

Conceptualizing Policy-Driven Repository Interoperability (PoDRI) Using iRODS and Fedora

David Pcolar

Carolina Digital Repository (CDR)
UNC Chapel Hill
david_pcolar@unc.edu

Daniel W. Davis

Cornell Information Sciences (CIS)
DuraSpace Affiliate
dwdavis@cs.cornell.edu

Bing Zhu

Data Intensive Cyber Environments
(DICE)
University of California: San Diego
bizhu@ucsd.edu

Alexandra Chassanoff

School of Information & Library
Science (SILS)
UNC Chapel Hill
achass@email.unc.edu

Chien-Yi Hou

Sustainable Archives & Leveraging
Technologies (SALT)
UNC Chapel Hill
chienyi@unc.edu

Richard Marciano

Sustainable Archives & Leveraging
Technologies (SALT)
UNC Chapel Hill
richard_marciano@unc.edu

Abstract

Given the growing need for cross-repository integration to enable a trusted, scalable, open and distributed content infrastructure, this paper introduces the Policy-Driven Repository Interoperability (PoDRI) project investigating interoperability mechanisms between repositories at the policy level. Simply moving digital content from one repository to another may not capture the essential management policies needed to ensure its integrity and authenticity. This project is focused on integrating policy-aware object models, including policy expressions, and a distributed architecture for policy-driven management, demonstrated using iRODS and Fedora as representative open source software products. Using iRODS and its Rules engine, combined with Fedora's rich semantic object model for digital objects, enables use of the best features of both products.

Index Keyword Terms—iRODS, Fedora, Preservation, Policy Management

1. Introduction

This paper introduces the Policy-Driven Repository Interoperability (PoDRI) project, investigating interoperability between repositories at the policy level. PoDRI is led by the University of North Carolina at UNC, with units ranging from SALT (Sustainable Archives & Leveraging Technologies), RENCI (Renaissance Computing Institute), SILS (School of Information and Library Science), and the Libraries/CDR (Carolina Digital Repository). Key partners include Bing Zhu at UCSD (DICE, Data Intensive Cyber Environments) and Daniel Davis at DuraSpace (combining DSpace and Fedora Commons) and Cornell Information Sciences. The project is sponsored by an Institute of Museum and Library Services (IMLS) National Leadership grant and is

motivated by the growing need to create a scalable, open, and distributed infrastructure that provides durable, trusted access and management of our valuable digital content of all kinds (e.g. research data sets, documents, video, metadata).

Simply replicating digital content from one repository, with or without any associated metadata, may not capture the essential management policies that ensure integrity and authenticity, a critical requirement for establishing a trust model. "A policy is typically a rule describing the interactions of actions that take place within the archive, or a constraint determining when and by whom an action may be taken." [1]. A distributed policy management architecture is an essential component in realizing a trust mechanism for repository interoperability. The PoDRI project investigates the requirements for policy-aware interoperability and demonstrates key features needed for its implementation. The project is focused on integrating object models, including interoperable policy expressions, and a policy-aware distributed architecture that includes both repositories and middleware services.

The PoDRI project addresses the following research problem: **What is the feasibility of repository interoperability at the policy level?** Research questions to be addressed are:

- Can a preservation environment be assembled from two existing repositories?
- Can the policies of the federation be enforced across repositories?
- Can policies be migrated between repositories?
- What fundamental mechanisms are needed within a repository to implement new policies?

iRODS, the Integrated Rule-Oriented Data System [2, 3] and the Fedora Repository [4, 5] will be used as representative open source software to demonstrate the PoDRI architecture. Combining iRODS and Fedora

enables use of the best features of both products for building sustainable digital repositories. iRODS provides an integrated rule engine, distributed virtual storage, the iCAT (iRODS Metadata Catalog)¹, and Micro-services². Fedora offers a rich semantic object modeling for digital objects, extensible format-neutral metadata and a flexible service mediation mechanism.

2. Rationale for Integrating Fedora and iRODS

Early in 2006, the DART [6] project created an Storage Resource Broker (SRB) storage interface for Fedora that allows all Fedora digital content, including Fedora Digital Objects (FDO) and their Datastreams, to be stored in SRB distributed repositories. Similarly, a storage module was developed by Aschenbrenner and Zhu [7] for iRODS. Using the Fedora-iRODS storage module, iRODS can act as a back-end for Fedora, and thus provide opportunities for Fedora to use iRODS capabilities such as virtual federated storage, micro-services and the rules engine.

iRODS offers an appealing platform for implementing a distributed policy-driven management architecture. The integrated rules engine can be used to invoke a range of rules including policy expressions and, through the use of micro-services, can execute code for those policies in a distributed environment. Rules can act as simple workflows, performing a sequence of pre-defined actions. iRODS rules can be executed explicitly, triggered by external conditions or events, and executed at timed intervals. For example, iRODS can implement a replication policy, geographically disbursing file copies across the network. Micro-services can be written for feature extraction, format migration, integrity checks and other preservation services.

While used to efficiently hold and query structured data and metadata, the iCAT relational database is not optimal for handling the complex, variable metadata needed for preservation and curation. Indeed, any relational database will require considerable coding to support complex metadata schemas, making the use of unstructured data (files) possibly in combination with XML databases or semantic triplestores a more flexible alternative [8].

Fedora is file-centric; all Fedora data and metadata is stored in files [9]. The Fedora Digital Object (FDO), a kind of compound digital object, provides the organizing metadata used to “make sense” of itself and other resources. It uses the FOXML schema to encapsulate

metadata, and to reference other files or web resources. Since the FDO is a file, it can be stored in iRODS like any other file.

Digital content (or user-defined metadata) managed by the FDO is stored in one or more separate files – each registered in a FOXML element called a Datastream. Datastreams can also capture relationships to other objects and external resources. Users may add metadata to the FDO or add additional metadata Datastreams (to be stored like any other file.

This means, however, that metadata is stored in an unstructured, often XML or RDF way, and requires external indices to support queries such as search engines, semantic triplestores, XML databases, and now the iCAT. Fedora’s approach provides a format-neutral, extensible framework for representing data and metadata.

The rich metadata environment provided by the FDO can augment the structured metadata found in the iCAT. Metadata can be copied from the iCAT into a more easily preserved unstructured file format, as demonstrated by Bing Zhu and colleagues [10]. Critical data can be copied from the FDO, or as user metadata files (Datastreams), so they can be queried from the iCAT. With suitable metadata, both the iCAT and Fedora could be entirely rebuilt from files if the indices were lost or corrupted.

Fedora has a set of “front-end” APIs that provide the means to ingest and manipulate FDOs (CRUD). iRODS is capable of calling these APIs to perform operations from micro-services. Fedora also provides an extensible mechanism to add custom functionality called “services” that are executed within the context of the FDO. Services act as extensions to the “front-end” API of the object. Fedora mediates the service request calling the appropriate “back-end” functionality. The back-end functionality can be a Web service, in this case potentially provided by iRODS. Custom Fedora services provide another mechanism to interact with iRODS. Since iRODS can interact with Fedora’s “front-end” APIs, “back-end” services, and the Fedora-iRODS storage module one may picture iRODS wrapping around Fedora.

3. First Steps Toward a Policy-driven Management Architecture

To demonstrate distributed policy-driven management architecture, we plan to implement the following operational scenarios:

- Integrate views of content, original arrangement (hierarchy) and metadata
- Create an audit trail of policy execution events and related provenance information
- Manage policies through Fedora
- Show iRODS invoking policies from Fedora

¹ iCAT is the iRODS Metadata Catalog that stores metadata about all objects in iRODS in a relational database.

² Micro-services are function snippets or executables that can be used to perform a distinct task using well-defined input information structures.

Both iRODS and Fedora fully support distributed computing installations. In effect, both products can be characterized as virtualization middleware for storage, access, and service execution. The products, however, have very different operational paradigms which must be accommodated, but provide complementary strengths that can be exploited when used together.

The virtual file system in iRODS makes it the logical choice for all storage (including FDOs). In addition, the iRODS rules engine and micro-services provide an effective means for policy invocation. Fedora's capabilities are especially powerful for handling variable content and metadata formats, to flexibly relate resources, to facilitate presentation, and its mediation capabilities make it appealing for supporting systems that are "designed for change."

A policy-driven management architecture requires that policy expressions be persistent somewhere. Fedora could be used to create FDOs containing policy expressions, which are subsequently loaded into actionable form and invoked in iRODS. As policies are part of the provenance, Fedora can relate the policy FDOs to the content items to which they apply. Since policy invocation is performed by iRODS, audit records of the execution must be created by iRODS; this will likely be done by creating FDOs (and relating them to the FDOs containing the content and policy expressions).

iRODS does not currently generate audit data in a format compliant with the PREMIS schema. The CDR implements auditing of objects via a PREMIS.XML file for each iRODS data object. This method may not be sustainable for repositories containing millions of objects. Preservation activities, such as replication or fixity checks, generate large amounts of log entries over time and potentially exceed the byte size of the original object. Discussions between CDR and iRODS developers suggest multiple methods for retaining and aggregating various component logs for translation into PREMIS-compliant events. Do we continue to store these events with the individual objects or as an aggregate? Do we generate specific PREMIS information upon request? In the case of replicas residing on disparate nodes in a data grid, auditable events will occur that differ from those affecting the original object. How do we reconcile these events in a singular view of the object?

Users and user applications will still need to interact with Fedora or iRODS directly. This is particularly true of research (grid) applications having large datasets. Selected metadata will need to be duplicated in both products to access content, represent arrangements, and preserve integrity and authenticity. Direct interaction by users or user applications with either Fedora or iRODS

might require both products to synchronize or update metadata.

These interactions may trigger policy invocations. For example, Fedora may trigger policy invocation indirectly when interacting with a file (CRUD) or directly through a Fedora custom service. Conversely, iRODS' micro-services can call Fedora services to provide feedback in the system.

A more comprehensive "Concept of Operations" document will be prepared as part of the PoDRI project. The following set of questions is drawn from our current understanding of the operational scenarios:

- How will the collection structure be represented in the two products?
- How will Fedora be initialized for existing content in iRODS?
- How will Fedora be informed of content or metadata changes initiated directly in iRODS?
- How can content or metadata from Fedora be accessed by iRODS services?

4. Enabling Use Cases

Five enabling use cases have been identified for the Fedora-iRODS integration. These use cases are:

1. New content ingest via Fedora
2. New content ingest via iRODS
3. Bulk registration from iRODS into Fedora
4. Update of content or metadata via Fedora
5. Update of content or metadata via iRODS

We describe the first two use cases in this paper; a full discussion of all the use cases is beyond the scope of this paper, and will be developed and documented throughout the project's lifecycle. While these use cases do not, by themselves represent policy management operations, they are prerequisites for enabling policy-driven operations and represent demonstrations of policy interoperability between repositories.

4.1 New Content Ingest via Fedora

Current users of Fedora will want to continue ingesting into Fedora. Users are also likely to use Fedora features to add and relate rich metadata including policy, provenance and authenticity information. As shown in Figure 1, when new content is ingested into Fedora, it is able to capture the metadata it needs for its operation. Digital content (or user-defined metadata) is either pulled in by Fedora or pushed to Fedora and stored in individual files. The file containing the FDO (FOXML) and the content files are subsequently stored in iRODS with no storage directly managed by Fedora.

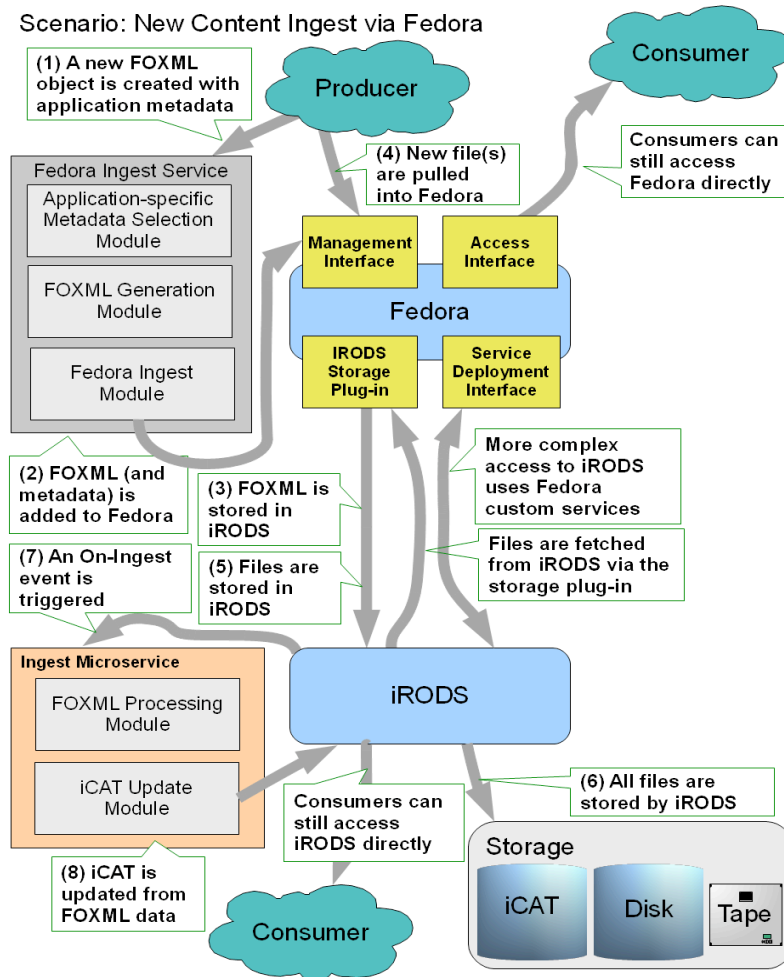


Figure 1: New Content Ingest via Fedora

Selected metadata is collected by Fedora during the ingest process and stored in an internal system index implemented using a relational database. This database is used only to remove latency (speed up) access to content or bindings to services (formerly called disseminators). Optionally, metadata or notifications can be sent to index services such as semantic triplestores, search engines and OAI-PMH harvesters.

The Carolina Digital Repository (CDR) is using Solr/Lucene as the indexing and search engine for discovery of ingested content. Metadata is extracted during the ingest process from MODS and FOXML files.

Objects ingested via Fedora and stored in iRODS do not, by default, retain the logical tree structure of the original file system. Instead, CDR preserves the hierarchal structure of the file system via relations in the RDF triple store.

The arrangement of objects is achieved by created FDOs representing the parent and child. The relationship is recorded in RDF (within the RELS-EXT Datastream) using the "isMemberOf" asserted in the child to the parent. The obverse relation "hasMember" is implied but

could be stated explicitly in the parent. These two relations provide a way to build a hierarchical structure for all objects, collections and files. In Fedora, these relations form a "graph" and objects may participate in any number of graphs using other relations and, therefore, are not limited to a single hierarchy. Relationship information can be accessed by introspecting on the FDO or the relations can be indexed into a RDF triplestore [11] and queried by applications to extract a graph for navigating from parent to children as people usually do for a tree structure. Similar methods can be used to navigate any relationship graph.

How will the metadata in iRODS be updated in this use case? Two alternatives being considered are: (1) call a Fedora custom service to update the iCAT; (2) when the FOXML file is ingested, a monitoring rule can trigger an iRODS micro-service to introspect on the FDO to extract the metadata.

4.2 New Content Ingest via iRODS

Current iRODS users will likely want to continue to use iRODS directly to store data objects, particularly in research settings where direct access to storage is

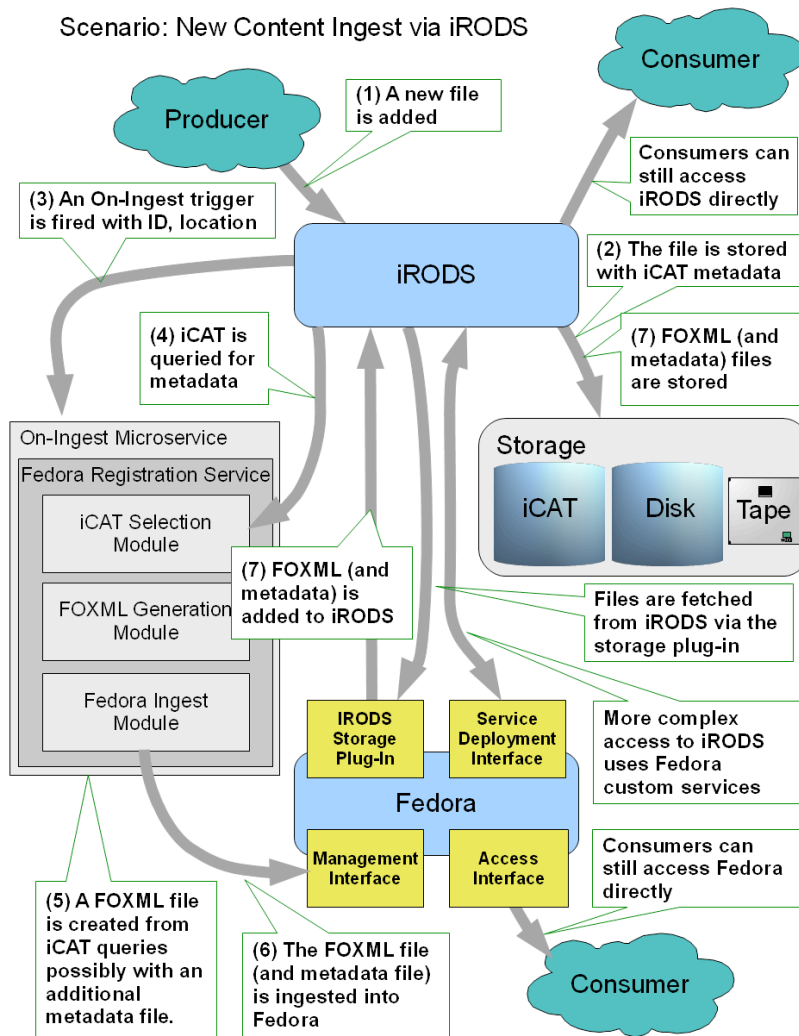


Figure 2: New Content Ingest via iRODS

desired. The digital content (data object) is typically ingested into iRODS as a file operation. In iRODS, the hierarchical relation of a data object and its ancestors are encoded and described explicitly in its global object name. Two questions arise from this scenario. First, how will Fedora be notified of arrival of the new data object? Second, how will an analog to its iRODS hierarchy be represented in Fedora?

A utility is needed to register iRODS files into Fedora. A micro-service could call this utility when triggered by a monitoring rule on the storage operation which would create the FDO for the data object and ingest it into Fedora. The micro-service can be deployed as a rule under the iRODS rule event, 'acPostProcForPut'. Once this rule is activated in an iRODS server, the micro-service can be triggered after each new iRODS data object is created in a specified collection in the iRODS Content Store (see iRODS Storage Module), as depicted in Figure 2. It will create

pre-ingest FOXML for the new data object, querying the iCAT for additional metadata as needed. Within the FOXML, it will create a Datastream containing a reference to the location of the data object within iRODS. It will then ingest the FOXML using Fedora's API-M to create the FDO. This rule is activated once placed in the rule configuration file of an iRODS server. It monitors all file activities in the iCAT catalog and creates an FDO for any newly created iRODS file.

When using iRODS for back-end storage, all FDOs and Datastreams are stored in iRODS as files in one of two collections: FOXML Object Store and iRODS Content Store. Therefore, users can directly access the files containing Fedora metadata through the iRODS interface. On the other hand, files stored in iRODS, whether for an FDO or a Datastream, have both an independent set of iRODS system metadata as well as a set of user-defined metadata. The system metadata contains important information for each replica of an

iRODS file, including the file's location, storage type, audit trail, and associated iRODS rules. The two sets of metadata can be represented as external Datastreams in FOXML and generated dynamically when accessed using the Fedora-iRODS storage module.

As described above, Fedora uses RDF relations to describe the arrangement of objects. This requires the creation of FDOs representing each hierarchical level which has the advantage of enabling the participation of iRODS in the semantic network functionality provided by Fedora. Since iRODS can create a virtual hierarchy, it may not be desirable to instantiate corresponding FDOs. Users can create custom Datastreams as "finding aids"; the virtual hierarchy can be encoded using RDF or any other desired format. Similar to iRODS, parent-child relationships can be modeled as path metadata and stored in the custom Datastream. An application or a Fedora custom service can be used to interpret the format of the Datastream to display the hierarchy [12].

Many of the CDR's core constituencies are the special collections in our libraries. These collections tend to have rich metadata associated with them and have usually undergone preliminary curation. The longer term goal of the repository is to harvest content directly from research-based iRODS data grids. Metadata quality and quantity is typically limited in these collections. Repository outreach and development is concerned not only with identifying and preserving "at risk" collections, but cultivating metadata collection and data curation proactively throughout the research lifecycle.

5. Additional Utilities

We plan to implement two utilities in addition to the functionality described above. First is an updated storage module as an iRODS-specific plug-in to replace Fedora's Low-level Store. Second is a harvester utility which can be used in both bulk registration and for disaster recovery.

5.1 iRODS Storage Module

We plan to store all files in iRODS. This will require an update of the existing iRODS-Fedora Storage Module or build a new module potentially using the Fedora Commons Akubra interface. If a new module is built, using Jargon is being considered. Building a new module would permit research on using it as a feedback path for policy operations including security policies.

When iRODS serves as a storage module for Fedora, current thought is to use two iRODS collections: (1) Fedora Digital Objects (FOXML) in the FOXML Object Store, and (2) content objects (Datastreams) in the iRODS Content Store. They are accessed through a single curator user account in iRODS. This makes it easier to distinguish between policies related to FDOs from those operating on content objects (Datastreams).

This approach, however, differs from the Fedora/Jargon default of storing objects in folders based on timestamp. For the CDR and other existing implementations, a restructuring of objects into the segregated object store will be required. This will alter iRODS based failure recovery mechanisms and integrity audits.

5.2 iRODS Data Harvester for Fedora

The iRODS Data Harvester is an adaptive version of the Data Rebuilder in Fedora. It is used to re-build the object indices from the FOXML Object Store and iRODS Content Store. It does not create any new FOXML objects; rather, it surveys all the objects stored within the FOXML Object Store, verifies the Datastreams inside the iRODS Content Store, and creates the indices in the database used by the Fedora server. The iRODS Data Harvester also builds the necessary RDF data to be stored in the RDF triplestore for the navigation of hierarchical structure.

6. Policy Federation and Migration

The iRODS rule engine provides the capability to apply rules on the data grid side to implement the policies. The Distributed Custodial Archival Preservation Environments (DCAPE) project [13] aims to work with a group of archivists to develop a set of rules to automate many of the administrative tasks associated with the management of archival repositories and validation of their trustworthiness. These DCAPE rules could be applied to different repositories based on the institution's policies. We plan to provide the functionality for users to manage the policies through the Fedora interface and be able to check what rules are in action.

Current implementations, even in data grid environments, depend on local enforcement of policies and typically do not consider the larger framework of uniform policy implementation across heterogeneous repositories. If policies are expressed in the language of ISO-MOIMS or DCAPE criteria, we have a clear model for identification of machine-actionable rules.

Stored as Fedora Service Definitions, the policies will have unique service deployment bindings for each data storage system. Our demonstration storage implementation is iRODS, but other storage environments may be supported by changing deployment mechanisms.

The CDR is developing a policy management framework based on a machine interpretable series of actions across repositories in a data grid. Implementation of new policy requires identification of machine-actionable components and mapping to specific, testable deployment mechanisms.

7. Summary

In this paper, we introduced the Policy-Driven Repository Interoperability (PoDRI) project investigating interoperability mechanisms between repositories at the policy level. The rationale for using iRODS and Fedora to demonstrate key features of a distributed policy-driven management architecture was described. Four scenarios that will be demonstrated as part of the project were enumerated. We have identified five enabling use cases and described two that are needed for the demonstration scenarios along with two key utilities planned for development. We also introduced work on policy federation and migration. PoDRI is an applied research project and its details will change as we develop a greater understanding of the methods for policy-driven interoperability.

8. Acknowledgements

This project is funded by IMLS grant LG-06-09-0184-09 as part of the 2009 National Leadership Grants NLG Library-Research and Demonstration, awarded to the University of North Carolina at Chapel Hill. Project Director is Richard Marciano. Collaborators at UNC / SILS include: Alex Chassanoff, Chien-Yi Hou, Reagan Moore, and Helen Tibbo. At UNC / Libraries: Steve Barr, Greg Jansen, Will Owen, and Dave Pcolar. At UNC / RENCI: Leesa Brieger. At UCSD: Bing Zhu. At DuraSpace and Cornell Information Sciences: Daniel Davis and Sandy Payette. Finally, at the University of Maryland iSchool: Bruce Ambacher.

9. References

- [1] DuraSpace, "PLEDGE Project," <http://fedora-commons.org/confluence/x/WSDS>
- [2] iRODS: Data Grids, Digital Libraries, Persistent Archives, and Real-time Data Systems. <http://www.irods.org>
- [3] R. Moore, A. Rajasekar, M. Wan, and W. Schroeder, "Policy-Based Distributed Data Management Systems," The 4th International Conference on Open Repositories, Atlanta, Georgia, May 19, 2009.
- [4] Fedora Commons, <http://www.fedora-commons.org>
- [5] Fedora Commons, "Fedora Repository Documentation," <http://fedora-commons.org/confluence/x/AgAU>
- [6] DART, University of Queensland, "Fedora-SRB Database integration module," <http://www.itee.uq.edu.au/~eresearch/projects/dart/outcomes/FedoraDB.php>
- [7] A. Aschenbrenner, B. Zhu, iRODS, "iRODS-Fedora Integration," <http://www.irods.org/index.php/Fedora>
- [8] M. Hedges, A. Hazan, and T. Blanke, "Management and Preservation of Research Data with iRODS," *Proceedings of the ACM first workshop on CyberInfrastructure: information management in eScience*, Lisbon, Portugal, pp. 17-22, 2007 doi: <http://doi.acm.org/10.1145/1317353.1317358>
- [9] DuraSpace, "The Fedora Digital Object Model," <http://fedora-commons.org/confluence/x/dgBI>

[10] B. Zhu, R. Marciano, and R. Moore, "Enabling Inter-repository Access Management between iRODS and Fedora," The 4th International Conference on Open Repositories, Atlanta, Georgia, May 19, 2009.

[11] Wikipedia, "Triplestore,"

<http://en.wikipedia.org/wiki/Triplestore>

[12] DuraSpace, "The Content Model Architecture,"

<http://fedora-commons.org/confluence/x/gABI>

[13] DCAPE, "Distributed Custodial Archival Preservation Environments", an NHPRC-funded project, <http://dcape.org>