Developing Preservation Metadata for Use in Grid-Based Preservation Systems

Arwen Hutt*, Brad Westbrook*, Ardys Kozbial*, Robert McDonald, Don Sutton*****

*UCSD Libraries
University of California, San Diego
9500 Gilman Drive, #0175
La Jolla, CA 92037
ahutt@ucsd.edu,
bdwestbrook@ucsd.edu,
akozbial@ucsd.edu

** IUB Libraries
Indiana University Bloomington
1320 East 10th St.
Bloomington, IN 47405
robert@indiana.edu

*** San Diego Super Computer Center
University of California, San Diego
9500 Gilman Drive, #0505
La Jolla, CA 92037
minor@sdsc.edu, suttond@sdsc.edu

Abstract

Establishing metadata requirements is a key challenge for any attempt to implement a digital preservation repository. The repository's capacity to provide cost-effective, trustworthy services largely derives from the metadata it uses. This paper describes the metadata posited to support services the Chronopolis preservation system will offer at the conclusion of its first year of development.

Chronopolis Overview

The Chronopolis Digital Preservation Framework [1] is a collaborative partnership between the San Diego Supercomputer Center (SDSC), the University of California, San Diego, Libraries; (UCSDL), The National Center for Atmospheric Research (NCAR), and The University of Maryland Institute for Advanced Computer Studies (UMIACS) to establish a digital preservation system within a grid-based network. During the 2008-09 fiscal year, The Library of Congress' National Digital Information Infrastructure and Preservation Program (NDIIPP)[2] awarded funding to the Chronopolis Consortium to build a demonstration preservation data grid containing up to 50 terabytes of heterogeneous data at each Chronopolis node (SDSC, NCAR, UMIACS). The long term goal is to develop a trustworthy digital preservation system offering a spectrum of reliable services to data producers. Short term goals for the first and current development phase include:

- build system infrastructure at three sites (physical machines, software installation, security, software configuration)
- transfer data from depositors
- replicate acquired data across three sites
- develop preservation services utilizing advantages of grid-based networks
- define metadata required to satisfy services

Services

In this first phase Chronopolis project staff is developing basic archiving services, chief of which are:

1. provide replication of files in multiple and geographically dispersed locations
2. provide regular monitoring to identify non-authentic files
3. develop mechanisms for replacing non-authentic files
4. deliver files back to the depositor on request

During this current phase, the team will not implement any of the following services:

1. allow modification of files on our servers
2. provide end user access
3. validate and / or migrate file formats

From a depositor perspective, Chronopolis will provide a data archive that will protect against data loss due to bit decay, system malfunction, natural disaster and vandalism. This will be accomplished by using replication and redundant storage techniques in a grid environment.

Data providers

Data providers for the Chronopolis project include the California Digital Library (CDL), the Inter-university Consortium for Political and Social Research (ICPSR) at the University of Michigan, North Carolina State University (NCSU) and the Scripps Institution of Oceanography (SIO) at UCSD. All of the data providers are also NDIIPP Partners and the data being ingested into Chronopolis are related to other 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 6 terabytes of data to Chronopolis from its Web-at-Risk project, which has been composed 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.

ICPSR is submitting its whole collection of data, consisting of approximately 12 terabytes of data. This

collection includes 40 years of social science research data comprised of millions of small files.

NCSU's data in Chronopolis include approximately 5 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 2 terabytes 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 cumulative amount of digital content transferred to Chronopolis' custody is approximately 25 terabytes. These data present themselves in a wide variety of file formats, and the content includes web crawls, geospatial data, social science data and atmospheric/oceanographic data. The Chronopolis team 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 was charged with developing metadata specifications for the first phase of Chronopolis development. These metadata specifications have several requirements, they must:

1. Support the services Chronopolis implements in its first phase.
2. Be conformant with community metadata standards.
3. Be extensible to support future development of services.
4. Promote trust between the customer and Chronopolis.

Metadata requirements have been established by working back from services to the events that trigger the services, which is discussed more fully in the ensuing sections.

Workflow Events & Associated Metadata

While it is anticipated that more services will be added in the future, the workflow currently in place within Chronopolis is, in broad strokes, one which the project team expects to follow going forward. The essential stages to the present system are ingest, replication, asset management, and asset retrieval (i.e., delivery back to the customer). These represent broad areas of an object's life cycle, as well as the rudimentary stages of the Reference Model for an Open Archival Information System (OAIS)[3]. A representation of the current system work flow is provided in Figure 1.

Figure 1:

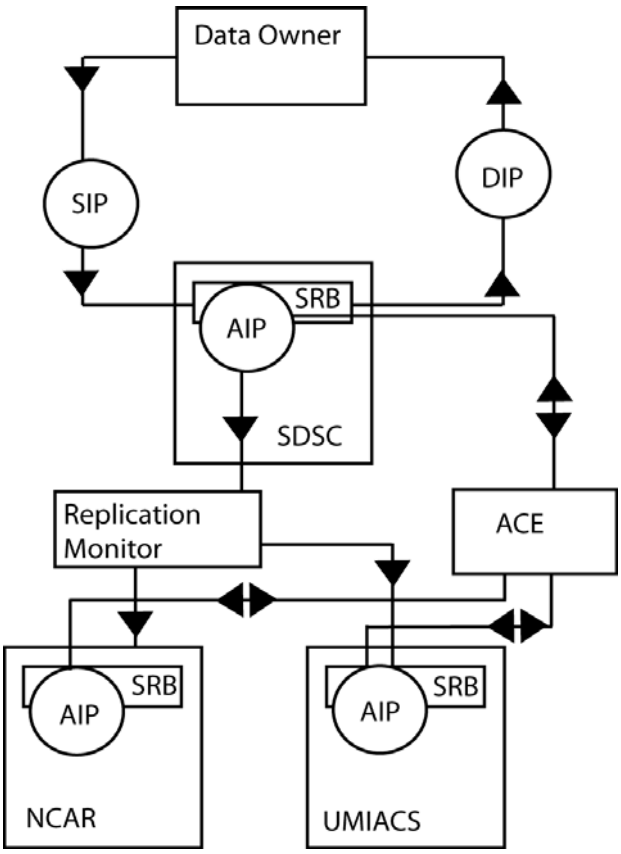

Ingest

Pre-Ingest

Within the Chronopolis project, pre-ingest has required determining the materials to be deposited, agreeing on the format(s) in which they will be submitted, and establishing secure and efficient transfer mechanisms. In addition, part of the pre-ingest process entails configuring the Storage Resource Broker (SRB)[4], the data grid management system used within Chronopolis, to host the submitted content. This involves establishing a collection or hierarchy of collections associated with the depositor.

A characteristic of the current pre-ingest process is that, beyond a few very core pieces of information, there is no need for submitted data to be compliant to a standardized Submission Information Package (SIP)[3] stipulated by Chronopolis. This imposes a non-trivial burden on the Chronopolis system, as there is less control over the form of the submission, as well as on the presence of necessary metadata. The absence of a standardized SIP results in the need to normalize or, where necessary, create the core metadata to enable Chronopolis management services. Without such actions, the Chronopolis system would need to define processes to manage each submitted collection individually, a position that is obviously not scalable or sustainable. That said, the absence of a standardized SIP reduces what might otherwise be a significant barrier for many potential customers. Certainly, its absence enabled the

import of a significant quantity of diverse data within a relatively short time frame.

Metadata
- Depositor name
- Collection name
- Collection structure

Transfer
The process of transferring data is an important component of the overall workflow. 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, affects the methods of transfer. Since, as described above, the collections being deposited are extremely diverse, careful attention must be given to the different storage locations and their specifications.

In addition, the submissions have varied in not only their transfer form, but whether the objects are deposited via a push method by the depositor, a pull method by the repository, or a combination of the two. The BagIt standard [5] recently developed by CDL and the Library of Congress, to "simplify large scale data transfers between cultural institutions" [6] is the submission format used for deposit by both CDL and NCSU, and it accounts for 17 of the 25 terabytes deposited 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 also 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).

Metadata

- File location (when files are transferred via a pull mechanism)

Verification
Regardless of the method by which content is transferred, all files are placed within the staging SRB instance 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.

Metadata
- Original file identifier
- Number of files in the collection
- Size of file
- Checksum algorithm
- Checksum for file

Registration
Once this quality assurance has been accomplished, files are registered within the receiving SRB instance. In most cases metadata is stored within the MCAT, the database system for managing the SRB, but it is not stored as a first class object, like the primary content files themselves. The deposited file's associated MCAT record is supplemented with system level data required for the management of that object, resulting in the creation of the Archival Information Package (AIP)[3], the object to be managed over time.

Metadata
- File identifier
- Date of deposit
- User who uploaded the file(s)
- User's associated group
- Size of file
- Checksum algorithm
- Checksum for file
- Resource where file is stored (information needed so SRB knows how to talk to the resource)
    - Type of resource (e.g., disc, tape)
    - OS resource is running
    - IP address of resource

Archival Storage

There are a number of threats posed to the long term preservation of digital objects. It is possible for problems to be introduced during a processing event, such as transfer to a repository, migration to new media or even delivery back to the data depositor. Failures of media or hardware can cause data loss. Natural disasters can cause catastrophic data loss for an entire repository. And either through error or malicious attack, humans can threaten the integrity of digital objects. There are two important components of protecting digital objects from all of these threats--replication and auditing.

Replication
The Chronopolis Network supports two levels of replication; replication between nodes of the network, also called mirroring, and replication within each node.
At present, mirroring between the Network partners provides copies of archived data in three dispersed geographic regions within the United States (the West Coast, East Coast and Rocky Mountains). This level of replication provides protection against data loss through natural disaster, large scale media or hardware failure and human error or attack. Mirroring occurs after ingest is complete, when the AIP is replicated at the other nodes within the network. This process then requires an additional round of quality assurance auditing to insure that all files are present and uncorrupted, and modification of some system level metadata to reflect the content's presence at the replicated node.

In addition, each node can create local replicas of the content managed within the SRB infrastructure. This local redundancy could provide a more efficient protection against data loss due to communication errors in transfer to new media and / or limited media or hardware failure.

This process is facilitated in part by the Replication Monitor[7], a tool developed at the University of Maryland. 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.

Metadata

Data which will match that of the SRB/MCAT from which the data is being replicated from

- Size of file
- Checksum for file
- Checksum algorithm
- Number of files in the collection

Data which will be unique within each node

- Resource where file is stored (information needed so SRB knows how to talk to the resource)
    - Type of resource (e.g., disc, tape)
    - OS resource is running
    - IP address of resource

Data related to replicas

- Date of replication
- File replicated from (node and resource location)
- File replicated to (node and resource location)

Auditing

The second component of archival storage is regular and ongoing monitoring of the files to identify any errors or failures. Regular, scheduled audits are necessary as depositor access to files is infrequent within an archive of this type and so cannot be relied upon for uncovering problems. Auditing allows the identification of data loss in a timely manner so action can be taken to repair or replace the damaged object.

Within Chronopolis this is being done using the Auditing Control Environment (ACE), also developed at the UMD. ACE is a policy driven environment for verifying the integrity of an archives' 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 (See [8][9] for more information about ACE). An important characteristic of ACE is that it is run independently of the archive, 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 an archive's, integrity.

Metadata

- Checksum for file
- Version number
- Checksum algorithm
- Last integrity token
- Time stamp
- Aggregation proof
- Last summary information

Dissemination

Within the current project it is expected that Chronopolis will be able to deliver materials back to the depositor in the same form as they were initially submitted. Additionally, Preservation Description Information (PDI)[10] will be provided to document the authenticity of the files. These deliverables will constitute the content of the Dissemination Information Package (DIP)[3].

Metadata

For file submitted

- Size of file
- Checksum algorithm
- Checksum for file

For file returned

- Size of file
- Checksum algorithm
- Checksum for file
- Audit trail documenting events in file's history
    - Deposit
    - Replication
    - Verification
    - Recovery (with a replica when a verification fails)
    - Dissemination

Metadata Packages

Work is now progressing on development of metadata specifications for the AIP and two DIPs. These are focused on documentation of metadata to be collected, created and retained.

As described above, the AIP is composed of metadata elements contributed by the depositor and created by SRB, ACE or the Replication Monitor. These elements are primarily stored within the MCAT database, but also depend upon data within ACE, and so the AIP is not truly a single 'package' in a physical sense, but a logical

one. The system dependencies and distributed nature of the AIP data, necessitates that reference is made to the internal metadata elements for the relevant systems, not that encoding of AIP metadata elements according to an external standards, such as the PREservation Metadata: Implementation Strategies (PREMIS) standard [11], is needed. But while *encoding* according to an external standard is not appropriate, indicating how the AIP metadata meets the requirements established by the community is. Within this context PREMIS is important for its detailed treatment of the metadata elements needed for preservation management, and its grounding within the OAIS framework.

In contrast, it is expected that the Preservation Description Information portion of the DIP will be expressed according to the PREMIS data dictionary and schemas. This package will contain much of the same data elements which make up the AIP, although there will be some variance between the data in the two packages. The DIP must thoroughly document the provenance of the digital object from its ingest into the repository to its dissemination to the depositor.

A mapping of Dissemination Information Package metadata for a file to PREMIS is presented in the chart in Figure 2. It should be noted that this includes the primary metadata which supports Chronopolis services as outlined thus far; it is not intended to be exhaustive of all elements which would be present in a DIP.

Figure 2:

DIP Metadata	PREMIS Elements
Object Entities	
Collection name	linkingIntellectualEntityIdentifier
Original file ID	originalName
Size of file	size
Checksum (pre-ingest)	messageDigestAlgorithm messageDigest messageDigestOriginator=Depositor
Checksum (post-ingest)	messageDigestAlgorithm messageDigest messageDigestOriginator=Repository
Cryptographic Summary	messageDigestAlgorithm messageDigest messageDigestOriginator=Audit control software
File identifier	objectIdentifier
Resource Type	storageMedium
Resource IP	contentLocation
Replica of file	relationshipType=replication relationshipSubType=is equal relatedObjectIdentification
Agent Entities	
Depositor	agentIdentifier agentName agentType=organization
Repository	agentIdentifier agentName agentType=organization
Networked repository	agentIdentifier agentName agentType=organization
Replication monitor software	agentIdentifier agentName agentType=software
Audit control software	agentIdentifier agentName agentType=software
Initiator of file recovery	agentIdentifier agentName agentType=person
Event Entities	
Deposit	eventType=ingestion eventDateTime eventOutcomeInformation
Replication	eventType=replication eventDateTime eventOutcomeInformation
Verification	eventType=fixity check eventDateTime eventOutcomeInformation
Recovery	eventType=replacement eventDateTime eventOutcomeInformation
Dissemination	eventType=dissemination eventDateTime eventOutcomeInformation

Development of the DIP specifications will build on work done during a previous NDIIPP project, Data Center for Library of Congress Digital Holdings: A Pilot Project, a one-year demonstration project to test the feasibility of engaging external partners as service providers to fill digital management needs. During this project, a prototype DIP for transferring preservation responsibility for an object was developed. Chronopolis will expand on that work by modeling an encoding for a more complete audit trail, including representation of mirrored sites, exploring other package formats, and updating the mapping to comply with the recently released PREMIS 2.0 [12].

Conclusion

Implementing a federated digital preservation repository network has required us to closely examine the services to be supported and what metadata is needed to enable them. It is expected that this first phase of development will provide a strong technological, policy and trust foundation upon which Chronopolis can build.

References

[1] Chronopolis Digital Preservation Framework. http://chronopolis.sdsc.edu/

[2] The Library of Congress' National Digital Information Infrastructure and Preservation Program. http://www.digitalpreservation.gov/

[3] Consultative Committee for Space Data Systems. 2002. Reference Model for an Open Archival Information System (OAIS). http://public.ccsds.org/publications/archive/650x0b1.pdf

[4] Storage Resource Broker. http://www.sdsc.edu/srb/index.php/Main_Page

[5] Kunze, J.; Littman, J. and Madden, L. 2008. The BagIt File Package Format (V0.95) http://www.cdlib.org/inside/diglib/bagit/bagitspec.html

[6] The Library of Congress. 2008. Library Develops Format for Transferring Digital Content, News and Events. http://www.digitalpreservation.gov/news/2008/20080602news_article_bagit.html

[7] SRB Replication Monitor V2.0. http://narawiki.umiacs.umd.edu/twiki/bin/view/Main/SrbRepMon2

[8] ACE: Audit Control Environment.http://narawiki.umiacs.umd.edu/twiki/bin/view/Main/ACEOverview

[9] Song, S. and JaJa, J. 2007. ACE: a Novel Software Platform to Ensure the Long Term Integrity of Digital Archives, Proceedings of the Archiving 2007 Conference, May 2007, Washington, DC. http://adaptwiki.umiacs.umd.edu/twiki/pub/Lab/Papers/rad71E67.pdf

[10] Caplan, P. 2006. DCC Digital Curation Manual: Installment on Preservation Metadata, http://www.dcc.ac.uk/resource/curation-manual/chapters/preservation-metadata/preservation-metadata.pdf

[11] PREMIS Working Group. 2008. PREMIS Data Dictionary for Preservation Metadata version 2.0. http://www.loc.gov/premis/v2/premis-2-0.pdf

[12] Lavoie, B. 2008. PREMIS with a Fresh Coat of Paint: Highlights from the Revision of the PREMIS Data Dictionary for Preservation Metadata, D-Lib Magazine, Vol.14 No.5/6. http://www.dlib.org/dlib/may08/05contents.html