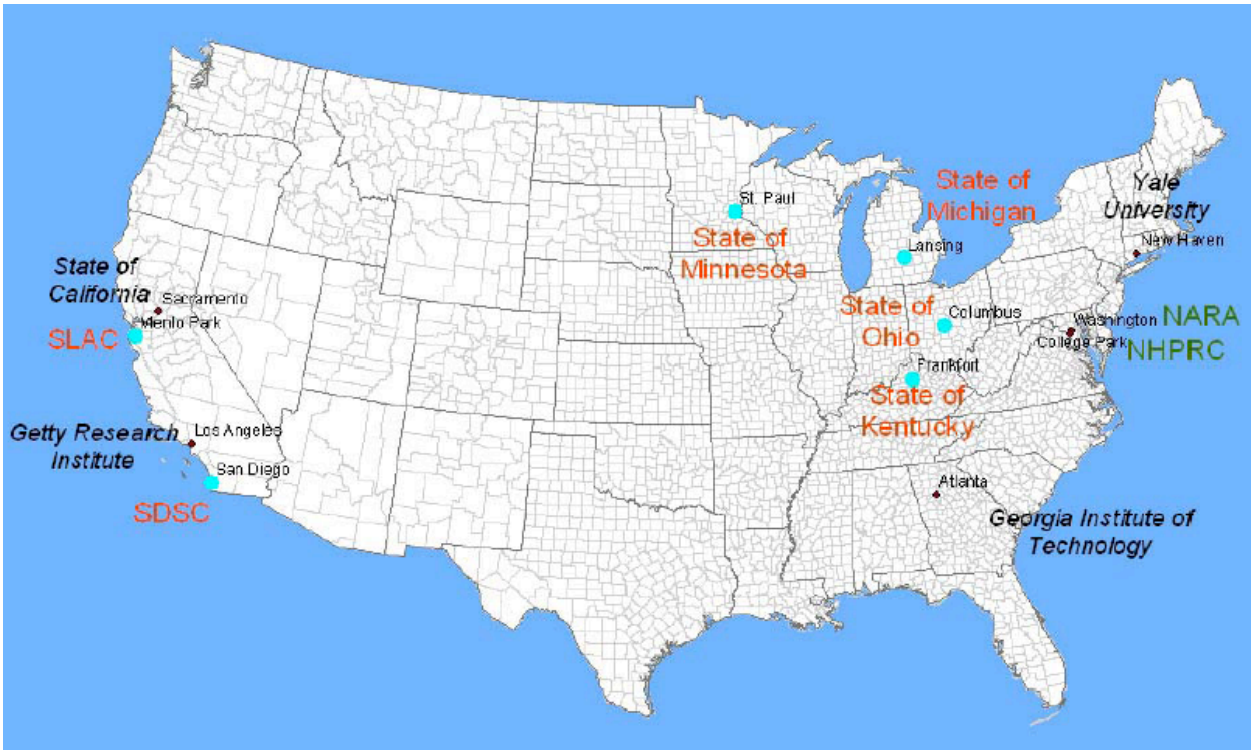

Spatial distribution of PAT collaborating institutions

One of the principal goals of the Persistent Archives Testbed (PAT) project was to explore the feasibility of a "distributed community model" for electronic records management[1] The participating institutions included the San Diego Supercomputer Center, the Stanford Linear Accelerator Archives and History Office, the Kentucky Department for Libraries and Archives, the Ohio Historical Society, the Minnesota Historical Society, the Michigan Department of History, Arts and Libraries, and the National Archives and Records Administration. Observers included researchers from the UCSD Libraries (Chris Frymann, Larry Cruse), the Getty Research Institute (Karim Boughida, David Farneth, Mahnaz Ghaznavi), the California State Archives (Linda Johnson, Renee Vincent-Finch), the California State Library (Janet Coles), the California Environmental Resources Evaluation System--CERES (David Harris), the Georgia Institute of Technology (William Underwood), and the Yale Manuscripts and Archives (Kevin Glick, Stephen Yearl, Derek Merieaux).

[1] For more information, please visit the project website, http://www.sdsc.edu/PAT

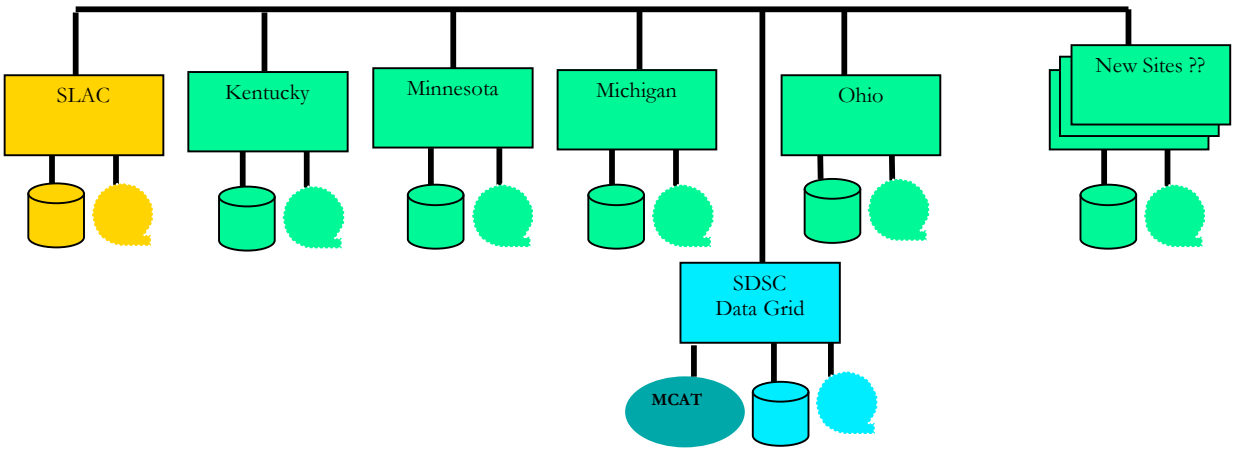

Figure 1: PAT Community Grid

The testbed was implemented using the Storage Resource Broker[2] (SRB) data grid technology developed at SDSC. As shown in Figure 1., the architecture consists of one data grid zone, one centralized metadata catalog (PAT MCAT) located at SDSC, and remote distributed SRB storage caches (SRB grid bricks) at each of the participating sites. The grid bricks were implemented using commodity disks, at an effective cost of about $2,000 per Terabyte. The archival collections for each site were replicated between their local grid brick storage and an archival repository at SDSC. Thus the architecture included the ability to manage multiple copies of the records to minimize risk of data loss, while asserting the consistency constraints needed to track chain of custody and manage authenticity and integrity. Preservation metadata were stored in the MCAT metadata catalog for each record. The metadata catalog and archival storage systems were administered by SDSC. The components of the system that were managed by each collaborating site were mainly the local SRB server and collection interface.

Each site implemented a different preferred interface for interacting with the archival collection and linking the distributed data to holdings at their institution. Each site tested one or more archival functions - appraisal, accessioning, arrangement, description, preservation, and access – with their archival collections as shown in Figure 2. The capabilities required in the SRB data grid to accomplish this included the use of a logical file name space to identify records, the organization of this logical file name space as a collection hierarchy, and the association of required provenance metadata with each sub-collection. Even though shared infrastructure was being used, each site could separately implement their preferred descriptive metadata, and separately control validation of integrity and authenticity.

	Kentucky (Web)	*Michigan (Databases)*	*Minnesota (Spatial)*	*Ohio (E-mail)*	*SLAC (e-documents)*
Appraisal				X	
Accessioning	X			X	X
Arrangement	X	X		X	X
Description	X	X	X	X	X
Preservation	X	X	X		X
Access	X	X	X		X

Figure 2: Archival Processes Explored within the Persistent Archive Testbed

[2] For more information, please visit the website, http://www.sdsc.edu/srb/index.php/Main_Page.

In this architecture, each site controls access and update permissions for their preservation environment independently of the other participants. At the same time, they are able to leverage common software and hardware resources for the management of the data and metadata. This lowers their cost of participation and reduces the level of expertise required at each site for managing the archives. In effect, each site outsourced archival repository (storage system) support issues to SDSC in order to focus on management of their archival collections (preserved records). We consider this project to be quite innovative and unique in its scope and level of interaction between archivists and IT staff. Many of the lessons learned from the project have to do with this interaction[3]

Characterization of the Processes at each Site

The processes that were used by each site varied, depending upon the types of records. The processes are loosely categorized in Figure 3.

Processes	KY	MI	MN	OH	SLAC
1. Installation of a SRB server at the sites	no	YES	YES	YES	no
2. Registration of collections into the SRB data grid using file names that are appropriate for the fonds.	YES	YES	local	YES	YES
3. Replication of digital collections onto a geographically remote storage repository for disaster recovery.	YES	YES	"yes"	YES	YES
4. Registration of metadata as part of the description process	YES				YES
5. Creation of a collection hierarchy as part of the arrangement process	YES	YES		YES	YES
6. Extraction of metadata as part of the description process	YES			YES	YES
7. Development of a validation process for assessing the authenticity of records in each collection.		YES	YES		
8. Development of a template to extract descriptive metadata and populate the MCAT catalog.				YES	YES
9. Initial automation of an archival process at each site	YES	YES	YES	YES	YES
10. Refinement of the automation of archival processes based upon the changing archival expectations derived from the work in initial automation		YES			YES
11. Development of web-based demonstrations of the results of the partial automation of archival processes	YES	YES		YES	YES

Figure 3. Preservation processes conducted at each site

[3] For a fuller discussion of these lessons learned, see: PAT Project Lessons Learned: Archivists' Perspectives
Mark Conrad, Jean Deken, Bob Horton, Richard Marciano, et al.
http://www.archivists.org/periodicals/ao_backissues/AO-Nov05.pdf (page 10+), and
PAT Project Lessons Learned, Part 2: IT Professionals' Perspectives
Mark Conrad et al. http://www.archivists.org/periodicals/ao_backissues/AO-Mar06.pdf (page 8+)

1. **Installation**: SRB storage nodes ("grid bricks") were built and shipped to MI, OH, and MN. MN's configuration stored 2 TeraBytes, whereas OH's and MI's had only 1TeraBtye of storage. KY shared a local web server host since they had minimal storage requirements. SLAC used a pre-existing SRB system (used in physics experiments).
2. **Registration**: Web harvesting processes were used at KY and SLAC to automatically register crawls into the SRB. This required the physical copying of data into the SRB data grid as well as the registration of the records into the MCAT catalog. MN registered local files directly into the MCAT catalog. Unlike other sites, this avoided the need to move the content. The registration procedure was automated.
3. **Replication**: Replication was done at all sites automatically for disaster recovery purposes. Copies of records were made at SDSC on both tape and disk storage systems. The copies were accessible through the same interfaces used to access records on the local grid bricks. MN's content was shipped through postal mail on 4 x 500GB disks. A script at SDSC was used to ingest all the content and have it appear as a "replica" of MN's content. Automated checksum computing and validation were carried out on all replicas.
4. **Metadata**: Descriptive metadata was associated with SRB ingested files and collections at SLAC, KY and OH. Each institution implemented a different set of descriptive attributes through use of the SRB "user-defined" metadata. The attributes could be queried.
5. **Arrangement**: Each record collection was separately organized as a collection folder hierarchy, uniquely designed for each site. Each site had the ability to extend the collection hierarchy, add new records, and establish soft-links between collections for re-purposed records.
6. **Metadata Extraction**: Scripts to automatically harvest and register metadata into the MCAT catalog were written and executed at OH and SLAC. This included support for bulk loading of metadata attributes.
7. **Validation**: Routines to automatically validate the authenticity of records were created at MN and MI. At MN, systematic checksum computing and comparing were done. At MI, an end-to-end workflow was written that integrated all of the legacy content into a common database that drove their web access interface.
8. **Metadata Templates**: For SLAC and OH, metadata templates and pattern matching techniques were developed to extract descriptive attributes from records and load the metadata into the MCAT catalog.
9. **Initial Automation**: Processes for recursive manipulation of directory structures are supported by the SRB data grid and were used at all sites.
10. **Refined Automation**: SLAC and MI developed multiple levels of automation, as the management policies for their collections evolved. Each new management decision was turned into a process that could be applied across the entire record collection.
12. **Web Interfaces**: Site specific interfaces were developed for MI, KY, SLAC, OH, but not for MN as they already had an access interface from their website. The types of interfaces ranged from Windows browsers, to web browsers, to Perl scripts.

As noted above, even though a common data grid was used that spanned all of the participating institutions, each site was able to implement independently the preservation processes required for accessioning their records, managing their records, and accessing their records. The ability to leverage expertise at a central location for archiving records and maintaining descriptive metadata minimized the labor requirements at each site. The ability to automate manipulation of their collections also minimized the level of effort required to manage the local preservation environment.

Unique Contributions of the Digital Curators to the Infrastructure

Feedback from the digital curators in this project led to improvements in the SRB infrastructure. Thus the project resulted in improved infrastructure that could then be leveraged by all users of the data grid infrastructure. This positive feedback also occurs with other communities that manage distributed data collections. The infrastructure used to support the preservation environments is generic software that also is applied in support of digital libraries, real-time data stream collections, and distributed data analysis platforms. Some of the specific improvements that were driven by the PAT project include:

- Windows-based SRB clients and Servers were deployed, something which had not been done in production before. This led to the creation of a Perl for Windows client library, used extensively across the project and later with NARA collections.
- Bulk operations were developed, tested, and refined. Specific operations were bulk registration and accession of records, bulk metadata extraction from records, bulk metadata loading, and bulk validation of data movement into, out of, and within the system (through use of checksums).
- End-to-end workflows were developed from accessioning of records, to replication of records.
- All SRB data grid bugs that were discovered were fixed, leading to a more reliable product.
- The MCAT catalog was ported to the mySQL open source database. This increased the number of databases that could be used to manage SRB metadata. Previously only Oracle, DB2, Sybase, and Informix could be used.
- Feedback from the project influenced the design and implementation of a wiki for posting documentation. Communication within the project used the SRB-chat e-mail list to discuss problems and experiences.
- Problems with registration of filenames with unusual characters were discovered and fixed. The character sets used by Windows, Unix, and Mac operating systems differ. Thus the naming conventions used on one operating system could cause problems when storing records on storage controlled by another type of operating system.
- Suggestions were made to simplify governance issues tied to particular types of data management. Governance policies were managed externally to the SRB software, and applied by invoking SRB operations at periodic intervals or after external events. The need for expressing such policies as rules that could be applied by the data management system itself became evident. The development of the next generation of data grid technology was initiated, in the product called iRODS (integrated Rule-Oriented Data System). The iRODS system automates the application of management policies as rules controlling the execution of remote micro-services. Each preservation process is expressed as a set of micro-services. Each micro-service is implemented as a set of operations that can be performed using a remote storage system access protocol. Each management policy is expressed as a set of rules. The results of applying the rules are stored as persistent state information that can be queried to verify assertions about the properties of the preservation environment.

Conclusions

The PAT project suggests that sustainability is probably beyond the capability of most individual archival repositories. The demands of tracking new types of technology, the expertise required to manage new technology, the costs of the storage systems and databases, and the expertise necessary to manage multiple types of storage systems were feasible with some of the sites and beyond what could be sustained for other sites. However, the PAT model strongly

suggests the feasibility of distributed multi-repository partnerships or consortia, with shared costs and responsibilities. For example, at the end of the project, an independent MCAT catalog was ported at KY on top of an open source mySQL database, allowing the partnership to be independent of SDSC. The collaboration can be easily extended and scaled to other archival repository partners, either through extension of a single data grid, or through federation of multiple independent data grids. Federation is the controlled registration of data and metadata into multiple MCAT catalogs. Federating such collaborative networks was not evaluated here, but is being pursued by NARA within the Transcontinental Persistent Archive Prototype. This system integrates multiple independent data grids located at the University of Maryland, UCSD, Renci/University of North Carolina, Georgia Tech, and NARA.

Acknowledgements
This project was partially funded by the National Historical Publications and Records Commission (NHPRC) under grant number 2004-008. The project was also partially funded by the National Archives and Records Administration (NARA). Special thanks to all the digital curators and archivists involved: Lucy Barber, Mike Meier, Pari Swift, Jean Deken, Glen McAninch, Bob Horton, and to their talented staff members who participated in this project both from the IT and digital curation perspectives.

References

1. Moore, R., M. Smith, "Automated Validation of Trusted Digital Repository Assessment Criteria", submitted to Journal of Digital Information, January 2007.
2. Smith, M., R. Moore, "Digital Archive Policies and Trusted Digital Repositories", proceedings of The 2nd International Digital Curation Conference: Digital Data Curation in Practice, November 2006, Glasgow, Scotland.
3. Moore, R., A. Rajasekar, M. Wan, W. Schroeder, R. Marciano, A. Jagatheesan, "On Building Trusted Digital Preservation Repositories," 5th e-Science All Hands Meeting, Sept. 2006, Nottingham, UK.
4. Moore, R., "Building Preservation Environments with Data Grid Technology", American Archivist, vol. 69, no. 1, pp. 139-158, July 2006.
5. Moore, R., M. Smith, "Assessment of RLG Trusted Digital Repository Requirements," JCDL on "Digital Curation & Trusted Repositories: Seeking Success", June 2006, Chapel Hill, North Carolina.
6. Rajasekar, A., M. Wan, R. Moore, W. Schroeder, "A Prototype Rule-based Distributed Data Management System", HPDC workshop on "Next Generation Distributed Data Management", May 2006, Paris, France.
7. Rajasekar, A., R. Moore, F. Berman, B. Schottlaender, "Digital Preservation Lifecycle Management for Multi-media Collections, Lecture Notes in Computer Science, vol. 3815/2005, pp. 380-384, November 2005.
8. Moore, R., R. Marciano, "Technologies for Preservation", chapter 6 in "Managing Electronic Records", edited by Julie McLeod and Catherine Hare, Facet Publishing, UK, October 2005.
9. Moore, R., R. Marciano, "Prototype Preservation Environments", Library Trends Volume 54, p. 144, Number 1 (Summer 2005), "Digital Preservation: Finding Balance", Johns Hopkins University Press.

10. Moore, R., J. JaJa, R. Chadduck, “Mitigating Risk of Data Loss in Preservation Environments”, NASA / IEEE MSST2005, Thirteenth NASA Goddard / Twenty-second IEEE Conference on Mass Storage Systems and Technologies, April 2005, pp. 39-48.
11. Moore, R., A. Rajasekar, M. Wan, “Data Grids, Digital Libraries and Persistent Archives: An Integrated Approach to Publishing, Sharing and Archiving Data”, Special Issue of the Proceedings of the IEEE on Grid Computing, Vol. 93, No.3, pp. 578-588, March 2005.
12. Baru, C., R, Moore, A. Rajasekar, M. Wan,"The SDSC Storage Resource Broker,” Proc. CASCON'98 Conference, Nov.30-Dec.3, 1998, Toronto, Canada, p. 5.