Requirements and Characteristics of a Preservation Quality Information Management System

Gabrielle V. Michalek
Carnegie Mellon University
October, 2010

Introduction

The proliferation of digital materials has changed not only how information is presented but also how people expect information to be available. People want access to all forms of information, from simple text to complex multimedia. Whether or not the items in question were created digitally, they can be made to behave digitally through scanning and conversion. This improves access, but makes preservation more difficult because of the rapid rate of obsolescence of formats, hardware and software systems.

In the early days of the digital age, the gap between librarians and the people working in information technology was vast. The past twenty years, however, have seen thousands of digitization projects that include scanning paper analog resources and making them available in digital format. Such work has coupled the worlds of tech and libraries. This cross pollination has resulted in rich, robust online resources worthy of preservation. Libraries have a role in preservation of such resources, much as they have had a role in preservation of the hard-copy word.

The Carnegie Mellon University Libraries has been at the forefront where digitization of rare and primary sources is concerned. Our pioneering efforts in digitization have also involved a creation of preservation strategies for digital content.

Strategies used by Carnegie Mellon University Libraries in System Development

This paper focuses on the requirements of an information management system that is designed to ensure the long term preservation of the information managed and delivered by that system. It is not an all inclusive list or description, but rather, an overview of the approaches and methodologies that have been used by the Carnegie Mellon University Libraries as we have built, developed, and now, acquired, information management systems that have the onerous task of supporting collections we have promised to keep in perpetuity.

This paper will focus on those approaches, activities and philosophies maintained at Carnegie Mellon University Libraries to ensure long term preservation is part of the design and management of our systems. These activities break down into four categories: adherence to standards for image and metadata creation, management of data, system architecture and security and auditing.

Use of Standards for Image and Metadata Creation

Use of standards is basic library practice and one that has been accepted by those tech specialists who partner with libraries. By now, most viable digital library programs have gotten onboard with the standards that pertain to image creation, derivative format creation, and some metadata creation. Most programs have embraced the standards that support image creation simply because they lead to the production of better image quality, better reproduction quality and longevity and interoperability of the data. Carnegie Mellon University Libraries participated in the work by the Digital Library Federation for the creation of their 2002 publication, "Benchmark for Faithful Digital Reproduction of Monographs and Serials".

<http://www.diglib.org/standards/bmarkfin.htm> The Carnegie Mellon University Libraries have been following these recommendations along with the standards set forth by The Library of Congress, <http://memory.loc.gov/ammem/about/techStandards.pdf>.

In addition, standards are beginning to evolve around derivative formats. JPEG2000 and PDF-A are two of these formats. The PDF-A format is good for archiving since it is 100% self-contained and even allows for the embedding of metadata into the file. Most files that began life as textual documents can be converted to PDF-A format making it an attractive choice for professionals concerned about long term preservation. Carnegie Mellon University Libraries maintains PDF-A files for most of our digital collections where the PDF format makes sense. Textual documents, and JPEG2000 format for image collections, are examples.

One of the most important methods of supporting long term preservation for digital collections is through the creation and/or acquisition of reliable metadata. There are several different types of metadata, each representing a different purpose such as descriptive, and technical. If each of the metadata types is used correctly and if they are used in concert with one another, chances of long term preservation increase significantly.

The descriptive metadata describes the collection. The standards for descriptive metadata that we use most frequently at the Carnegie Mellon University Libraries include, MARC for library materials, Dublin Core for library materials and historic image collections, Encoded Archival Description (EAD) for archival collections, and Visual Resources Association Data Core (VRA) for art image collections. These metadata schemas are in XML format and work as part of the XML platform that makes up the system architecture.

Technical metadata for each digital collection is also created. This metadata describes the attributes of the file, such as which software was used to create the file, the date the file was created, the file size, which software was used to create the image, etc... Much of this metadata is created by the software used to create the digital image. This information is captured and stored as an XML file and is used to manage the collection over time. It is especially useful when the collection is being migrated to a new platform.

Management of Data

Libraries and Archives have long been involved in management of data sources. The digital world spend its first few decades not taking advantage of the data management expertise so easily found in the local library. Much as digitization projects have brought IT methods into the library world, so library sciences such as data management have gradually become incorporated into the IT mindset. To help ensure both long term preservation of data and persistent integrity of that data, Carnegie Mellon University Libraries has put into place a program, mostly administered by System Managers, to develop tools and techniques to manage data effectively over time. One technique is the daily back up of scanned images to a remote server. This automated process includes email alerts that notify the System Mangers of the progress and of transfer problems. The data are then periodically backed up to tape. According to industry standards and best practices, these tapes have a life expectancy of up to 25 years. However, because tape capacities continue to increase, each time we change to a new type of tape with a larger capacity format, all of the tapes are read back and re-written to the new backup tapes. At a minimum, tapes are spot checked annually. In essence, what we have here is a disk to disk tape backup strategy which parallels not only the method of data preservation of Carnegie Mellon's Computing Services Department, but also mirrors best practices of the IT industry at large.

The utility used for disk to disk copying is an open-source program called *rsync*. Rsync first compares the files (via an rsync algorithm) on the source and the destination, and only copies files which are newer or updated. Rsync effectively transfers only the changes across the network and merges them with the last full backup on the remote machine thus providing full backups of the local server with the added efficiency of incremental backups. There are other ways to reliably copy and check large volumes of data, however, this is the method selected by the Carnegie Mellon University Libraries. When the data is moved to backup tape, the Libraries uses a reliable utility and format called tar, (tape archival).

Post processed "fixed content" is stored on networked volumes that are part of our 10 Terabyte Storage Area Network (SAN) which is located in a climate controlled server room in Hunt Library. This SAN resides on a private "storage network" and it is only accessible by hosts that are also on this private storage network. This is a security precaution. The data are further protected by RAID 5 (distributed parity) technology. If there is any hardware failure or pending hardware failure, the system managers are notified via email alerts from the RAID array. In such a case, the spare disk is written to and the bad disk is immediately replaced thus restoring active RAID protection. All data have now been copied to our new SAN which has RAID protection.

To help ensure long term preservation, Managers must verify the integrity of the data within the repository. The method we use at Carnegie Mellon University libraries is by adding checksum capability at those locations where data is transferred from one place to another. The Managers compute checksums as early as possible and store the checksum alongside the data. Later on the manager runs the checksum again and compares the checksum with the results. If the sums do not match the file has been corrupted in transmission or storage. This capability helps verify the integrity of the data and that the

data have not been corrupted over time. This not only has the advantage of catching network problems, but also tape problems when data is brought back from tape.

For over a decade Carnegie Mellon University Libraries have been using persistent URLs as a way of ensuring that links to our digital objects do not break over time. This is especially important when scholars cite digital resources in their research. Broken links make it impossible for future scholars to retrace or continue the work. The persistent URL is actually a tool in the shape of a database that directs a request to the current URL for a resource. That link will always be associated with the resource even when data gets moved to different servers over time. This functionality is critical to preservation activities and it is essential that it is included as part of a digital library program.

System Architecture

Extensible Mark-up Language, XML, is an open standard, and a precise way of storing and communicating data. XML is easy to work with and aids preservation because it is human readable. It is especially good at representing complex data structures and hierarchies such as Encoded Archival Description (EAD), MARC and Dublin Core as described in the section on metadata creation. Several application programmable interface's (API's) have already been developed to explore and extract data stored in XML format making it robust and giving the specification long term viability. XML is platform independent so it allows a project to create data that can be read by applications on other platforms. This improves the odds of preservation by eliminating the concern that applications and platforms will become obsolete.

The Open Archives Initiative, OAI-PMH, is a method used to expose structured metadata such as EAD and Dublin Core to web crawlers. The metadata can then be discovered by internet search engines used by harvesters and aggregators such as OCLC's OAIster, or the National Science Digital Library (NSDL) in the United States, making it accessible to the world. In addition, the metadata from one set of collections can then be added to a database with metadata from other collections making the collections interoperable. Interoperability allows content from a diverse group of collections from multiple organizations to be organized, managed, and disseminated in ways that promote new learning to a wide range of communities. Since the EAD metadata schema is XML encoded, there are many ways in which to use the OAI standard to communicate the data between repositories and harvesters. Making data interoperable and sharing it with other repositories helps preserve information by distributing it across locations and platforms.

Security and Auditing

It is important that systems include software tools that allow for content versioning and the ability to conduct a digital audit. These two pieces of functionality will indicate if data has been modified, when the changes occurred, and who made the changes. This helps ensure accuracy of the information making researchers more confident in the integrity of the information. It also helps System Managers track down any problems

with the data and locate where in the data stream the problem may have occurred. While, DIVA, the current home grown information system used by the Carnegie Mellon University Libraries has rudimentary content versioning to produce an audit trail, a more robust version of this functionality is required as the Libraries moves to a new commercial grade information management system.

To aid in the creation of an audit trail, managers must include computer security that provides administrative control. This will also protect the data from theft, corruption and disaster. Users may have read permission and can access documents, but they may not alter the document or metadata in any way. Personnel working with data may have limited permissions to work within the system. They must be required to log in to add, delete or modify the data. System Managers and Administrators may have access to all of the data within the system to be able to change the data, and also to modify the system itself. However, they must be forced to log in as well. This, too, is functionality required of the new system Carnegie Mellon University Libraries is purchasing.

Good security controls authorization right down to the object level. For the end user, this security determines whether a user can access a folder, or a file, and if the end user can move from file to file, or folder to folder. Security control also determines if the end user can see all of the metadata associated with a file or just a limited set of metadata. For personnel, this determines whether the user can access, create, modify or delete a folder or file. For a System manager or Administrator, it determines if the user can do all of the above as well as change permissions down to the file level. This functionality is especially important as “born digital” materials are ingested to the system. As materials are digitized we have the luxury of reviewing each document to determine if it should be scanned and added to the system. This may not be the case with born digital materials such as email. It may be that Managers would like to make some of the emails available to end users because of reasons of confidentiality. This functionality would support these objectives.

Conclusion

For the past 20 years the Carnegie Mellon University Libraries has been a leader in the field of library digitization. Whenever the Libraries commit to creating and maintaining a digital collection it is making a promise to preserve that digital collection in perpetuity, or for as long as it is technically feasible. In the process of making these unique library resources available online we have gained experience and knowledge concerning preservation of these digital objects. This paper has described some of the tools and approaches our managers have used to aid in the long term preservation of the digital library. While this paper does not purport to be a comprehensive list of all of the technical approaches necessary to ensure long term preservation, our techniques are easily replicated and represent some of the best practices in the field.

References

Caplan, Priscilla. Understanding PREMIS. Washington, DC: Library of Congress, 2009.
<http://www.loc.gov/standards/premis/understanding-premis.pdf>.

Federal Agencies Digitization Guidelines Initiative. "Digital Conversion--Documents and Guidelines: A Bibliographic Reference." Updated August 28, 2009.
http://www.digitizationguidelines.gov/stillimages/documents/Guidelines_Bibliography-2009rev.pdf

Peterson, Zachary N.J., Randal Burns, Guiseppe Ateniese, and Stephen Bono. "Design and Implementation of Verifiable Audit Trails for a Versioning File System." *Proceedings of the Conference on File and Storage Technologies, USENIX*. (2007): 93-106.

Rose, Kenneth L. *Preserving Our Digital Collections*. Presentation for Cultural Memory Class, Carnegie Mellon University, Pittsburgh, PA, February 25, 2009.

Rumsey, Abby S., ed. Sustainable Economics for a Digital Planet: Ensuring Long-Term Access to Digital Information. Technical Report, Blue Ribbon Task Force on Sustainable Digital Preservation and Access, 2010.
http://brtf.sdsc.edu/biblio/BRTF_Final_Report.pdf