




Comparative Analysis of Distributed Digital Preservation (DDP) Systems

Chronicles in Preservation Project

Authors: Matt Schultz and Katherine Skinner

2 April 2014

Publication Notes

Title: Comparative Analysis of Distributed Digital Preservation (DDP) Systems

Authors: Matt Schultz and Katherine Skinner

Publisher: Educopia Institute, 1230 Peachtree Street, Suite 1900, Atlanta, GA 30309.

Copyright: 2014

This publication is covered by the following Creative Commons License:

Attribution-NonCommercial-NoDerivs 4.0

You are free to copy, distribute, and display this work under the following conditions:



Attribution – You must attribute the work in the manner specified by the author or licensor (but not in any way that suggests that they endorse you or your use of the work). Specifically, you must state that the work was originally published as the Comparative Analysis of Distributed Digital Preservation (DDP) Systems, and you must attribute the copyright holder as the Educopia Institute.



Noncommercial – You may not use this work for commercial purposes.



No Derivative Works – You may not alter, transform, or build upon this work.

Any of these conditions can be waived if you get permission from the copyright holder. Your fair use and other rights are in no way affected by the above.

The above is a human-readable summary of the full license, which is available at the following URL:

<http://creativecommons.org/licenses/by-nc-nd/4.0>

1 Introduction

The National Endowment for the Humanities (NEH)-funded *Chronicles in Preservation* project (<http://metaarchive.org/neh/>) completed this *Comparative Analysis* of three Distributed Digital Preservation systems to analyze their underlying technologies and methodologies:

- Chronopolis using iRODS (<http://chronopolis.sdsc.edu/>).
- University of North Texas using Coda (<http://www.library.unt.edu/>).
- MetaArchive Cooperative using LOCKSS (<http://metaarchive.org/>).

Chronicles in Preservation is a three-year effort to study, document, and model techniques for the preservation of digital newspaper collections. This *Comparative Analysis* is based on a series of test exchanges between academic libraries curating digital newspaper collections and the three above-mentioned DDP systems.

The workflow for the test exchanges consisted of the following steps:

- Each DDP representative met with collection curators to organize the transfer.
- The collection curators used BagIt to create a SIP for each collection.
- The collection curators staged SIPs for ingest by each DDP system.
- Each DDP system generated AIPs from all received SIPs.
- Each DDP system returned DIPs to the collection curators for verification.

While this workflow was developed specifically for digital newspaper collections, it could be adapted to many other types of content genres. In the *Comparative Analysis Discussion* section, digital newspaper exchange details are **bolded** for attention.

The systems were analyzed on the basis of nine metrics broken down further in the *table* below:

- *Ingest* – What is each systems’ typical ingest approach? How does each system ingest a SIP created by the project’s *Interoperability Tools*?¹
- *Subsequent Ingests* – How are changes or updates to AIPs handled?
- *Data Models* – How do typical workflows impact SIPs as they are transformed into AIPs?
- *Storage Environment* – How is storage technology designed to handle preservation practices after ingest?
- *Monitoring* – How are preserved AIPs audited?
- *Security* – What security technologies are employed to ensure reliable ingest, monitoring, and recovery?

¹ The project *Interoperability Tools* are a series of scripted utilities to facilitate BagIt-based exchanges. Available here: <https://github.com/metaarchive>.

- *Recovery* – How is a DIP generated? How do typical workflows impact DIPs as they are transformed from AIPs?
- *Scalability* – How does each system handle growing collections, increasing replication, and expanding organizations?
- *Costs* – How are operating costs affected by actions in the other categories?

1.1 How to Read This Document

It is important to understand at the outset that this Comparative Analysis is not intended to designate any of the Distributed Digital Preservation (DDP) systems as superior or inferior to one another in any of the areas disclosed. On the contrary, digital preservation is often best served by maintaining a variety of solutions, and each of the three DDP systems have partnered actively with one another on several digital preservation initiatives and are learning constantly from one another's approaches. The *Chronicles in Preservation* project, and more specifically, this *Comparative Analysis*, has been undertaken by these three systems in order to test, document, and refine their processes, not in isolation, but as a collaborative effort. We believe this project serves as an important model of information exchange and disclosure. Most importantly, the intent of the *Comparative Analysis* is to demonstrate that each system employs a range of methodologies—some held in common and some distinct—to accomplish DDP activities for these content types. It is hoped that readers will gain insight into the mature workings of DDP for both digital newspaper and other collections, and the project encourages other preservation systems and/or service providers to make use of this comparison methodology to share details with the broader community about their own infrastructures and workflows.

To that end, the Educopia Institute will host the open data-gathering instrument developed to facilitate the *Comparative Analysis* within the *Chronicles in Preservation* project. We invite other preservation systems and/or service providers to consider sharing details about their own methodologies using this tool. Please see the project wiki for instructions and a link to the open data-gathering instrument: [http://metaarchive.org/neh/index.php/Comparative_Analysis_of_Distributed_Digital_Preservation Frameworks](http://metaarchive.org/neh/index.php/Comparative_Analysis_of_Distributed_Digital_Preservation_Frameworks).

The *Comparative Analysis Table*, along with the corresponding *Comparative Analysis Discussion* that follows, documents all findings in the context of the digital newspaper collection exchanges that were performed between January-August 2013. All *Table* responses were supplied by the respective DDP systems. Where possible, similar terminology has been normalized across the document. Sites can be contacted if additional details or clarifications are needed or desired. See below for contact information.

Contacts

Chronopolis: David Minor, Program Manager (dminor@ucsd.edu)

MetaArchive Cooperative: Matt Schultz, Program Manager (matt.schultz@metaarchive.org)

University of North Texas Libraries: Mark Phillips, Assistant Dean for Digital Libraries
(mark.phillips@unt.edu)

2 Comparative Analysis Table

	Chronopolis	MetaArchive Cooperative	UNT-Coda
2.1 Ingest What is each system's typical ingest approach? How does each system ingest a SIP created by the project's Interoperability Tools? See Discussion			
2.1.1 Supported Ingest Methods	- Content provider staging server (via secure protocols) - Upload to local mount points - Upload to sftp server - Upload to SDSC cloud - Transfer of Bags via email	- Content provider staging server (via secure protocols) - Secure Amazon EC2/S3 instance (via scp, rsync or sftp upload) - Member-hosted staging server (via http, https, or shipped hard-drives, flash-drives, CD/DVD media)	- Transfer of Bags via removable storage media (hard-drives, flash-drives, CD and DVD) - Content provider staging server (via http, https, rsync)
2.1.2 SIP Package Support	- BagIt - Uncompressed data directories - ZIP and TAR compressed directories	- BagIt - Bag Split Utility (Interoperability Tool) - Uncompressed data directories - ZIP and TAR compressed directories (lists of per-file checksums also requested)	- BagIt - Uncompressed data directories - ZIP and TAR compressed directories
2.2 Subsequent Ingests How are changes or updates to AIPs handled? See Discussion			
2.2.1 Versioning	Yes – full Bag re-submissions	Yes (change deltas only)	Yes – full Bag re-submissions
2.2.2 Versioning Scope	- Additions - Deletions - Content changes - Filename changes	- Additions - Deletions - Content changes - Filename changes	- Additions - Deletions - Content changes - Filename changes
2.2.3 Versioning Methods	Manual uploads	- wget w/Last-Modified - bag-diff.py (Interoperability Tool)	Versioned instances of items are ingested as new object in Coda
2.2.4 Supported Versioning Schedules	When desired	2 weeks 1, 2, 6 months	When desired

	Chronopolis	MetaArchive Cooperative	UNT-Coda
2.3 Data Models How do typical workflows impact SIPs as they are transformed into AIPs? See Discussion			
2.3.1 SIP Elements	- Content provider files - Content provider checksums (if available) - Manifest (if available)	- Content provider files - Collection metadata (if available) - Content provider checksums (if available) - Content provider bag-info.txt (if available) - Archival Unit (AU) designations	- Content provider files - Content provider checksums - Content provider bag-info.txt
2.3.2 Additional AIP Elements	- Chronopolis supplied checksum (if needed) - Manifest of submitted files	- Collection manifest (LOCKSS permission page) - Collection harvester rules (Plugin) - Archival Unit Identifiers (AUIDs) - Collection identifier - Collection metadata	- METS record creation - PREMIS Object for each file - JHOVE output for each file - PREMIS event for each item ingested - bag-Info.txt file per item
2.4 Storage Environment How is storage technology designed to handle preservation practices after ingest? See Discussion			
2.4.1 Storage Technologies	Low- to medium- cost enterprise servers, provisioned differently at each storage node	Low-cost enterprise servers, SANs	Low-cost enterprise storage arrays (iSCSI currently)
2.4.2 Geographic Replication	Three managed copies of each collection (U.S.)	Seven geographically distributed copies per collection (Int'l)	Two copies in physically separated data centers (Texas)
2.4.3 Formats Accepted	All, format agnostic	All, format agnostic	All, format agnostic
2.4.4 Data Deletion Methods	- Only by request of data owner - Manual process	Retired collections are phased out at request of data owner and upon hardware retirement	- Only by request of data owner. - Manual process

Chronopolis		MetaArchive Cooperative	UNT-Coda
2.5 Monitoring			
How are preserved AIPs audited? See Discussion			
2.5.1 Monitoring Technologies	- All content verified via checksums every 4-6 weeks via ACE - All content verified via checksums separately at each node	- Content verified via ongoing checksum comparisons between distributed copies - Collection database - Member Annual Reports - LOCKSS Manifest Validator & PREMIS Event Service	- Content verified via checksums annually and upon replication - PREMIS Event Service created for replication, migration, fixity check
2.5.2 Monitoring Schedules	Content verified via checksums every 4-6 weeks via ACE	- Regularly audited by LOCKSS - Upon request - Annually	Content verified via checksums annually and upon replication
2.6 Security			
What security technologies are employed to ensure reliable ingest, monitoring, and recovery? See Discussion			
2.6.1 Ingest Security Protocols	Only via approved accounts with strong encrypted passwords	- Web server restrictions to known IP addresses (i.e., firewalls) - Via approved accounts with strong encrypted passwords	Only via approved accounts with strong encrypted passwords
2.6.2 Storage Security Protocols	- Strong access policies at each storage node, based on encompassing data center. - All data transferred via encrypted channels.	- Archive node communication encrypted and restricted to SSL peers and known IP addresses - Access Policy - Material Breach Agreements	- Access policies at each data center. - All data transferred via encrypted channels.
2.6.3 Monitoring Security Protocols	- Monitoring run at each node locally using restricted accounts. - Accounts not shared across nodes.	- Access Policy - Material Breach Agreements - Access restricted to known IPs - No central root user access	- Separate monitoring for each Coda instance. - Access restricted to known IP addresses (i.e., firewalls)
2.6.4 Recovery Security Protocols	Manual process conducted only by verified staff via secure connection methods	- Content Recovery & Access Policy - Secure temporary recovery logins - Access log reviews - Access restricted to known IPs - Material Breach Agreements	Manual process conducted only by verified staff via secure connection methods

Chronopolis		MetaArchive Cooperative		UNT-Coda
2.7 Recovery How is a DIP generated? How do typical workflows impact DIPs as they are transformed from AIPs? See Discussion				
2.7.1 DIP Elements	- Files as originally supplied - Chronopolis manifest	- Files as originally supplied - Collection metadata as originally supplied - Archival Unit (AU) designations - Collection manifest (LOCKSS permission page)	- Files as originally supplied - METS record creation - Per-file PREMIS Objects generated - Per-file JHOVE output generated - PREMIS events related to object - Updated bag-Info.txt - Coda manifest	
2.7.2 Supported Recovery Methods	Manual process tailored to data owner. Via secure methods	Secure download (see also 2.6.4)	- Manual process tailored to data owner - Large data transfer via removable storage (hard-drives, flash-drives)	
2.7.3 DIP Package Support	BagIt by default, or package supplied by content provider	- ZIP (w/ BagIt if supplied at ingest) - Bag Unsplit Utility (Interoperability Tool) - WARC	BagIt by default, custom packaging for content provider as requested	
2.8 Scalability How does each system handle growing collections, increasing replication, and expanding organizations? See Discussion				
2.8.1 Collection Growth Management	Consultation of program director with staff and node managers	- Annual Steering Committee Reviews - Storage Server/Disk Expansion Policies & Practices	- Annual UNT Libraries infrastructure support - Project based funding pooled for shared storage purchases	
2.8.2 Replication Growth Management	N/A	- Similarly sized storage infrastructure always available at seven locations - Storage Server/Disk Expansion Policies & Practices	Similarly sized storage infrastructure available at two locations	
2.8.3 Organization Growth Management	Consultation of program director with node managers	- Annual Steering Committee Reviews - Technical staff hires (as needed)	Based on growth of UNT Libraries Digital Collections	

	Chronopolis	MetaArchive Cooperative	UNT-Coda
2.9 Costs How might operating costs be affected by increases or decreases of actions in the other categories? See Discussion			
2.9.1 Ingest Costs	Currently not affected. Ingest processes are fixed with a few options, but these options don't change cost. We are working toward more automation, but any cost changes will be minimal and not seen by customers.	- Costs may increase under shipped-drive and secure FTP methods (impacts staging costs and staff time) - Costs can decrease using secure wget methods more frequently - especially via BagIt (saves staff time)	- Costs remain constant for valid SIP submission. - Costs can increase due to invalid submissions requiring manual investigation.
2.9.2 Subsequent Ingest Costs	Not affected. Theoretically could be under extreme cases of high numbers of ingests, but highly unlikely. If happened, costs would be for labor to process. We've born significant ingest events w/ no cost impact.	- Costs may increase if only semi-automating versioning (impacts staff time) - Costs decrease when automatically versioning and using BagIt Interoperability Tools (saves staff time)	Costs remain consistent with standard ingest costs for subsequent ingest and versioning.
2.9.3 Data Models	Currently only have one data model. We've worked with other data models in the past, but costs did not change.	- Costs may increase when crafting ingests from repository software systems (impacts staff development time) - Costs decrease when using BagIt Interoperability Tools (saves staff time)	- UNTL SIP Bag format establishes base costs. - Non-UNTL SIP Bag requires conversion and adds to cost
2.9.4 Storage Environment	Storage environments at the local nodes are driven by the budget available, not the other way around. Biggest chance for change would be the addition or departure of a node, which would affect overall costs. This scenario hasn't happened so far in our history.	- Costs in-line with data growth (via storage fees) - Costs can increase when network and content expands, but mitigation is built into the model via membership fees and storage fees (impacts staff time)	- Costs in-line with data growth (impacts storage costs) - Costs will increase with number of unique copies stored (impacts storage costs)

Chronopolis		MetaArchive Cooperative		UNT-Coda
Costs (continued)				
How might operating costs be affected by increases or decreases of actions in the other categories? See Discussion				
2.9.5 Monitoring	Our monitoring tool (ACE) does not have a direct cost associated with it. Only related cost is development time, but this is assumed as part of normal operating processes and an increase in the work would not equal an increase in cost. There is no relationship between number of items monitored and cost.	- Costs may increase with higher rates of request for on-demand audit (impacts staff time) - Costs will level or decrease with more frequent content provider monitoring via UI methods (saves staff time)	Monitoring provided by UI methods, which present level costs for current and planned growth	
2.9.6 Security	Currently born by local nodes as part of their default operating environment. Not a direct cost for our network.	Costs may increase under conditions of multiple concurrent collection recovery, involves multiple archive node troubleshooting (impacts staff time)	Increased number of replicas adds to potential security issues. (impacts staff time)	
2.9.7 Recovery	Same as 2.9.2 above	Costs may increase under conditions of multiple concurrent collection recovery (impacts staff time). See also 2.9.6 above	Large scale recovery and delivery of items cause increase in costs (impacts staff time)	
2.9.8 Scalability	Significant scaling events would likely impact timelines and deliverables more than cost. We assume a linear connection between scale and income, so the related costs should be flat.	Costs may increase as the number of archive node numbers expand – mitigated by increased income (impacts staff time)	- Costs increase as additional archival copies are stored. - Cost increases for current and planned growth are constant with storage increase for current master and replica copies (more data leads to higher absolute cost though constant or lower per TB cost)	

3 Comparative Analysis Discussion

3.1 Ingests

Return to [Table](#)

Each of the three DDP systems offers multiple pathways for ingesting data from a content provider. BagIt was used as a common experimental ingest packaging method for the various digital newspaper providers in the *Chronicles in Preservation* project. **All of the content providers had success packaging their digital newspaper collections as Bags.** MetaArchive Cooperative developed a Bag Split Utility that facilitated the formation of smaller archival units of BagIt content prior to ingest—this enables optimal ingest, replication, and on-going integrity checking for that DDP system. **Network transfers completed most efficiently and effectively when Bags were able to be further packaged using ZIP or TAR (uncompressed data directories) prior to transfer. This proved especially helpful when attempting to transfer Bagged digital newspaper collections that consisted of thousands of small born-digital web files.** In MetaArchive, with its use of HTTP web harvesting, any received ZIP or TAR packages can be staged, unzipped, audited with BagIt, and then ingested so that content receives per-file auditing going forward. If necessary the Bag Split Utility can be used to properly aggregate the data in order to manage network transfers involving large numbers of small files, which can sometimes cause persistent server response timeouts when using HTTP.

3.2 Subsequent Ingests

Return to [Table](#)

Digital newspaper collections are added to and remediated over time just like all digital collections. Each of the three DDP systems has developed successful strategies for absorbing and documenting changes to previously ingested collections. Chronopolis and UNT perform a full resubmission of the original Bag along with any changes that need to be accounted for. Previously ingested Bags that may contain outdated or duplicated files are aged out of the storage environment according to internal policies and schedules. MetaArchive has the capability to revisit an active source for content and ingest only the changed content. That changed content does not replace the original, but is stored alongside it as a dated version. For those collection owners that make use of BagIt, MetaArchive has also developed a Bag-Diff Utility that can scan an updated set of Bag content for changes and add only those changes to the base collection as a separate Bag of changed data. Again, this Bag is stored alongside, rather than replacing, the original version by default.

3.3 Data Models

Return to [Table](#)

Each of the three DDP systems has a range of options for SIPs. They can each receive content provider data as Bags or as uncompressed data directories on various stationary or transportable media. MetaArchive works with content providers to think through how best to package their collections as

long term archival units and works to achieve these units prior to ingest. **The use of Bags for digital newspaper collections makes the process of forming archival units very flexible.** In terms of rounding out SIPs to produce AIPs, there are a number of different data management elements that get generated and functionalized for the packages that are specific to the DDP environment itself. Chronopolis creates checksums and manifest files for collections. MetaArchive puts in place the bridge elements to accomplish the ingest (plugins and manifest pages) and the data elements for identifying, monitoring, and managing the collection. UNT performs a number of technical and structural metadata operations to create METS and PREMIS records at the object level for any received content. Each DDP system's data model is determined by the overall distributed hardware/software infrastructures that undergird the systems.

3.4 Storage Environments

Return to [Table](#)

In terms of storage technologies, the three DDP systems all intentionally make use of standard enterprise servers. It is important to note that these servers are typically sourced from multiple vendors and/or provisioned differently at the various replicated storage nodes for each of these DDP networks. Heterogeneity is a priority for each of the systems. Replication of storage nodes within Chronopolis is currently based on regional distribution across the US. MetaArchive's replication is international in extent and based on statistical auditing protocols (achieving seven distributed copies in conjunction with LOCKSS). UNT's Coda repository is replicated locally within the Denton, Texas area across multiple data centers. **Each of the DDP storage environments can readily accept digital newspaper collections.** Moreover they are format agnostic and can store data of any type. None of these infrastructures are operated by proprietary software and all have proven extensible. The storage arrays are accessible only through secure permissions. Data management tasks, including deletion, can take place as needed according to internal management and retention policies.

3.5 Monitoring

Return to [Table](#)

Each of the three DDP systems has developed monitoring technologies and capabilities that suit its organizational and technical environments. Chronopolis has leveraged its partners' programming and technology resources to deploy the ACE (Audit Control Environment) Audit Manager and the Data Portal—both of which are detailed here: <http://chronopolis.sdsc.edu/infrastructure/>. The MetaArchive Cooperative created the Conspectus (<http://conspectus.metaarchive.org>) for monitoring, which provides central staff and MetaArchive content providers with the ability to publish, describe, and monitor their collections on an on-going basis. **Thanks to the success of the digital newspaper BagIt ingests in the Chronicles in Preservation project, the Conspectus is being enhanced with an on-request audit feature that will facilitate the comparison of a Bag manifest with a corresponding set of requested checksums. In addition, the collaborations in the project have enabled the deployment of UNT's PREMIS Event Service for MetaArchive to help monitor and track ingest and replication completion and success.** UNT's Coda repository has scheduled auditing tasks and records pass/fail

events around activities such as replication, migration, and fixity, also using the PREMIS Event Service, which is now publically available here: <https://github.com/unt-libraries/django-premis-event-service>.

3.6 Security

Return to [Table](#)

Security has a number of dimensions in each of the three DDP systems, including ingest, storage, monitoring, and recovery. Due to the highly networked infrastructures, ingests and exchanges of data from content providers are handled using restricted firewalls and/or password encrypted accounts. In MetaArchive, communications for the purposes of auditing or content repairs are carried out through the use of SSL encryption. In both MetaArchive and Chronopolis, partner staff do not have root access to external archive nodes (i.e., each server in the network is under different administrative control). University of North Texas restricts access to Coda to known IP addresses within its library repository unit at the UNT Libraries. At each of Coda's three DDP sites, systems monitoring is restricted by firewalls to known application servers. Content recovery is handled through secure channels established directly with a content provider. No outside or third party intermediaries are involved. In the case of MetaArchive, these actions are governed by reference to a set of Content Recovery & Access Policies and all access and recovery situations fall under observance of a Material Breach Agreement.

3.7 Recovery

Return to [Table](#)

Recovery scenarios are described with respect to security protocols in 3.6. Upon recovery, each of the three DDP systems returns data to the content provider as it was originally provided. Chronopolis and MetaArchive are both dark archives and recovery is restricted to circumstances of data loss at the content provider's location. UNT typically takes full on-going authority control and stewardship of all content submitted to UNT Libraries for deposit into Coda, but it can return data to a content provider upon request. For both UNT and Chronopolis, preferred return package formats are Bags, but any original packaging that is present at deposit (e.g., as ZIP or TAR files) can be reconstituted. In the case of MetaArchive, standard exports are built as uncompressed ZIP packages. If BagIt and the Bag Split Utility were used to construct multiple archival unit Bags for the purposes of ingest, then the Bag Unsplit Utility can be used to reconstitute the original Bag (see 3.1).

3.8 Scalability

Return to [Table](#)

Scalability for each of the three DDP systems can be viewed with respect to collections, replication, and organizational expansion. In both Chronopolis and MetaArchive, the on-going provision of storage resources for collection growth management is handled through both scheduled and as-needed consultations with partners and management. Replication growth in Chronopolis and UNT will be maintained at its current levels for the foreseeable future. MetaArchive actively seeks membership growth, both internationally and within the US. With each new member, it extends the geographical distribution of its archive nodes to assist with replication and storage availability. For each of the three

DDP systems, the need for additional staff is evaluated on an as-needed and/or annual basis and is primarily tied to collection growth or requirements for maintaining increasing numbers of replications.

3.9 Costs

Return to [Table](#)

All of the necessary details for understanding how costs are perceived to increase/decrease for each of the above metrics under various conditions are contained in Table 2.9. The *Chronicles in Preservation* project is extremely grateful to each of the three DDP sites for their willingness to document these details—all of which are subject to change/update as time passes. Chronopolis, MetaArchive, and UNT have different and unique cost models specific to their missions and operations. The digital preservation community would benefit from further transparency in the areas of cost, and the project encourages other systems and/or service providers to publish similar data here: http://metaarchive.org/neh/index.php/Comparative_Analysis_of_Distributed_Digital_Preservation_Frameworks.