Preservation Environments

Reagan W. Moore
San Diego Supercomputer Center
University of California, San Diego
9500 Gilman Drive, MC-0505
La Jolla, CA 92093-0505
moore@sdsc.edu
tel: +1-858-534-5073
fax: +1-858-534 5152

Abstract:
The long-term preservation of digital entities requires mechanisms to manage the authenticity of massive data collections that are written to archival storage systems. Preservation environments impose authenticity constraints and manage the evolution of the storage system technology by building infrastructure independent solutions. This seeming paradox, the need for large archives, while avoiding dependence upon vendor specific solutions, is resolved through use of data grid technology. Data grids provide the storage repository abstractions that make it possible to migrate collections between vendor specific products, while ensuring the authenticity of the archived data. Data grids provide the software infrastructure that interfaces vendor-specific storage archives to preservation environments.

1. Introduction

A preservation environment manages both archival content (the digital entities that are being archived), and archival context (the metadata that are used to assert authenticity) [8]. Preservation environments integrate data storage repositories with information repositories, and provide mechanisms to maintain consistency between the context and content. Preservation systems rely upon software systems to manage and interpret the data bits. Traditionally, a digital entity is retrieved from an archival storage system, structures within the digital entity are interpreted by an application that issues operating system I/O calls to read the bits, and semantic labels that assign meaning to the structures are organized in a database. This process requires multiple levels of software, from the archival storage system software, to the operating system on which the archive software is executed, to the application that interprets and displays the digital entity, to the database that manages the descriptive context. A preservation environment assumes that each level of the software hierarchy used to manage data and metadata will change over time, and provides mechanisms to manage the technology evolution.

A digital entity by itself requires interpretation. An archival context is needed to describe the provenance (origin), format, data model, and authenticity [9]. The context is created by archival processes, and managed through the creation of attributes that describe the knowledge needed to understand and display the digital entities. The archival context is organized as a collection that must also be preserved. Since archival storage systems manage files, software infrastructure is needed to map from the archival repository to the preservation collection. Data Grids provide the mechanisms to manage collections that are preserved on vendor-supplied storage repositories [7].

Preservation environments manage collections for time periods that are much longer than the lifetime of any storage repository technology. In effect, the collection is held invariant while the underlying technology evolves. When dealing with Petabyte-sized collections, this is a non-trivial problem. The preservation environment must provide mechanisms to migrate collections onto new technology as it becomes available. The driving need behind the migrations is to take advantage of lower-cost storage repositories that provide higher capacity media, faster data transfer rates, smaller foot-print, and reduced operational maintenance. New technology can be more cost effective.

2. Persistent Archives and Data Grids

A persistent archive is an instance of a preservation environment [9]. Persistent archives provide the mechanisms to ensure that the hardware and software components can be upgraded over time, while maintaining the authenticity of the collection. When a digital entity in migrated to a new storage repository, the persistent archive guarantees the referential integrity between the archival context, and the new location of the digital entity. Authenticity also implies the ability to manage audit trails that record all operations performed upon the digital entity, access controls for asserting that only archivists performed the operations, and checksums to assert the digital entity has not been modified between applications of archival processes.

Data grids provide these data management functions in addition to abstraction mechanisms for providing infrastructure independence [7]. The abstractions are used to define the fundamental operations that are needed on storage repositories to support access and manipulation of data files. The data grid maps from the storage repository abstraction to the protocols required by a particular vendor product. By adding drivers for each new storage protocol as they are created, it is possible for a data grid to manage digital entities indefinitely into the future. Each time a storage repository becomes obsolete, the digital entities can be migrated onto a new storage repository. The migration is feasible as long as the data grid uses a logical name space to create global, persistent identifiers for the digital entities. The logical name space is managed as a collection, independently of the storage repositories. The data grid maps from the logical name space identifier to the file name within the vendor storage system.

Data grids support preservation by applying mappings to the logical name space to define the preservation context. The preservation context includes administrative attributes (location, ownership, size), descriptive attributes (provenance, discovery attributes), structural attributes (components within a compound record), and behavioral attributes (operations that can be performed on the digital entity). The context is managed as metadata in a database. An information repository abstraction is used to define the operations required to manipulate a collection within a database, providing the equivalent infrastructure independence mechanisms for the collection.

Archivists apply archival processes to convert digital entities into archival forms. Similar ideas of infrastructure independence can be used to characterize and manage archival processes. The application of each archival process generates part of the archival context. By creating an infrastructure independent characterization of the archival processes, it becomes possible to apply the archival processes in the future. An archival form can then consist of the original digital entity and the characterization of the archival process. Virtual data grids support the characterization of processes and on demand application of the process characterizations. A reference to the product generated by a process can result in direct access to the derived data product, or can result in the application of the process to create the derived data product. Virtual data grids can characterize and apply archival processes.

Data grids provide the software mechanisms needed to manage the evolution of software infrastructure [7] and automate the application of archival processes. The standard capabilities provided by data grids were assessed by the Persistent Archive Research Group of the Global Grid Forum [8]. Five major categories were identified that are provided by current data grids:

1. Logical name space; a persistent and infrastructure independent naming convention
2. Storage repository abstraction; the operations that are used to access and manage data
3. Information repository abstraction; the operations that are used to organize and manage a collection within a database
4. Distributed resilient architecture; the federated client-server architecture and latency management functions needed for bulk operations on distributed data
5. Virtual data grid; the ability to characterize the processing of digital entities, and apply the processing on demand.

The assessment compared the Storage Resource Broker (SRB) data grid from the San Diego Supercomputer Center [18], the European DataGrid replication environment (based upon GDMP, a project in common between the European DataGrid [2] and the Particle Physics Data Grid [15], and augmented with an additional product of the European DataGrid for storing and retrieving meta-data in relational databases called Spitfire and other components), the Scientific Data Management (SDM) data grid from Pacific Northwest Laboratory [20], the Globus toolkit [3], the Sequential Access using Metadata (SAM) data grid from Fermi National Accelerator Laboratory [19], the Magda data management system from Brookhaven National Laboratory [6], and the JASMine data grid from Jefferson National Laboratory [4]. These systems have evolved as the result of input by user communities for the management of data across heterogeneous, distributed storage resources.

EGP, SAM, Magda, and JASMine data grids support high energy physics data. The SDM system provides a digital library interface to archived data for PNL and manages data from multiple scientific disciplines. The Globus toolkit provides services that can be composed to create a data grid. The SRB data handling system is used in projects for multiple US federal agencies, including the NASA Information Power Grid (digital library front end to archival storage) [11], the DOE Particle Physics Data Grid (collection-based data management) [15], the National Library of Medicine Visible Embryo project (distributed data collection) [21], the National Archives Records Administration (persistent archive research prototype) [10], the NSF National Partnership for Advanced Computational Infrastructure (distributed data collections for astronomy, earth systems science, and neuroscience) [13], the Joint Center for Structural Genomics (data grid) [5], and the National Institute of Health Biomedical Informatics Research Network (data grid) [1].

The systems therefore include not only data grids, but also distributed data collections, digital libraries and persistent archives. Since the core component of each system is a data grid, common capabilities do exist across the multiple implementations. The resulting core capabilities and functionality are listed in Table 1.

These capabilities should encompass the mechanisms needed to implement a persistent archive. This can be demonstrated by mapping the functionality required by archival processes onto the functionality provided by data grids.

Core Capabilities and Functionality
Storage repository abstraction
Storage interface to at least one repository
Standard data access mechanism
Standard data movement protocol support
Containers for data
Logical name space
Registration of files in logical name space
Retrieval by logical name
Logical name space structural independence from physical file
Persistent handle
Information repository abstraction
Collection owned data
Collection hierarchy for organizing logical name space
Standard metadata attributes (controlled vocabulary)
Attribute creation and deletion
Scalable metadata insertion
Access control lists for logical name space
Attributes for mapping from logical file name to physical file
Encoding format specification attributes
Data referenced by catalog query
Containers for metadata
Distributed resilient scalable architecture
Specification of system availability
Standard error messages
Status checking
Authentication mechanism
Specification of reliability against permanent data loss
Specification of mechanism to validate integrity of data
Specification of mechanism to assure integrity of data
Virtual Data Grid
Knowledge repositories for managing collection properties
Application of transformative migration for encoding format
Application of archival processes

Table 1. Core Capabilities of Data Grids

3. Persistent Archive Processes

The preservation community has identified standard processes that are applied in support of paper collections, listed in Table 2. These standard processes have a counterpart in the creation of archival forms for digital entities. The archival form consists of the original bits of the digital entity plus the archival context that describes the origin (provenance) of the data, the authenticity attributes, and the administrative attributes. A preservation environment applies the archival processes to each digital entity through use of a dataflow system, records the state information that results from each process, organizes the state information into a preservation collection, transforms the digital entity into a sustainable format, archives the original digital entity and its transforms, and provides the ability to discover and retrieve a specified digital entity.

Archival Process	Functionality
Appraisal	Assessment of digital entities
Accession	Import of digital entities
Description	Assignment of provenance metadata
Arrangement	Logical organization of digital entities
Preservation	Storage in an archive
Access	Discovery and retrieval

Table 2. Archival process functionality for paper records

To understand whether data grids can meet the archival processing requirements for digital entities, scenarios are given below for the equivalent operations on digital entities. The term record is used to denote a digital entity that is the result of a formal process, and thus a candidate for preservation. The term fonds is used to denote a record series.

Appraisal is the process of determining the disposition of records and in particular which records need long-term preservation. Appraisal evaluates the various terms and conditions applying to the preservation of records beyond the time of their active life in relation to the affairs that created them. An archivist bases an appraisal decision on the uniqueness of the record collection being evaluated, its relationship to other institutional records, and its relationship to the activities, organization, functions, policies, and procedures of the institution.

Data grids provide the ability to register digital entities into a logical name space organized as a collection hierarchy for comparison with other records of the institution that have already been accessioned into the archives. The logical name space is decoupled from the underlying storage systems, making it possible to reference digital entities without moving them. The metadata associated with those other collections assist the archivist in assessing the relationship of the records being appraised to the prior records. Queries are made on the descriptive and provenance metadata to identify relevant records. The data grid supports controlled vocabularies for describing provenance and formats. This metadata also provides information that helps the archivist understand the relevance/importance/value of the records being appraised for documenting the activities, functions, etc. of the institution that created them. The activities of the institution can be managed as relationships maintained in a concept space, or as process characterizations maintained in a procedural ontology. By authorizing archivist access to the collection, and providing mechanisms to ensure authenticity of the previously archived records, the preservation environment maintains an authentic environment.

Accessioning is the formal acceptance into custody and recording of an acquisition. Data Grids control import by registering the digital entities into a logical name space organized as a collection/sub-collection hierarchy. The records that are being accessioned can be managed as a collection independently of the final archival form. By having the data grid own the records (stored under a data grid Unix ID), all accesses to the records can be tracked through audit trails. By associating access controls with the logical name space, all references to the records can be authorized no matter where the records are finally stored.

Data grids put digital entities under management control, such that automated processing can be done across an entire collection. Bulk operations are used to move the digital entities using a standard protocol and to store the digital entities in a storage repository. Digital entities may be aggregated into containers (the equivalent of a cardboard box for paper) to control the data distribution within the storage repository. Containers are used to minimize the impact on the storage repository name space. The metadata catalog manages the mapping from the digital entities to the container in which they are written. The storage repository only sees the container names. Standard clients are used for controlling the bulk operations.

The information repository supports attribute creation and deletion to preserve record or fonds specific information. In particular, information on the properties of the records and fonds are needed for validation of the encoding formats and to check whether the entire record series has been received. The accession schedule may specify knowledge relationships that can be used to determine whether associated attribute values are consistent with implied knowledge about the collection, or represent anomalies and artifacts. An example of a knowledge relationship is the range of permissible values for a given attribute, or the expected number of records in a fonds. If the range of values do not match the assertions provided by the submitter, the archivist needs to note the discrepancy as a property of the collection.

Bulk operations are needed on metadata insertion when dealing with collections that contain millions of digital entities. A resilient architecture is needed to specify the storage system availability, check system status, authenticate access by the submitting institution, and specify reliability against data loss. At the time of accession, mechanisms such as checksums, need to be applied to be able to assert in the future that the data has not been changed.

The Open Archival Information System (OAIS) specifies submission information packages that associate provenance information with each digital entity [14]. While OAIS is presented in terms of packaging of information with each digital entity, the architecture allows bulk operations to be implemented. An example is bulk loading of multiple digital entities, in which the provenance information is aggregated into an XML file, while the digital entities are aggregated into a container. The XML file and container are moved over the network from the submitting site to the preservation environment, where they are unpacked into the storage and information repositories.

The integrity of the data (the consistency between the archival context and archival content) needs to be assured, typically by imposing constraints on metadata update.

When creating replicas and aggregating digital entities into containers, state information is required to describe the status of the changes. When digital entities are appended to a container, write locks are required to avoid over-writes. When a container is replicated, a synchronization flag is required to identify which container holds the new digital entities, and synchronization mechanisms are needed to update the replicas.

The accession process may also impose transformative migrations on encoding formats to assure the ability to read and display a digital entity in the future. The transformative migrations can be applied at the time of accession, or the transformation may be characterized such that it can be applied in the future when the digital entity is requested.

In order to verify properties of the entire collection, it may be necessary to read each digital entity, verify its content against an accession schedule, and summarize the properties of all of the digital entities within the record series. The summarization is equivalent to a bill of lading for moving the record series into the future. When the record series is examined at a future date, the archivist needs to be able to assert that the collection is complete as received, and that missing elements were never submitted to the archive. Summarization is an example of a collection property that is asserted about the entire record series. Other collection properties include completeness (references to records within the collection point to other records within the collection), and closure (operations on the records result in data products that can be displayed and manipulated with mechanisms provided by the archive). The closure property asserts that the archive can manipulate all encoding formats that are deposited into the archive.

Arrangement is the process and result of identification of documents for whether they belong to accumulations within a fonds or record series. Arrangement requires organization of both metadata (context) and digital entities (content). The logical name space is used as the coordination mechanism for associating the archival context with the submitted digital entities. All archival context is mapped as metadata attributes onto the logical name for each digital entity. The logical name space is also used as the underlying naming convention on which a collection hierarchy is imposed. Each level of the collection hierarchy may have a different archival context expressed as a different set of metadata. The metadata specifies relationships of the submitted records to other components of the record series. For a record series that has yearly extensions, a suitable collection hierarchy might be to organize each year's submission as a separate sub-collection, annotated with the accession policy for that year. The digital entities are sorted into containers for physical aggregation of similar entities. The expectation is that access to one digital entity will likely require access to a related digital entity. The sorting requires a specification of the properties of the record series that can be used for a similarity analysis. The container name in which a digital entity is placed is mapped as an administrative attribute onto the logical name. Thus by knowing the logical name of a digital entity within the preservation environment, all pertinent information can be retrieved or queried.

The process of arrangement points to the need for a digital archivist workbench. The storage area that is used for applying archival processes does not have to be the final storage location. Data grids provide multiple mechanisms for arranging data, including soft-links between collections to associate a single physical copy with multiple sub-collections, copies that are separately listed in different sub-collections, and versions within a single sub-collection. Data grids provide multiple mechanisms for managing data movement, including copying data between storage repositories, moving data between storage repositories, and replicating data between storage repositories.

Description is the recording in a standardized form of information about the structure, function and content of records: Description requires a persistent naming convention and a characterization of the encoding format, as well as information used to assert authenticity. The description process generates the archival context that is associated with each digital entity. The archival context is includes not only the administrative metadata generated by the accession and arrangement processes, but also descriptive metadata that are used for subsequent discovery and access.

Preservation Function	Type of information
Administrative	Location, physical file name, size, creation time, update time, owner, location in a container, container name, container size, replication locations, replication times
Descriptive	Provenance, submitting institution, record series attributes, discovery attributes
Authenticity	Global Unique Identifier, checksum, access controls, audit trail, list of transformative migrations applied
Structural	Encoding format, components within digital entity
Behavioral	Viewing mechanisms, manipulation mechanisms

Table 3. Archival context managed for each digital entity

The description process can require access to the storage repository to apply templates for the extraction of descriptive metadata, as well as access to the information catalog to manage the preservation of the metadata. The description process should generate a persistent handle for the digital entity in addition to the logical name. The persistent handle is used to assert equivalence across preservation environments. An example of a persistent handle is the concatenation of the name of the preservation environment and the logical name of the entity, and is guaranteed unique as long as the preservation environments are uniquely named. The ability to associate a unique handle with a digital entity that is already stored requires the ability to apply a validation mechanism such as a digital signature or checksum to assert equivalence. If a transformative migration has occurred, the validation mechanism may require access to the original form of the digital entity.

Preservation is the process of protecting records of continuing usefulness: Preservation requires a mechanism to interact with multiple types of storage repositories, mechanisms for disaster recovery, and mechanisms for asserting authenticity.

The only assured mechanism for guaranteeing against content or context loss is the replication of both the digital entities and the archival metadata. The replication can implement bit-level equivalence for asserting that the copy is authentic. The replication must be done onto geographically remote storage and information repositories to protect against local disasters (fire, earthquake, flood). While data grids provide tools to replicate digital entities between sites, some form of federation mechanism is needed to replicate the archival context and logical name space. One would like to assert that a completely independent preservation environment can be accessed that replicates even the logical names of the digital entities. The independent systems are required to support recovery from operation errors, in which recovery is sought from the mis-application of the archival procedures themselves.

The coordination of logical name spaces between data grids is accomplished through peer-to-peer federation. Consistency controls on the synchronization of digital entities and metadata between the data grids are required for the user name space (who can access digital entities), the resources (whether the same repository stores data from multiple grids), the logical file names (whether replication is managed by the systems or archival processes), and the archival context (whether insertion of new entities is managed by the system or archival processes). Multiple versions of control policies can be implemented, ranging from automated replication into a union archive from multiple data grids, to simple cross-registration of selected sub-collections.

Data grids use a storage repository abstraction to manage interactions with heterogeneous storage systems. To avoid problems specific to vendor products, the archival replica should be made onto a different vendor's product from the primary storage system. The heterogeneous storage repositories can also represent different versions of storage systems and databases as they evolve over time. When a new infrastructure component is added to a persistent archive, both the old version and new version will be accessed simultaneously while the data and information content are migrated onto the new technology. Through use of replication, the migration can be done transparently to the users. For persistent archives, this includes the ability to migrate a collection from old database technology onto new database technology.

Persistence is provided by data grids through support for a consistent environment, which guarantees that the administrative attributes used to identify derived data products always remain consistent with migrations performed on the data entities. The consistent state is extended into a persistent state through management of the information encoding standards used to create platform independent representations of the context. The ability to migrate from an old representation of an information encoding standard to a new representation leads to persistent management of derived data products. It is worth noting that a transformative migration can be characterized as the set of operations performed on the encoding syntax. The operations can be applied on the original digital entity at the time of accession or at any point in the future. If a new encoding syntax standard emerges, the set of operations needed to map from the original encoding syntax to the new encoding syntax can be defined, without requiring any of the intermediate encoding representations. The operations needed to perform a transformative migration are characterized as a digital ontology [8].

Authenticity is supported by data grids through the ability to track operations done on each digital entity. This capability can be used to track the provenance of digital entities, including the operations performed by archivists. Audit trails record the dates of all transactions and the names of the persons who performed the operations. Digital signatures and checksums are used to verify that between transformation events the digital entity has remained unchanged. The mechanisms used to accession records can be re-applied to validate the integrity of the digital entities between transformative migrations. Data grids also support versioning of digital entities, making it possible to store explicitly the multiple versions of a record that may be received. The version attribute can be mapped onto the logical name space as both a time-based snapshot of a changing record, and as an explicitly named version.

Access is the process of using descriptive metadata to search for archival objects of interest and retrieve them from their storage location. Access requires the ability to discover relevant documents, transport them from storage to the user, and interact with storage systems for document retrieval. The essential component of access is the ability to discover relevant files. In practice, data grids use four naming conventions to identify preserved content. A global unique identifier (GUID) identifies digital entities across preservation environments, the logical name space provides a persistent naming convention within the preservation environment, descriptive attributes support discovery based on attribute values, and the physical file name identifies the digital entity within a storage repository. In most cases, the user of the system will not know either the GUID, logical name or physical file name, and discovery is done on the descriptive attributes.

Access then depends upon the ability to instantiate a collection that can be queried to discover a relevant digital entity. A knowledge space is needed to define the semantic meaning of the descriptive attributes, and a mechanism is needed to create the database instance that holds the descriptive metadata. For a persistent archive, this is the ability to instantiate an archival collection from its infrastructure independent representation onto a current information repository. The information repository abstraction supports the operations needed to instantiate a metadata catalog.

The other half of access is transport of the discovered records. This includes support for moving data and metadata in bulk, while authenticating the user across administration domains. Since access mechanisms also evolve in time, mechanisms are needed to map from the storage and information repository abstractions to the access mechanism preferred by the user.

4. Preservation Infrastructure

The operations required to support archival processes can be organized by identifying which capability is used by each process. The resulting preservation infrastructure is shown in Table 4. The list includes the essential capabilities that simplify the management of collections of digital entities while the underlying technology evolves.

The use of each capability by one of the six archival processes is indicated by an x in the appropriate row. The columns are labeled by App (Appraisal), Acc (Accessioning), Arr (Arrangement), Des (Description), Pres (Preservation), and Ac (Access). Many of the data grid capabilities are required by all of the archival processes. This points out the difficulty in choosing an appropriate characterization for applying archival processes to digital entities. Even though we have shown that the original paper-oriented archival processes have a counterpart in preservation of digital entities, there may be a better choice for characterizing electronic archival processes.

Core Capabilities and Functionality	App	Acc	Arr	Des	Pres	Ac
Storage repository abstraction		x	x		x	x
Storage interface to at least one repository		x	x	x	x	x
Standard data access mechanism		x	x	x	x	x
Standard data movement protocol support		x	x	x	x	x
Containers for data		x	x		x	x
Logical name space	x	x	x	x	x	x
Registration of files in logical name space	x	x	x	x	x	
Retrieval by logical name		x	x		x	x
Logical name space structural independence from physical file	x	x	x	x	x	x
Persistent handle		x	x	x	x	x
Information repository abstraction	x	x	x	x	x	x
Collection owned data	x	x	x	x	x	x
Collection hierarchy for organizing logical name space	x	x	x	x		
Standard metadata attributes (controlled vocabulary)	x	x	x	x	x	x
Attribute creation and deletion	x	x	x	x	x	
Scalable metadata insertion		x	x	x	x	
Access control lists for logical name space	x	x	x	x	x	x
Attributes for mapping from logical file name to physical file		x	x		x	x
Encoding format specification attributes	x	x		x	x	x
Data referenced by catalog query						x
Containers for metadata		x	x	x	x	x
Distributed resilient scalable architecture	x	x	x	x	x	x
Specification of system availability		x			x	x
Standard error messages		x	x	x	x	x
Status checking		x	x	x	x	x
Authentication mechanism	x	x	x	x	x	x
Specification of reliability against permanent data loss	x	x	x	x	x	
Specification of mechanism to validate integrity of data		x	x	x	x	x
Specification of mechanism to assure integrity of data	x	x	x	x	x	x
Virtual Data Grid		x	x	x	x	x
Knowledge repositories for managing collection properties	x	x	x	x	x	x
Application of transformative migration for encoding format		x	x	x	x	x
Application of archival processes		x	x	x	x	x

Table 4. Data Grid capabilities used in preservation environments

5. Persistent Archive Prototype

The preservation of digital entities is being implemented at the San Diego Supercomputer Center (SDSC) through multiple projects that apply data grid technology. In collaboration with the United States National Archives and Records Administration (NARA), SDSC is developing a research prototype persistent archive. The preservation environment is based on the Storage Resource Broker (SRB) data grid [17], and links three archives at NARA, the University of Maryland, and SDSC. For the National Science Foundation, SDSC has implemented a persistent archive for the National Science Digital Library [12]. Snapshots of digital entities that are registered into the NSDL repository as URLs are harvested from the web and stored into an archive using the SRB data grid. As the digital entities change over time, versions are tracked to ensure that an educator can find the desired version of a curricula module.

Both of these projects rely upon the ability to create archival objects from digital entities through the application of archival processes. We differentiate between the generation of archival objects through the application of archival processes, the management of archival objects using data grid technology, and the characterization of the archival processes themselves, so that archived material can be re-processed (or re-purposed) in the future using virtual data grids.

The San Diego Supercomputer Center Storage Resource Broker (SRB) is used to implement the persistent archives. The SRB provides mechanisms for all of the capabilities and functions listed in Table 2 except for knowledge repositories. The SRB also provides mechanisms for the extended features listed in section 3, such as soft-links, peer-to-peer federation of data grids, and mapping to user-preferred APIs. The SRB storage repository abstraction is based upon standard Unix file system operations, and supports drivers for accessing digital entities stored in Unix file systems (Solaris, SunOS, AIX, Irix, Unicos, Mac OS X, Linux), in Windows file systems (98, 2000, NT, XP, ME), in archival storage systems (HPSS, UniTree, DMF, ADSM, Castor, Dcache, Atlas Data Store), as binary large objects in databases (Oracle, DB2, Sybase, SQLServer, PostgresSGL), in object ring buffers, in storage resource managers, in FTP sites, in GridFTP sites, on tape drives managed by tape robots, etc. The SRB has been designed to facilitate the addition of new drivers for new types of storage systems. Traditional tape-based archives still remain the most cost-effective mechanism for storing massive amounts of data, although the cost of commodity-based disk is approaching that of tape [17]. The SRB supports direct access to tapes in tape robots.

The SRB information repository abstraction supports the manipulation of collections stored in databases. The manipulations include the ability to add user-defined metadata, import and export metadata as XML files, support bulk registration of digital entities, apply template-based parsing to extract metadata attribute values, and support queries across arbitrary metadata attributes. The SRB automatically generates the SQL that is required to respond to a query, allowing the user to specify queries by operations on attribute values.

Version 3.0.1 of the Storage Resource Broker data grid provides the basic mechanisms for federation of data grids [16]. The underlying data grid technology is in production use at SDSC and manages over 90 Terabytes of data comprising over 16 million files. The ultimate goal of the NARA research prototype persistent archive is to identify the key technologies that facilitate the creation of a preservation environment.

5. Summary

Persistent archives manage archival objects by providing infrastructure independent abstractions for interacting with both archival objects and software infrastructure. Data grids provide the abstraction mechanisms for managing evolution of storage and information repositories. Persistent archives use the abstractions to preserve the ability to manage, access and display archival objects while the underlying technologies evolve.

The challenge for the persistent archive community is the demonstration that data grid technology provides the correct set of abstractions for the management of software infrastructure. The Persistent Archive Research Group of the Global Grid Forum is exploring this issue, and is attempting to define the minimal set of capabilities that need to be provided by data grids to implement persistent archives [8]. A second challenge is the development of digital ontologies that characterize the structures present within digital entities. The Data Format Description Language research group of the Global Grid Forum is developing an XML-based description of the structures present within digital entities, as well as a description of the semantic labels that are applied to the structures. A third challenge is the specification of a standard set of operations that can be applied to the relationships within an archival object. A preservation environment will need to support operations at the remote storage repository, through the application of a digital ontology.

6. Acknowledgements

The concepts presented here were developed by members of the Data and Knowledge Systems group at the San Diego Supercomputer Center. The Storage Resource Broker was developed principally by Michael Wan and Arcot Rajasekar. This research was supported by the NSF NPACI ACI-9619020 (NARA supplement), the NSF NSDL/UCAR Subaward S02-36645, the DOE SciDAC/SDM DE-FC02-01ER25486 and DOE Particle Physics Data Grid, the NSF National Virtual Observatory, the NSF Grid Physics Network, and the NASA Information Power Grid. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the National Science Foundation, the National Archives and Records Administration, or the U.S. government. This document is based upon an informational document submitted to the Global Grid Forum. The data grid and Globus toolkit characterizations were only possible through the support of the following persons: Igor Terekhov (Fermi National Accelerator Laboratory), Torre Wenaus (Brookhaven National Laboratory), Scott Studham (Pacific Northwest Laboratory), Chip Watson (Jefferson Laboratory), Heinz Stockinger and Peter Kunszt (CERN), Ann Chervenak (Information Sciences Institute, University of Southern California), Arcot Rajasekar (San Diego Supercomputer Center). Mark Conrad (NARA) provided the archival process characterization.

7. References

1. Biomedical Informatics Research Network, http://nbirn.net/
2. EDG – European Data Grid, http://eu-datagrid.web.cern.ch/eu-datagrid/
3. Globus – The Globus Toolkit, http://www.globus.org/toolkit/
4. Jasmine – Jefferson Laboratory Asynchronous Storage Manager, http://cc.jlab.org/scicomp/JASMine/
5. Joint Center for Structural Genomics, http://www.jcsg.org/
6. Magda – Manager for distributed Grid-based Data, http://atlassw1.phy.bnl.gov/magda/info
7. Moore, R., C. Baru, "Virtualization Services for Data Grids", Book chapter in "Grid Computing: Making the Global Infrastructure a Reality", John Wiley & Sons Ltd, 2003.
8. Moore, R., A. Merzky, "Persistent Archive Concepts", Global Grid Forum Persistent Archive Research Group, Global Grid Forum 8, June 26, 2003.
9. Moore, R., "The San Diego Project: Persistent Objects," Archivi & Computer, Automazione E Beni Culturali, l'Archivio Storico Comunale di San Miniato, Pisa, Italy, February, 2003.
10. NARA Persistent Archive Prototype, http://www.sdsc.edu/NARA/Publications.html
11. NASA Information Power Grid (IPG) is a high-performance computing and data grid, http://www.ipg.nasa.gov/
12. National Science Digital Library, http://nsdl.org
13. NPACI National Partnership for Advanced Computational Infrastructure, http://www.npaci.edu/
14. OAIS - Reference Model for an Open Archival Information System (OAIS). submitted as ISO draft, http://www.ccsds.org/documents/pdf/CCSDS-650.0-R-1.pdf, 1999.
15. Particle Physics Data Grid, http://www.ppdg.net/
16. Peer-to-peer federation of data grids, http://www.npaci.edu/dice/srb/FedMcat.html
17. Rajasekar, A., M. Wan, R. Moore, G. Kremenek, T. Guptil, "Data Grids, Collections, and Grid Bricks", Proceedings of the 20th IEEE Symposium on Mass Storage Systems and Eleventh Goddard Conference on Mass Storage Systems and Technologies, San Diego, April 2003.
18. Rajasekar, A., M. Wan, R. Moore, "mySRB and SRB, Components of a Data Grid", 11th High Performance Distributed Computing conference, Edinburgh, Scotland, July 2002.
19. SAM – Sequential data Access using Metadata, http://d0db.fnal.gov/sam/.
20. SDM – Scientific Data Management in the Environmental Molecular Sciences Laboratory, http://www.computer.org/conferences/mss95/berard/berard.htm.
21. Visible Embryo Project, http://netlab.gmu.edu/visembryo.htm