A Prototype Platform for Policy-Based Archival Replication

by **Micah Altman**[1] (Senior Research Scientist, Institute for Quantitative Social Science, Harvard University) <Micah_Altman@harvard.edu> *http://maltman.hmdc.harvard.edu*

and **Bryan Beecher** (Director of Computing & Network Services, Inter-university Consortium for Political and Social Research, University of Michigan)

and **Jonathan Crabtree** (Assistant Director for Archives and Information Technology, Odum Institute for Research in Social Science, University of North Carolina, Chapel Hill)

with **Leonid Andreev**, **Ed Bachman**, **Adam Buchbinder**, **Steve Burling**, **Patrick King**, **Marc Maynard** *http://data-pass.org*

Overview

Many threats to preservation are ameliorated through replication of the materials to be preserved by independent *institutions,* combined with regular auditing of that replication.[2] **Data-PASS** partners, as well as others who archive social science data, are prototyping a "syndicated storage" platform that would assist them in such preservation-oriented replication. This system will serve two institutional goals. First, it will help each institution insure against media, software, hardware, and physical failure, since geographically distributed partners will keep separate archival copies. Second, it will help the partnership insure against institutional failure, since if one partner should suffer institutional failure, the partnership as a whole will still retain copies of the holdings of that failed partner, and will be able to redistribute those copies. Many "single points of failure" are eliminated when the institutions involved are diversified with respect to the legal regimes and economic models under which they operate, and the technical preservation strategies that they employ.

Another institutional issue we plan to address is the asymmetrical nature of storage needs among current and potential partners. How do we construct systems that serve both the technology and the business needs for a collective when some members may require an order of magnitude more resources than others? Our institutional policies are asymmetric; some institutional members contribute more resources, or have larger holdings, than others. (see Figure 1)

At this point we have built a prototype that allows us to audit a replication network to determine whether it conforms to our stated replication policies. This prototype system is built around a core of **Private LOCKSS[3] Networks (PLN)** technology; a schema to encapsulate inter-archival replication commitments; an automated schema-driven service that audits **PLN's**; and **Open Archives (OAI-PMH)**[4] clients to harvest data collections from the **Dataverse Network**[5] **(DVN)** and other repositories using the **Data Documentation Initiative**[6] **(DDI)** schema.

This work is conducted by the **Data Preservation Alliance for the Social Sciences** (**Data-PASS**). Five major American social science data archives have created the **Data-PASS** partnership to ensure the long-term preservation of our holdings and of materials as yet un-archived**.**[7] The partners are the **Inter-university Consortium for Political and Social Research**, **The Roper Center for Public Opinion Research**, **The Howard W. Odum Institute for Research in Social Science**, the electronic records custodial division of the **National Archives and Records Administration (NARA)**; and **The Henry A. Murray Research Archive**, with strong technology support from the **Harvard-MIT Data Center**.[8] We seek to acquire and preserve data at-risk of being lost to the research community, such as opinion polls, voting records, large-scale surveys, and other social science studies; develop joint best practices for data preservation; and develop open shared infrastructure for digital preservation.[9]

Design Goals

Our design has three main goals: The first goal is to *automate policy* — the behavior of the syndicated storage system should be automatically auditable by reference to archival policy. We describe these policies formally, using a metadata schema. These formal policies include systematically describing the commitment of resources each of the archives has made to preserve the contents of the other partners; the auditing commitments each has made to its depositors; and the legal policies supporting access to the data by other partners in the case of institutional failure. The system acts on this metadata by auditing the actual state of the replication network, and reporting any deviations from policy. These auditing reports are also schematized, so that they can be used manually (now) or automatically (in future) to initiate corrective action by hosts in the network in response to deviations from policy.

Also as a matter of policy, this schematized approach is integrated with TRAC.[10] The schema can be used to document the TRAC criteria associated with particular commitments, and also to provide evidence in support of a number of criteria related to managing holdings.

A second goal of the design is to keep consistent with the model of trust already among collaborating archives: Each partner is trusted to hold the public content of the others, and not to disseminate it improperly. (Legal agreements among the archives reinforce this trust.) But, no partner is trusted to be a "super-user" and thus to arbitrarily delete content from the network. An implication of the decentralization is that harvesting assignments are completed by "request" — with approval by the local host administrator and not through a "super user" with privileges on all hosts.

The third goal is to accommodate asymmetries in the archival commitments:

- Partners may vary by policy in their commitment of storage for the syndicated storage network.
- Partners may vary by policy in the size of the holdings that they require to be replicated.

In practice, these two asymmetries imply a third — each partner need not replicate the entire network, but only a portion thereof.
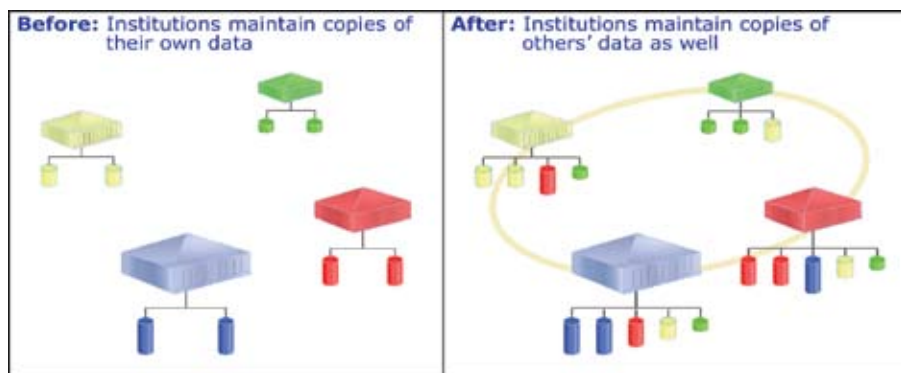

Figure 1: Conceptual Representation of Asymmetric Replication

A Prototype Platform ...
from page 44

Use Cases

Use cases describe the fundamental desired behavior of a system. We have identified two essential use cases for syndicated storage:

1. Auditing

Compare the actual state of the network to the policy schema (instance), and report any deviations, such as inadequate number of replicas, infrequency of updates, etc.

2. Recovery

a. If a host fails, and is replaced (with a new host having the same credentials), the replacement should be able to use the network to recover all lost content.

b. Any member that was able to harvest the content in the past should be able to recover a copy of content from that time, at any time in the future.

Satisfying these use cases enables each archive to be assured that archival replication policies are being maintained, and that these replications are sufficient for future recovery of content.

We are also investigating automation of the configuration and reconfiguration of the system in order to reduce the effort of managing the network. In theory, these uses cases, as well as off-the-path cases in which "requests" to a

CONTRIBUTING AUTHORS

LEONID ANDREEV is a Systems Engineer at the **Institute for Quantitative Social Science, Harvard University**.

ADAM BUCHBINDER is a Systems Engineer at the **Roper Center for Public Opinion Research, University of Connecticut**.

ED BACHMAN manages Demographic and Domestic Economic Data Services at the **Odum Institute, University of North Carolina, Chapel Hill**.

STEVE BURLING is a Systems Engineer at the **Inter-university Consortium for Political and Social Research, University of Michigan**.

PATRICK KING is IT Consultant at the **Odum Institute, University of North Carolina, Chapel Hill**.

MARC MAYNARD is Director of Technical Services at the **Roper Center for Public Opinion Research, University of Connecticut**. **Marc** serves on the Operations Committee of the **Data Preservation Alliance for the Social Sciences**, a partnership in the **Library of Congress' National Digital Information Infrastructure and Preservation Program** (NDIIPP).

host node are refused, could be automated, and would eliminate most of the manual maintenance of the network:

1. Initialization

Given a policy schema (instance), send a set of harvesting requests to each host, so that when completed the network will conform to the replication policy.

2. Add replicated collections/hosts

Add new replicated collections and hosts to the networks.

3. Grow hosts/collections

Allow hosts and collections to grow in their resource commitments. (As a design assumption we assume that resources grow monotonically.)

However, given automated auditing, and the limited number of institutions participating in the network, these cases can be readily resolved manually, through communication with the administrators of each host.

How it Works

The Syndicated Storage Platform (SSP) combines standard **LOCKSS** mechanisms; tools produced by the **LOCKSS** group for managing **PLN's**; and tools developed by the **Data-PASS** partners. An overview is shown in Figure 2:

Standard **LOCKSS** mechanisms and software are used for the harvesting framework (although like many others we have created our own plugin), integrity checks, and recovery. Harvesting is conducted by each host in the network using **LOCKSS** plugins. Hosts participate in polls of network content, and restoration occurs automatically when a host has a storage failure, or when a host fails entirely and is replaced with a new host that has the same configuration and credentials.

One should note that the limited trust model has a number of implications for implementation:

- For replication, hosts do not trust other single hosts to provide correct content, so each must be able to harvest its content directly from the source.
- For restores, hosts must have proven that they obtained the content previously, or be specially authenticated.
- No central authority can delete content from the network.
- No centralized authority can make arbitrary changes to any host, or access it through lower level (e.g., operating system) interfaces.

- The SSP also makes use of a **PLN** tool, the "cache manager," which is a tool supplied by the **LOCKSS** group for monitoring **PLN's**. The cache manager is used to gather information on the state of the network. To generate our auditing reports we run the cache manager and analyze the database it produces.

The **Data-PASS** partners have built three supplementary tools to manage the SSP. First, a commitment schema describes institutional replication requirements and commitments. Second, a schema-based auditing tool automatically audits the network against the schema. Third, a specialized harvesting plugin gathers content based on the partners catalog interoperation standards.

The schema describes the network, hosts, and collections (archival units) being replicated. At the network level, the schema describes the number of copies to be maintained and the frequency of updates. At the host level, the schema describes each participating host, and the size of resource commitment it is making available for replication. At the collection level, the schema describes the plugin to use for harvesting, the storage commitment that the organization has made to the network, and the desired frequency of harvesting.

In support of this, the SSP schema provides elements that may be used to include text, or reference external text that documents evidence of compliance with TRAC criteria. Some fields document *how the SSP itself* supports relevant TRAC criteria, while others document how the virtual organization (in this case, the **Data-PASS** partnership) responsible for the SSP is compliant with additional TRAC

criteria. Specific TRAC criteria are identified implicitly, and any TRAC criterion can be explicitly identified using schema attributes. The documentation describes each element's relevance to TRAC, and its mapping to particular TRAC criteria.

The auditing tool accesses the database maintained by the cache manager. It queries this database, aggregates the results, and compares them to the commitment schema. A report of all differences between actual and desired state is produced. This difference report can be presented in human-readable form and, more importantly, can serve as input to another stage of processing which outputs a set of changes to be made to the network.[11]

We have built a plugin to facilitate harvesting of content in our archives. Since the **Data-PASS** partners use OAI-PMH and DDI metadata to support a common catalog,[12] we have built upon this approach to replicate holdings. The DDI metadata is structural as well as descriptive, and contains links to each of the files in a research study. We have vastly updated and extended the OAI plugin so that it will handle this. In particular we have extended the plugin in the following ways:

- The plugin can now be configured to harvest a specific group of OAI "sets" corresponding to the archival units being shared by the partners.
- The plugin can now handle several sorts of authentication. HTTP "basic" authentication is now built-in, and other forms of authentication can be included via hooks to supplied libraries. Since many partners use the Dataverse Network,
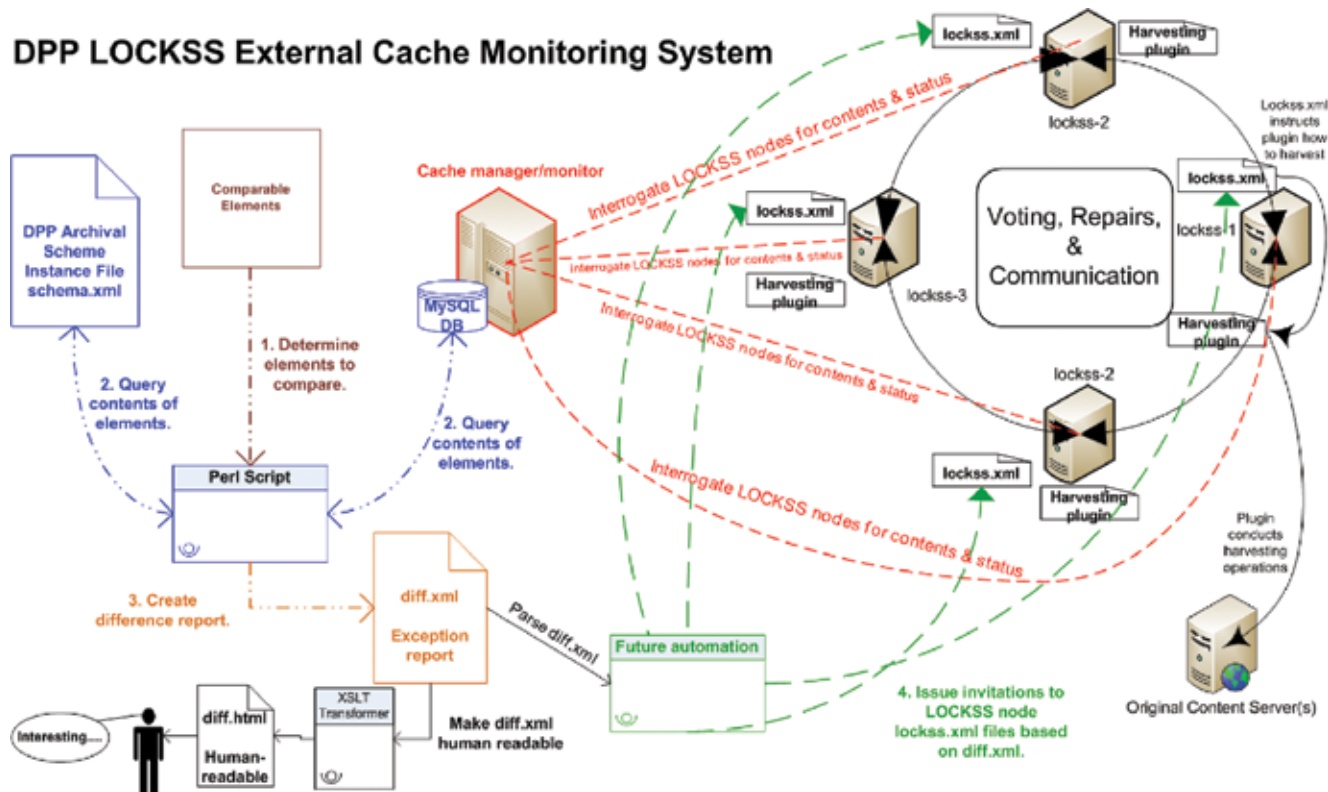
Figure 2: How the Syndicated Storage Platform Works

these hooks are designed to interface with the Dataverse libraries in order to navigate the click-through style authentication used in the DVN.

- The plugin can process metadata schemas other than OAI-DC, when an appropriate library is supplied. We have supplied a library for DDI version 2.0 metadata. This preserves the metadata, and harvests any URI's listed as resources in it. We have also included performance enhancements for handling portions of large metadata objects.

Discussion

The ability to express replication requirements and inter-archival replication commitments using a formal schema, and to automatically audit a **LOCKSS** network for dynamic consistency with these requirements is a significant advance. This provides an organization with automatic, continuous and compelling evidence that accurate, timely, and complete replicas are being maintained. Moreover this approach does not require a central administrator or homogenous configuration of the **LOCKSS** network, or create a single point of failure, either in terms of individual machines or entire institutions.

This work is in a prototype stage, and two questions remain before it could be used in production: First, the **LOCKSS** cache manager, which plays a much more prominent role in a **PLN** than it does it in the public **LOCKSS** network, is still in a "beta" stage, and in our experience, must be manually triggered regularly in order to update its state — this can trigger "false alarms" when the state of the cache manager database becomes stale, and does not reflect the actual network state. It is unclear when the cache manager will be robust enough to support automated auditing. Second, we are investigating the extent to which the **PLN** architecture can support reconfiguration of host nodes by a source that is not completely trusted.

In future work, we plan to investigate how the network might adapt automatically to changes to commitments through harvesting requests to participating hosts to perform additional harvesting. We also intend to identify ways to make the network self-repairing, so that deviations from policy commitments are repaired using the same request mechanism. That said, having the ability to audit the network against a formal policy is a useful innovation on its own. Our prototype serves as a proof-of-concept of the ability of **LOCKSS** to accommodate the institutional needs of archives as well as libraries. For more information on **Data-PASS's** approach to archival replication, including our policies and practices, software,[13] and schemas see the **Data-PASS** Website: *http://data-pass.org*. 🐾

Micah Altman, Bryan Beecher & Jonathan Crabtree

Micah Altman, Senior Research Scientist
Institute for Quantitative Social Science, Harvard University
<Micah_Altman@harvard.edu> • *http://maltman.hmdc.harvard.edu*

Bryan Beecher, Director of Computing & Network Services
Inter-university Consortium for Political and Social Research
University of Michigan

Jonathan Crabtree, Assistant Director for Archives and Information
Technology, Odum Institute for Research in Social Science
University of North Carolina, Chapel Hill

MICAH ALTMAN (Ph.D. California Institute of Technology) is Senior Research Scientist in the **Institute for Quantitative Social Science** in the Faculty of Arts and Sciences at **Harvard University**, Associate Director of the **Harvard-MIT Data Center**, and Archival Director of the **Henry A. Murray Research Archive**. **Dr. Altman** conducts research in social science informatics, social science research methodology, and American politics, focusing on the intersection of information, technology, and politics; and on the dissemination, preservation, and reliability of scientific knowledge. His current research interests include survey quality; computationally reliable and efficient statistical methods; the collection, sharing, citation and preservation of research data; the creation and analysis of networks of scientific knowledge; and computational methods of redistricting. His over thirty-five publications and five open-source software packages span informatics, statistics, computer science, political science, and other social-science disciplines.

WHERE DOES MICAH SEE THE INDUSTRY IN FIVE YEARS: The next five years will see a further application of **Melvin Kranzberg's** second law of technology: "Invention is the mother of necessity." We are now seeing a flood of new inventions for the individual and social creation, recording, and dissemination of knowledge. Information is becoming ambient — continually created, and recorded. Libraries, by necessity, will need to take a larger role in integrating with the workflows around this knowledge, and especially around tagging, social annotation, and other forms of social creation of secondary data.

BRYAN BEECHER is Director of Computing & Network Services in **ICPSR**. **Mr. Beecher** is responsible for the delivery of information technology systems to support the mission of **ICPSR**, and for the ongoing support of those systems. He supervises staff responsible for the day-to-day support of **Windows**, **Macintosh**, and **UNIX** computing systems, common business applications, and specialized data analysis applications, as well as staff responsible for designing and developing customized internal data systems, and Web-based software systems used by the member community.

JONATHAN CRABTREE is Assistant Director for Archives and Information Technology at the **Odum Institute for Research in Social Science** at **University of North Carolina, Chapel Hill**. **Crabtree** joined the institute 15 years ago and is responsible for designing and maintaining the technology infrastructure that supports the institute's wide array of services. Before moving to the social science side of campus he was an information systems technologist for the **UNC School of Medicine**. In addition to his work at the institute he is working part time on an advance degree in the **School of Information & Library Science** at **UNC**. 🐾

Endnotes for **A Prototype Platform ...** *from page 47*

1. Corresponding author, <Micah_Altman@harvard.edu>. This project was supported by an award (PA#NDP03-1) from the **Library of Congress** through its *National Digital Information Infrastructure and Preservation Program (NDIIPP)*.

2. See **David S. Rosenthal**, **Thomas Robertson**, **Tom Lipkis**, **Vicky Reich**, **Seth Morabito**. "Requirements for Digital Preservation: A Bottom-Up Approach," *D-Lib Magazine* 11 no. 11 (2005).

3. See **Victoria Reich**, and **David S. Rosenthal**, "LOCKSS (Lots Of Copies Keep Stuff Safe)," Preservation 2000, *The New Review of Academic Librarianship* 6: 155- 161 (2000).

4. See **Carl Lagoze**, **Herbert Van de Sompel**, **M. Nelson, M., & S. Warner**, "The Open Archives Initiative Protocol for Metadata Harvesting - Version 2.0," (2002). (Accessed March 24, 2009) *http://www.openarchives.org/OAI/openarchivesprotocol.html.*

5. See **Gary King** "An Introduction to the Dataverse Network as an Infrastructure for Data Sharing," *Sociological Methods and Research* 32 no 2 (2007): 173-199.

6. See **Blank**, **Grant** and **Karsten Boye Rasmussen**. The Data Documentation Initiative. The Value and Significance of a Worldwide Standard. In: *Social Science Computer Review*, Vol. 22, No. 3, 307-318 (2004).

7. For a description of **Data-PASS** collection development and its challenges, see **Myron Gutmann**, et. al, 2009, "From Preserving the Past to Preserving the Future: The **Data-PASS** Project and the challenges of preserving digital social science data." *Library Trend*s (in press).
The **Data-PASS** project Website is: *http://www.icpsr.org/DATAPASS/*.

8. Both the Harvard-MIT Data Center and the Henry A. Murray Research Archive are now part of the Institute for Quantitative Social Science, in the Faculty of Arts & Sciences at Harvard University.

9. **Data-PASS** has been, in part, funded by an award from the U.S. Library of Congress' *National Digital Information Infrastructure and Preservation Program* (NDIIPP).

10. **Data-PASS** is striving toward becoming a virtual organization conforming with preservation standards and practices, and in particular the TRAC (Trusted Repositories Audit & Certification) checklist. As such, it is a long-term goal that the virtual organization as a whole be able to demonstrate conformance with these standards, but not essential that every participating host of the SSP platform be conformant. Demonstrating conformance with these examples of digital preservation community standards and practice entails explicitly documenting the approach of a repository is addressing the requirements (mapping actions and developments to the requirements) and being able to provide evidence that the requirements are being addressed. The TRAC requirements incorporate the essential requirements of both the Trusted Digital Repositories and the OAIS documents.

11. Our model of changing network state is based on simple primitives. The tool uses the difference report to generate a set of requests of the form: *HOST_ID* [start|stop] *COLLECTION_ID* (with plugin parameters XYZ).
The early stages of this effort consist of sending the requests as e-mail messages to the administrators of the hosts requiring changes, and providing them with a tool to update their LOCKSS configuration based on the requests. We are investigating more automated approaches, however the LOCKSS PLN architecture does not currently offer hooks for automated remote management with restricted privileges, and allowing full access to automated clients is unacceptable from a security standpoint.

12. For a description of the common catalog and cataloging standards, see: **Altman**, et. al., 2009, "Digital Preservation Through Archival Collaboration," *The American Archivist,* (Forthcoming.)

13. With regard to the software used in our system, much of it is based on standard LOCKSS, or uses extensions to it, created in response to the requests from our projects and other users of PLN's. Much of the software we developed for our prototype system, such as the extensions to the harvesting plugins we describe above, has also now been contributed back to the LOCKSS project.

Pelikan's Antidisambiguation
from page 48

demographics of their audience: this detailed demographic data is likely a near-irresistible value-add to offer to the publishers in exchange for signing on to the **Kindle** distribution service.

Synchronize your page location between your **Kindle** and your **iPhone**? It's neat, I guess. Well actually, it's not really such a big deal to accomplish, but it does enrich **Amazon's** understanding of how the material they sell is consumed, where, when, over how long a period, even where, given the rudimentary GPS capabilities of the devices involved.

But this way of moving e-content around is transitional, folks. The **Kindle** is the 8-track tape player of the eBook age. I'm not saying that's bad — I'm just saying it's so.

Always remember: We like to think we're living in the Modern Age, but really we're living in the Old Days!

We're living back in the time when you had to download a book to read it — and not just that, but download it to a specific, licensed device, in a specified format, from a specific service, over a specific connection, provided by a specific vendor! (This attempt at lock-in kind of sounds like **iTunes** or the **iPhone app store**, doesn't it)?

Am I suggesting, throughout this column, that **Amazon** or **Sony** or **Google** don't deserve a mechanism for cost recovery? Certainly not! Thank goodness someone has finally achieved some traction in these arenas!

But imagine if **CNN** only let you see their Website if you used a computer you'd bought from **CNN**, using only the browser they sold you, and only over the Internet service they specified — and then made you pay by the item as well.

We're not done figuring all of this out yet, but at least we know who's paying for the R&D. 🌳

Rumors
from page 43