

Generating Metadata for Digital Preservation: the Chronopolis Scenario

Ardys Kozbial, UCSD Libraries, La Jolla, CA; Arwen Hutt, UCSD Libraries, La Jolla, CA; David Minor, San Diego Supercomputer Center, La Jolla, CA; Don Sutton, San Diego Supercomputer Center, La Jolla, CA; Bradley Westbrook UCSD Libraries, La Jolla, CA

Abstract

The Chronopolis Project Partners are building a digital preservation environment to manage and store data as a service to Data Providers. The system does not hold data from the partners. A subset of the project team, the Metadata Working Group made an analysis of the system by tracing a digital object through the system. System activities were divided into Event Types and metadata gathered by the digital object on its way through the system were documented.

Chronopolis Overview

Chronopolis [1] is a digital preservation environment as well as the name of a project currently funded by the National Digital Information Infrastructure Preservation Program (NDIIPP) [2] at the Library of Congress. The project is building this preservation environment within a grid-based network which holds data from institutions other than the Partner institutions. None of the data in the system belong to the partners. The Partners store and manage the data in the system as a service to the Data Providers. In the first phase of the project, the following goals were met:

- built system infrastructure at three sites (physical machines, software installation, security, software configuration);
- transferred data from the Data Providers;
- replicated transferred data across three sites;
- developed preservation services utilizing the advantages of grid-based networks;
- analyzed system data flow to determine the metadata created at various points.

The Partner institutions are UCSD (the San Diego Supercomputer Center (SDSC) and the UCSD Libraries), the University of Maryland Institute for Advanced Computer Studies (UMIACS) and the National Center for Atmospheric Research (NCAR). The Partner institutions are responsible for storage and network support, maintaining a copy of all the data at each institution, network testing, metadata expertise and support for the Storage Resource Broker (SRB) [3], the Replication Monitor [4], the Audit Control Environment (ACE) [5].

Data Providers

The Data Providers for Chronopolis include the California Digital Library (CDL) [6], the Inter-university Consortium for Political and Social Research (ICPSR) [7] at the University of Michigan, North Carolina State University (NCSU) [8], the Scripps Institution of Oceanography (SIO) [9] at UCSD and the MetaArchive Cooperative [10]. All of the Data Providers are also

NDIIPP Partners and all of the data ingested into Chronopolis are related to NDIIPP projects.

CDL, a department of the University of California's Office of the President (UCOP), provides centralized support for digital initiatives that serve all of the libraries in the University of California system. CDL contributed 5 terabytes of data to Chronopolis from its Web-at-Risk project, which has been comprised of web crawls of political and governmental web sites over the course of five years. The web crawler packages the data into files of uniform size. CDL anticipates transferring an additional 10-15 terabytes of web crawl data to Chronopolis in the spring of 2009.

Chronopolis is managing ICPSR's complete holdings, consisting of approximately 8 terabytes of data. This collection includes 40 years of social science research data comprised of millions of small files.

NCSU's data include approximately 6 terabytes of state and local geospatial data that were collected under the auspices of the North Carolina Geospatial Data Archiving Project, one of the initial eight NDIIPP projects. NCSU is also part of NDIIPP's new multistate effort, which is keenly interested in exchange of digital content among states.

SIO's approximately 1 terabyte of data are made up of data gathered from approximately 1,000 SIO research expeditions during the past 50 years. SIO was able to combine these data into one place with the help of a Digital Archiving (DigArch) research grant from NDIIPP.

The Chronopolis technical team is working with the MetaArchive Cooperative to transfer approximately 8 terabytes of cultural heritage data. MetaArchive is a digital preservation environment which uses a different infrastructure from Chronopolis. Figuring out a variety of methods of data exchange between the systems is one of the next steps for both projects.

By the end of spring 2009, Chronopolis will be nearly filled to its 50 terabyte capacity allotted in the first phase. These data present themselves in a wide variety of file formats, and the content includes web crawls, geospatial data, social science data atmospheric/oceanographic data and cultural heritage data. The Chronopolis Partners purposely solicited a diverse set of data content and types in order to develop and test Chronopolis' capacity to manage it efficiently and reliably.

Metadata Working Group

The Metadata Working Group, consisting of the metadata experts from the Libraries, members of the technical teams from SDSC and UMIACS and the project managers, was charged with

Chronopolis System Data Flow: Metadata Generation

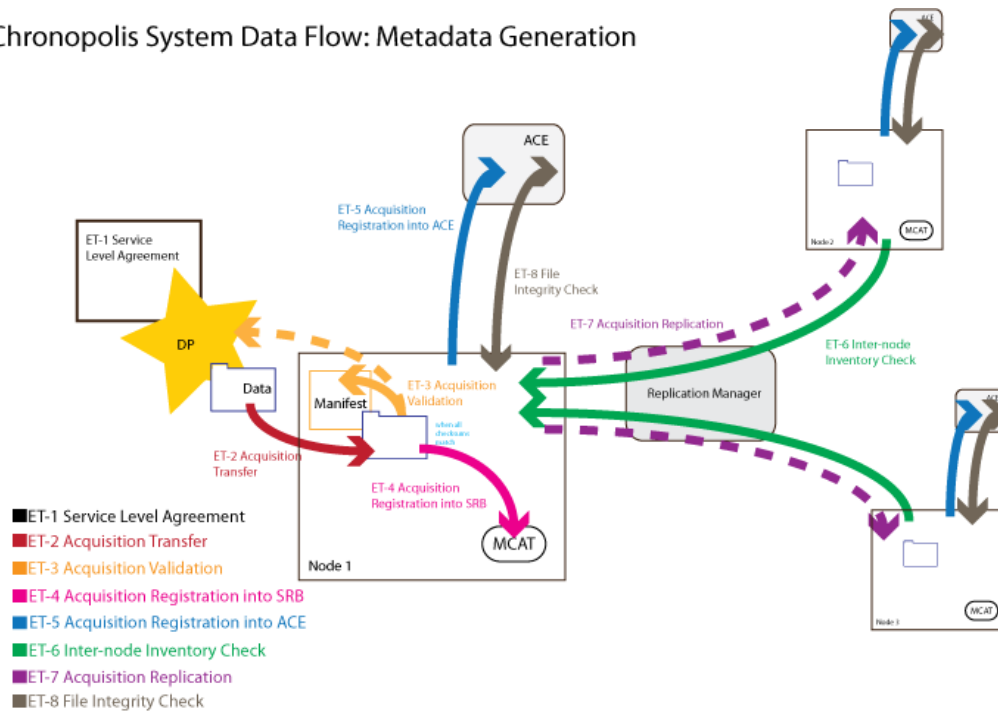


Figure 1

Chronopolis System Data Flow: Metadata Generation

developing metadata specifications that meet the following requirements:

1. Support Chronopolis' first phase services:
 - a. replicate files in multiple and geographically dispersed locations,
 - b. monitor the files regularly to identify non-authentic files,
 - c. develop mechanisms for replacing non-authentic files
 - d. deliver files back to the Data Provider upon request;
2. Be conformant with community metadata standards
3. Be extensible to support future development of services
4. Promote trust between the data providers and Chronopolis.

While the working group has yet to complete a specification, it has made an analysis of Chronopolis, determining which applicable metadata are created at certain places as the digital object travels through the system.

Methodology for Analysis

The group started with a series of questions and assumptions related to the requirements.

- Data Providers who transfer data into the system will eventually want to get it out. For example, a Data Provider's server fails and it needs to ingest the preserved data back into its system.

- Which metadata will assert that the data put into the system are the same data coming out?
- How does Chronopolis prove to Data Providers that the data being returned to them are unchanged from those ingested?
- How does Chronopolis monitor the health of its system in a transparent fashion?

The group began by tracking a single digital object through the system. As the object travelled, the group divided each step into an Event Type. Metadata were generated during each Event Type by either a human being (project manager, technical person) or a system (SRB, ACE, the Replication Monitor).

The group made a diagram of the data flow (Figure 1) and collected metadata for each Event Type in a spreadsheet. Currently the group is fairly settled on the Event Types, but is still in discussion, paring down the full metadata element list to those elements that support the requirements.

Finally, the metadata elements that do not currently exist in the system will be added by technical staff.

Event Types

The group divided the data flow into eight event types: Service Level Agreement, Acquisition Transfer, Acquisition Validation, Acquisition Registration into SRB, Acquisition Registration into ACE, Inter-node Inventory Check, Acquisition Replication and File Integrity Check. As the diagram shows, these events do not happen in a linear fashion.

Since the element list is long and in flux, for this paper, a subset, or sample, of the metadata elements will be listed.

The UCSD Libraries and SDSC have considerable experience collaborating on projects and both institutions have become adept at translating between library and computer science vocabulary. That said it is well worth noting that the terminology used in this diagram was achieved with lengthy debate and compromise.

Event Type 1: Service Level Agreement

The Service Level Agreement occurs outside the Chronopolis system. This is the point at which the Data Provider expresses interest in transferring data to Chronopolis. The Project Manager collects information about the data as well as contact information from the Data Provider and negotiates a Service Level Agreement.

It is important at this stage to determine the size of the submission(s), as well as whether it is composed of a few very large files or a great number of small files because both of these affect the methods of transfer. Since the collections being transferred are extremely diverse, careful attention must be given to the different storage locations and their specifications.

Sample Metadata

Object Type	Element Description
Service agreement type	Services provided (default/non)
	Non default services
	Annual fee
	Duration of contract
	Liability / responsibilities
	Rights / permissions
	Computer account needed
Contact information type	First name
	Last name
	Telephone contact
	Email address

Event Type 2: Acquisition Transfer

The Acquisition Transfer is negotiated by technical staff from the Data Provider and SDSC and the transfer protocol is decided at this point. (Technically, data could enter through SDSC, UMIACS or NCAR, but it is current Chronopolis policy to bring in all data through SDSC.) Two methods used by Chronopolis are an SRB to SRB transfer and BagIt.

ICPSR and SIO both use SRB as their storage system and transferring data was done using SRB commands to create a new collection on servers at SDSC. Data are moved into this new SRB instance.

The BagIt standard, recently developed by CDL and the Library of Congress to "simplify large scale data transfers between cultural institutions" [11] is the submission format used for deposit by both CDL and NCSU and it accounts for 11 terabytes of the data stored in Chronopolis. The BagIt standard is a simple format for transferring digital content focused on the core necessities for efficient and verifiable data transfer. As such, it allows packaging of digital objects with a small amount of accompanying metadata. The core of this metadata is an inventory of content files and a checksum value for each file. It is possible to point to content files

via URLs instead of packaging them 'within' the bag. This configuration is referred to as a 'holey bag' and is an example of a deposit which consists both of pushed content (files within the bag) and pulled content (files which are retrieved via URLs).

Sample Metadata

Object Type	Element Description
Acquisition information type	Name of collection or BagIt
	Path pointer to manifest file
	Date of transfer
	URL or IP of collection source
	Transfer speed
	Comments on transfer process

Event Type 3: Acquisition Validation

Regardless of the method by which content is transferred, all files are placed in SRB and are subject to an initial audit to assess how complete the transfer was and if all files transferred without corruption. This is done by comparing the transferred files to the manifest to verify that all files were received, and by calculating the checksum value for the file and comparing it to the checksum value calculated before transfer. These quality control procedures allow the identification of any corrupted transfers or missing files. The Data Provider can then be notified and the appropriate action(s) can be taken.

Sample Metadata

Object Type	Element Description
Validation information type	Name of collection or BagIt
	Path pointer to validation file
	Validation (checksum) method
	Date of validation
	Comments on validation process

Event Type 4: Acquisition Registration into SRB

After validation is complete, the data can be registered into SRB, the storage management system used by Chronopolis. The data are transferred from the Data Provider with metadata from the Data Provider's system. These metadata are stored in the MCAT, which is SRB's metadata catalog. They are supplemented with the metadata generated in Event Types 1-3 as well as SRB system level metadata required for the management of that object.

In the context of the Event Types presented here, most of the metadata elements are created during this Event.

Sample Metadata

Object Type	Element Description
MCAT register type	Internal data ID
	Data name
	Physical path name of data object
	Data modification time stamp

	Size of data
	Data creator name
	Position of data in container
MCAT physical resource type	Zone location ID
	Physical resource owner
	Physical resource class
	Physical resource lock type
	Physical resource free space size
MCAT collection type	Internal collection ID
	Parent collection name
	Collection creation timestamp
	Collection lock type
	Collection link

Event Type 5: Acquisition Registration into ACE

Regular, scheduled audits are necessary as access to files is infrequent within an archive of this type and cannot be relied upon for uncovering problems. Auditing allows the identification of data loss in a timely manner so that action can be taken to repair or replace the damaged object. Within Chronopolis audit is done using ACE.

ACE is a policy driven environment for verifying the integrity of a digital repository's holdings. ACE provides a two-tiered approach to integrity management. The first tier includes Integrity Tokens and Cryptographic Summary Information (CSI), and the second tier Witness values. An important characteristic of ACE is that it is run independently of the digital repository, which reduces the chance that a malicious file modification can go undetected since verification information will need to be changed in two independent, and independently administered, systems.

A file must first be registered with ACE. On this registration a token is created which documents integrity information for the file. This, in concert with the CSI and Witness values, is used to conduct regular evaluations of a file, and the digital repository's integrity.

Sample Metadata

Object Type	Element Name
ACE collection type	ID
	STORAGE
	CHECKPERIOD
	DIRECTORY
	PROXYDATA
ACE log event type	ID
	SESSION
	PATH
	DESCRIPTION
	DATE
	LOGTYPE
	COLLECTION_ID

Event Type 6: Inter-node Inventory Check

The Inter-node Inventory Check is where data replication happens. For a new data set in Chronopolis, the SDSC SRB looks

at the data in the UMIACS SRB and the NCAR SRB. If the data do not exist SRB replicates the data at each site, triggering the process and metadata described in Event Type 4.

Subsequent replication checks generate the actions and metadata found in Event Type 7.

Event Type 7: Acquisition Replication

This process is facilitated in part by the Replication Monitor, a tool developed at UMIACS. The tool automatically synchronizes collections between master and mirror sites and logs any actions or anomalies. The Replication Monitor is a tool built on top of the SRB and is a simple web application that watches designated SRB directories and ensures that copies exist at designated mirrors. The monitor stores enough information to know if files have been removed from the master site and when the last time a file was seen. In addition any action that the application takes on files is logged.

Sample Metadata

Object Type	Element Description
MCAT replication type	Location of the file (relative path)
	Parent path
	What type of item
	Date the item was last successfully seen intact
	Which collection the file is part of
	Current state of the item

Event Type 8: File Integrity Check

The File Integrity Check is similar in nature to Event Type 6 (Inter-node Inventory Check) except using ACE instead of the Replication Monitor. The Event itself is ACE making its audit. If checksums match and everything in the system is fine, the only metadata generated are the date of the check and an "OK." If the files are not fine, metadata from Event Type are generated.

Conclusion

The next steps for the Chronopolis team in terms of metadata are: adding the elements that do not exist in the systems (SRB, ACE, Replication Monitor) to them, paring the metadata down to those essential for proving the authenticity of the data and creating a metadata element set or a schema that can be used when transferring data from Chronopolis back to the Data Providers.

This is clearly a work in progress and the project team is very interested in questions, feedback and considerations that might have been missed.

References

- [1] Chronopolis Digital Preservation Project. <http://chronopolis.sdsc.edu/>
- [2] National Digital Information Infrastructure Preservation Program (NDIIPP). <http://www.digitalpreservation.gov/>
- [3] The Storage Resource Broker. http://www.sdsc.edu/srb/index.php/Main_Page
- [4] Replication Monitor. <https://wiki.umiacs.umd.edu/adapt/index.php/Replication:Main>
- [5] Audit Control Environment (ACE). <https://wiki.umiacs.umd.edu/adapt/index.php/Ace>
- [6] California Digital Library. <http://www.cdlib.org/>
- [7] Inter-university Consortium for Political and Social Research. <http://www.icpsr.umich.edu/>
- [8] North Carolina State University. <http://www.lib.ncsu.edu/>
- [9] Scripps Institution of Oceanography. <http://www.sio.ucsd.edu/>

- [10] MetaArchive Cooperative. <http://www.metaarchive.org/>
- [11] BagIt specification. <http://www.cdlib.org/inside/diglib/bagit/bagitspec.html>

Author Biography

David Minor (Head of Curation Services, SDSC) and Ardys Kozbial (Technology Outreach Librarian, UCSD Libraries) are the Project Managers for Chronopolis. Brad Westbrook is Head of the Metadata Analysis and Specifications Unit (MASU) in the Metadata Services Department of the UCSD Libraries. Arwen Hutt is a Metadata Librarian in MASU in the Metadata Services Department of the UCSD Libraries. Don Sutton is the Data Services Administrator on the Chronopolis Project at SDSC.