

Rule-Oriented Data Management Infrastructure

Reagan W. Moore
San Diego Supercomputer Center

moore@sdsc.edu

<http://www.sdsc.edu/srb>

Funding: NSF ITR / NARA

Distributed Data Management

- **Driven by the goal of improving access to data, information, and knowledge**
 - Data grids for sharing data on an international scale
 - Digital libraries for publishing data
 - Persistent archives for preserving data
 - Real-time sensor systems for recording data
 - Collections for managing simulation output
- **Identified fundamental concepts required by generic distributed data management infrastructure**
 - Data virtualization - manage properties of a shared collection independently of the remote storage systems
 - Trust virtualization - manage authentication, authorization, auditing, and accounting independently of the remote storage systems

Extremely Successful

- After initial design, worked with user communities to meet their data management requirements with the **Storage Resource Broker (SRB)**
 - Used collaborations to fund the continued development
 - Averaged 10-15 simultaneous collaborations for ten years
 - Worked with:
 - Astronomy Data grid
 - Bio-informatics Digital library
 - Ecology Collection
 - Education Persistent archive
 - Engineering Digital library
 - Environmental science Data grid
 - High energy physics Data grid
 - Humanities Data Grid
 - Medical community Digital library
 - Oceanography Real time sensor data
 - Seismology Digital library
 - ...

History - Scientific Communities

- **1995 - DARPA Massive Data Analysis Systems**
- 1997 - DARPA/USPTO Distributed Object Computation Testbed
- 1998 - NSF National Partnership for Advanced Computational Infrastructure
- 1998 - DOE Accelerated Strategic Computing Initiative data grid
- 1999 - NARA Transcontinental Persistent Archive Prototype
- 2000 - NASA Information Power Grid
- 2001 - NLM Digital Embryo digital library
- 2001 - DOE Particle Physics data grid
- 2001 - NSF Grid Physics Network data grid
- 2001 - NSF National Virtual Observatory data grid
- 2002 - NSF National Science Digital Library persistent archive
- 2003 - NSF Southern California Earthquake Center digital library
- 2003 - NIH Biomedical Informatics Research Network data grid
- 2003 - NSF Real-time Observatories, Applications, and Data management Network
- **2004 - NSF ITR, Constraint based data systems**
- 2005 - LC Digital Preservation Lifecycle Management
- 2005 - LC National Digital Information Infrastructure and Preservation program

Collaborations - Preservation

- 1. MDAS:** 1995-1997, DARPA - SDSC
Integration of DB and Archival Storage. Support for shared collections
- 2. DOCT:** 1997-1998, DARPA/USPTO - SDSC, SAIC, U Va, ODU, UCSD, JPL
Distributed object computation testbed. Creation of USPTO patent digital library.
- 3. NARA:** 1998 - , NARA - U Md, GTech, SLAC, UC Berkeley
Transcontinental Persistent Archive Prototype based on data grids.
- 4. IP2:** 2002-2006, NHPRC/SHRC/NSF - UBC and others.
InterPARES 2 collaboration with UBC on infrastructure independence
- 5. PERM:** 2002-2004, NHPRC - Michigan, SDSC
Preservation of records from an RMA. Interoperability across RMAs.
- 6. UK e-Science data grid:** 2003-present, - CCLRC, SDSC
Federation of independent data grids with a central archive repository
- 7. LoC:** 2003-2004, LoC - SDSC, LOC
Evaluation of use of SRB for storing America Memory collections
- 8. NSDL:** 2003-2007, NSF - Cornell, UCAR, Columbia, SDSC
Persistent archive of material retrieved from web crawls of NSDL URLs
- 9. ICAP:** 2003-2006, NHPRC - UCSD,UCLA,SDSC
Exploring the ability to compare versions of records, run historical queries
- 10. UCSD Libraries:** 2004- , - UCSD Libraries, SDSC
Development of a preservation facility that replicates collections
- 11. PAT:** 2004-2006, NHPRC - Mi,Mn,Ke,Oh,Slac,SDSC
Demonstration of a cost-effective system for preserving electronic records.

Collaborations - Preservation

12. **DSpace:** 2004-2005, NARA - MIT, SDSC, UCSD Libraries
Digital library. This is an explicit integration of DSpace with the SRB data grid.
13. **PLEDGE:** 2005-2006, NARA - MIT, SDSC, UCSD Libraries
Assessment criteria for trusted digital repositories.
14. **Archivist Workbench:** 2000-2003, NHPRC - SDSC
Methodologies for preservation & access of software- dependent electronic records
15. **NDIIPP:** 2005-2008, LoC - CDL, SDSC
Preservation of selected web crawls, management of distributed collections
16. **DIGARCH:** 2005-2007, NSF - UCTV, Berkeley, UCSD Libraries, SDSC
Preservation of video workflows
17. **e-Legislature:** 2005-2007, NSF - Minnesota, SDSC
Preserving the records of the e-Legislature
18. **VanMAP:** 2005-2006, UBC - UBC, Vancouver
Preserving the GIS records of the city of Vancouver
19. **Chronopolis:** 2005-2006, NARA - SDSC, NCAR, U MD,
Develop preservation facility for collections
20. **eLegacy:** 2006-2008, NHPRC - California
Preserving the geospatial data of the state of California
21. **CASPAR:** 2006 - , 17 EU institutions
Development of representation information for records stored in a SRB data grid.
22. **LoC Storage:** 2006-2007, LoC - SDSC, UCSD libraries
Demonstration of the systems needed to manage remote storage of digital data collections.
23. **IMLS:** 2006-2008, IMLS - UCHRI, SDSC
California's redlining archives testbed (under consideration for funding)

US Academic Institutions (2005)

Project	Institution
National Virtual Observatory	Caltech
Cooperative Institute for Research in Environmental Sciences /Center for Integrated Space Weather Modeling	Colorado University
Institute for Astronomy	Hawaii University
Common Instrument Middleware Architecture, National Middleware Initiative	Indiana University
Indiana University Cyclotron Facility	Indiana University
Dspace digital library	MIT
Atmospheric Sciences Data	NASA
NOAO data grid	National Optical Astronomy Observatory
Web-at-Risk National Digital Information Infrastructure and Preservation Program (CDL)	New York University Libraries
MPI-IO interface	Ohio State University
Computer Science	Oregon State University
BioPilot	Pacific Northwest National Laboratory
TeraGrid project	Purdue University
Fusion Portal	San Diego State University
SDSC Production SRB system	San Diego Supercomputer Center
Texas Advanced Computing Center	Texas University
Network for Earthquake Engineering Simulation	Texas University
NCAR Visualization	UCAR
Network for Earthquake Engineering Simulation	University at Buffalo

San Diego Supercomputer Center

University of California, San Diego

US Academic Institutions (2005)

Project	Institution
Database and Information Systems Laboratory	University of California Davis
Chemistry/Biochemistry	University of California Los Angeles
Consortium of Universities for the Advancement of Hydrologic Science, Inc., Digital Library Server	University of California Merced
Computer Science & Engineering	University of California San Diego
ITR - constraint based data management, Computer Science Department	University of California San Diego
Marine Physical Laboratory	University of California San Diego
National Center for Microscopy and Imaging	University of California San Diego
Cosmology, Physics Department	University of California San Diego
National Center for Microscopy and Imaging, TeleScience	University of California San Diego
University of Florida Research Grid (HPS)	University of Florida
Bioinformatics	University of Kansas
Department of Computer Science	University of Maryland
Network for Earthquake Engineering Simulation	University of Minnesota
Library archive	University of Pittsburgh
Rapid Unified Generation of Urban Databases (RUGUD)	US Army Research Activity
P2Tools Design & Development Team Leader	US Environmental Protection Agency
EPA Data Grid initiative	US Environmental Protection Agency
Government Agency	US Navy
Oceanography collections	Woods Hole Oceanographic Institute

International Institutions (2005)

Project	Institution
Data mangement project	British Antarctic Survey, UK
eMinerals	Cambridge e-Science Center, UK
Sickkids Hospital in Toronto	Canada
Welsh e-Science Centre	Cardiff University, UK
Visualization in scientific computing	Chinese Academy of Science, China
Australian Partnership for Advanced Computing Data Grid	Commonwealth Scientific and Industrial Research Organization, Australia
Consorzio Interuniversitario per il Calcolo Automatico dell'Italia Nord Orientale, HPC-EUROPA project	Italy
Center for Advanced Studies, Research, and Development	Italy
LIACS(Leiden Inst. Of Comp. Sci)	Leiden University, The Netherlands
Australian Partnership for Advanced Computing Data Grid	Melbourne, Australia
Monash E-Research Grid	Monash University, Australia
Computational Materials Science	Nanyang Technological University, China
Virtual Tissue Bank	Osaka University, Japan
Cybermedia Center	Osaka University, Japan
Belfast e-Science Centre	Queen's University, UK
Information Technology Department	Sejong University, South Korea
Nanyang Centre for Supercomputing	Singapore
National University (Biology data grid)	Singapore
Swiss Federal Institute (Ecole Polytechnique Federale de Lausanne)	Switzerland

San Diego Supercomputer Center

University of California, San Diego

International Institutions (2005)

Project	Institution
CERN- GridFTP	Switzerland
Protein structure prediction	Taiwan University, Taiwan
Trinity College High Performance Computing (HPC-Europa)	Trinity College, Ireland
National Environment Research Council	United Kingdom
Universidad Nacional Autonoma de Mexico Grid	Universidad Nacional Autonoma de Mexico
Parallab(HPC-EUROPA project)	University of Bergen, Norway
Physics Labs	University of Bristol, UK
Laboratory for Bioimages and Bioengineering	University of Genoa, Italy
Bio Lab	University of Genoa, Italy
School Computing	University of Leeds, UK
Dept. of Computer Science	University of Liverpool, UK
Worldwide Universities Network	University of Manchester, UK
Large Hadron Collider Computing Grid	University of Oxford, UK
Computational Modelling	University of Queensland, Australia
Instituto do Coracao	University of Sao Paulo, Brazil
White Rose Grid	University of Sheffield, UK
Australian Partnership for Advanced Computing Data Grid	University of Technology, Australia
Computational Chemistry environment	University of Zurich, Switzerland
Australian Partnership for Advanced Computing Data Grid	Victoria, Australia

Extremely Successful

- **Storage Resource Broker Production Environment**
 - Respond to user requests for help
 - SRB-chat Email
 - Email archive
 - Bugzilla bug/feature request list
 - Hot page for server status
 - Wiki web page with all documentation, user contributed software
 - Continue development of new features, ports
 - CVS repository for all source code changes
 - Daily build and test procedure
 - NMI testbed builds before each release
 - Average of four releases per year
 - Supporting projects now ending or have ended
 - (NSF ITR, DOE, NASA)
- **How can such systems be sustained for use by the academic community?**

Recent SRB Releases

- 3.4.2 June 26, 2006
- 3.4.1 April 28, 2006
- 3.4 October 31, 2005
- 3.3.1 April 6, 2005
- 3.3 February 18, 2005
- 3.2.1 August 13, 2004
- 3.2 July 2, 2004
- 3.1 April 19, 2004
- 3.0.1 December 19, 2003
- 3.0 October 1, 2003
- 2.1.2 August 12, 2003
- 2.1.1 July 14, 2003
- 2.1 June 3, 2003
- 2.0.2 May 1, 2003
- 2.0.1 March 14, 2003
- 2.0 February 18, 2003

Date	5/17/02		6/30/04			7/10/06		
Project	GBs of data stored	10000s of files	GBs of data stored	10000s of files	Users with ACLs	GBs of data stored	10000s of files	Users with ACLs
Data Grid								
NSF / NVO	17,800	5,139	51,380	8,690	80	106,070	14,001	100
NSF / NPACI	1,972	1,083	17,578	4,694	380	35,109	7,240	380
Hayden	6,800	41	7,201	113	178	8,013	161	227
Pzone	438	31	812	47	49	23,475	13,576	68
NSF / LDAS-SALK	239	1	4,562	16	66	143,429	165	67
NSF / SLAC-JCSG	514	77	4,317	563	47	17,595	1,814	55
NSF / TeraGrid			80,354	685	2,962	267,422	6,970	3,267
NIH / BIRN			5,416	3,366	148	17,155	16,116	385
Digital Library								
NSF / LTER	158	3	233	6	35	257	41	36
NSF / Portal	33	5	1,745	48	384	2,620	53	460
NIH / AfCS	27	4	462	49	21	733	94	21
NSF / SIO Explorer	19	1	1,734	601	27	2,653	1,159	27
NSF / SCEC			15,246	1,737	52	168,689	3,544	73
Persistent Archive								
NARA	7	2	63	81	58	2,999	2,033	58
NSF / NSDL			2,785	20,054	119	5,698	50,600	136
UCSD Libraries			127	202	29	190	208	29
NHPRC / PAT						1,888	521	28
TOTAL	28 TB	6 mil	194 TB	40 mil	4,635	804 TB	118 mil	5,417

Standards Effort

- **Global Grid Forum - Grid Interoperability Now**

- **Organizers:** Erwin Laure (Erwin.Laure@cern.ch)
Reagan Moore (moore@sdsc.edu)
Arun Jagatheesan (arun@sdsc.edu) - grid coordination
Sheau-Yen Chen (sheauc@sdsc.edu) - data grid administrator
Chien-Yi Hou (chienyi@sdsc.edu) - collection administrator

- **Goals:**

- Demonstrate federation of 17 SRB data grids (shared name spaces)
- Demonstrate replication of a collection

- **Global Grid Forum - Preservation Environments Research Group**

- **Organizers:** Reagan Moore (moore@sdsc.edu)
Bruce Barkstrom

- **Goals:**

- Demonstrate creation of preservation environments based on data grid technology
- Demonstrate federation of preservation environments

SRB Data Grid Federation Status

Data Grid	Country	SRB version	Demouser ggfsdsc	SRB Zone name	Storage Resource Logical Name	I/O MB/sec
APAC	Australia	3.4.0-P	yes	AU	StoreDemoResc_AU	3.9
NOAO	Chile/US	3.4.2	yes	noao-ls-t3-z1	noao-ls-t3-fs	
ChinaGrid	China	CGSP-II	(software)			
RNP	Brazil	3.4.1-P2	yes	GGF-RNP	demoResc	
UERJ	Brazil	3.4.1-P2	yes	UERJ-HERPGrid	demoResc	
IN2P3	France	3.4.0-P	yes	ccin2p3	LyonFS4	[25.]
DEISA	Italy	3.4.0-P	yes	DEISA	demo-cineca	
KEK	Japan	3.4.0-P	yes	KEK-CRC	rsr01-ufs	7.4
SARA	Netherlands	3.4.0-P	yes	SARA	SaraStore	
IB	New Zealand	3.4.1	yes	aucklandZone	aucklandResc	(0.3)
ASGC	Taiwan	3.4.0-P	yes	TWGrid	SDSC-GGF_LRS1	(0.1)
NCHC	Taiwan	3.4.0-P	yes	ecogrid	ggf-test	
CCLRC	UK	3.4.0-P	yes	tdmg2zone		
IB	UK	3.4.1	yes	avonZone	avonResc	
WunGrid	UK	3.3.1	(hardware)	SDSC-wun	sfs-tape	
LCDRG	US	3.4.1-P2	Yes	LCDRG-GGF	demoResc	
Purdue	US	3.4.0-P	yes	Purdue	uxResc1	(2.5)
Teragrid	US	3.4.1-P2	yes	SDSC-GGF	sfs-disk	
U Md	US	3.4.0-P	yes	umiacs	narasrb02-unix1	

Data Grid Federation

- **Builds on:**

- **Registry for data grid names** - ensures each data grid has a unique identity
- **Trust establishment** - explicit registration command issued by the data grid administrator of each data grid
- **Peer-to-peer server interaction** - each SRB server can respond to commands from any other SRB server, provided trust has been established between the data grids
- **Administrator controlled registration of name spaces** - each grid controls whether they will share user names, file names, replicate data, replicate metadata or allow remote data storage
- **Shibboleth style user authentication** - a person is identified by
/Zone-name/user-name.domain-name.
Authentication is done by the home zone. No passwords are shared between zones.
- **Local authorization** - operations are under the control of the zone being accessed, including controls on access to files, storage resources, metadata and user quotas. Owners of data can set access controls for other persons

Federation Between Data Grids

Data Access Methods (Web Browser, Scommands, OAI-PMH)

Data Collection A

Data Collection B

Data Grid

- Logical resource name space
- Logical user name space
- Logical file name space
- Logical context (metadata)
- Control/consistency constraints

Data Grid

- Logical resource name space
- Logical user name space
- Logical file name space
- Logical context (metadata)
- Control/consistency constraints

Access controls and consistency constraints on cross registration of name spaces

Observing Operations Implementation: ***EarthScope/USArray and ROADNet***

Future Proposals

LOOKING Review
Calit2, UCSD
5-7 July 2006

Frank Vernon
UCSD

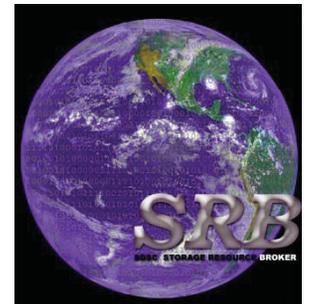
Real-time Observatory Cyberinfrastructure Challenges

- **Scalability**
 - Dynamic station deployment
 - Data integration with remote archives
- **Extensibility**
 - New sensor types
 - New data types
- **Operational Issues**
 - Multiple communication types
 - Dynamic IP assignment for instruments
 - Intermittent communications
- **Observatory interaction**
 - Real time data integration with other observatories

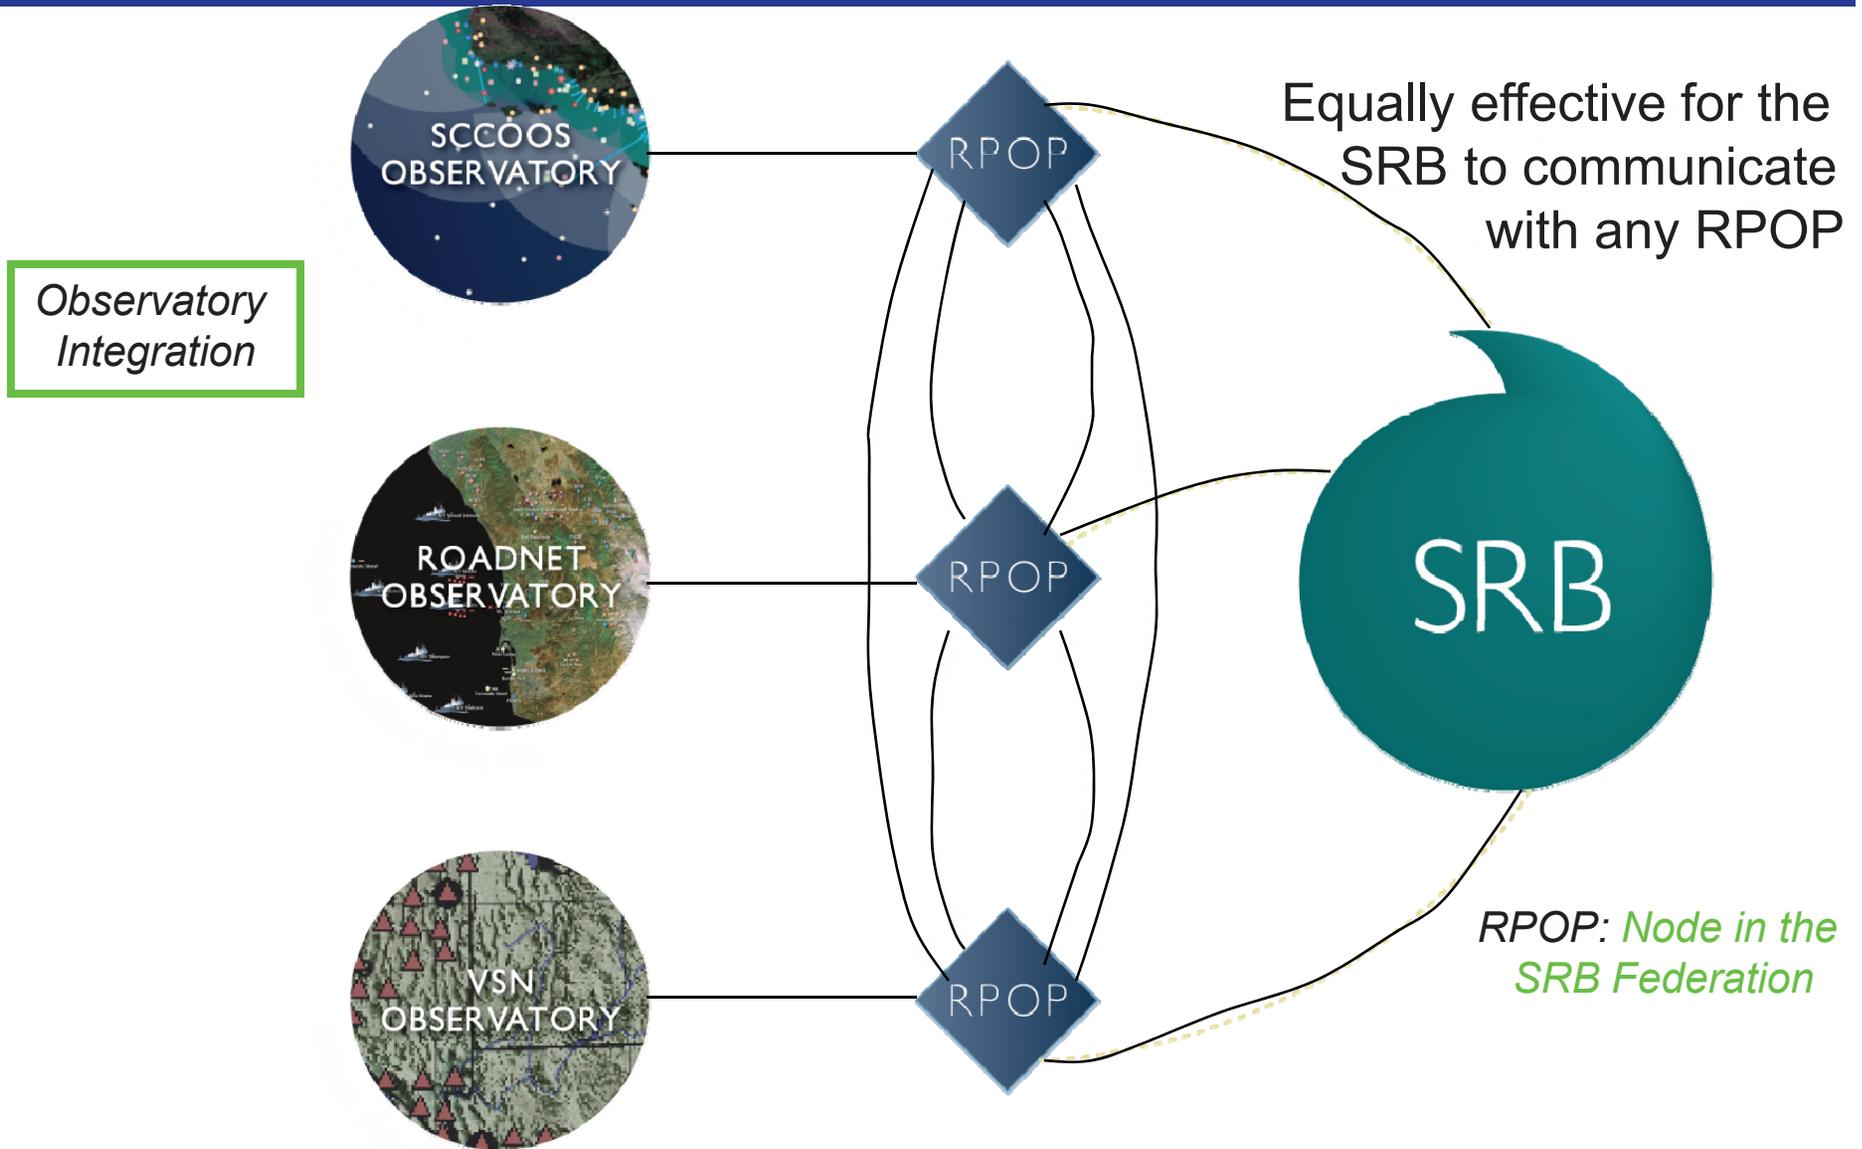
ROADNet Point Of Presence

- “RPOP”

- Embedded real-time processing system
- Integrated with Storage Resource Broker
- Sophisticated FEDERATION NODE
 - *Data Acquisition tools*
 - *Data concentration and distribution tools*
 - *Data processing tools*
- Sun Fire server machines
- Being installed on oceanographic research vessels



RPOP: multiple grid paradigms



Observatory
Integration

Equally effective for the
SRB to communicate
with any RPOP

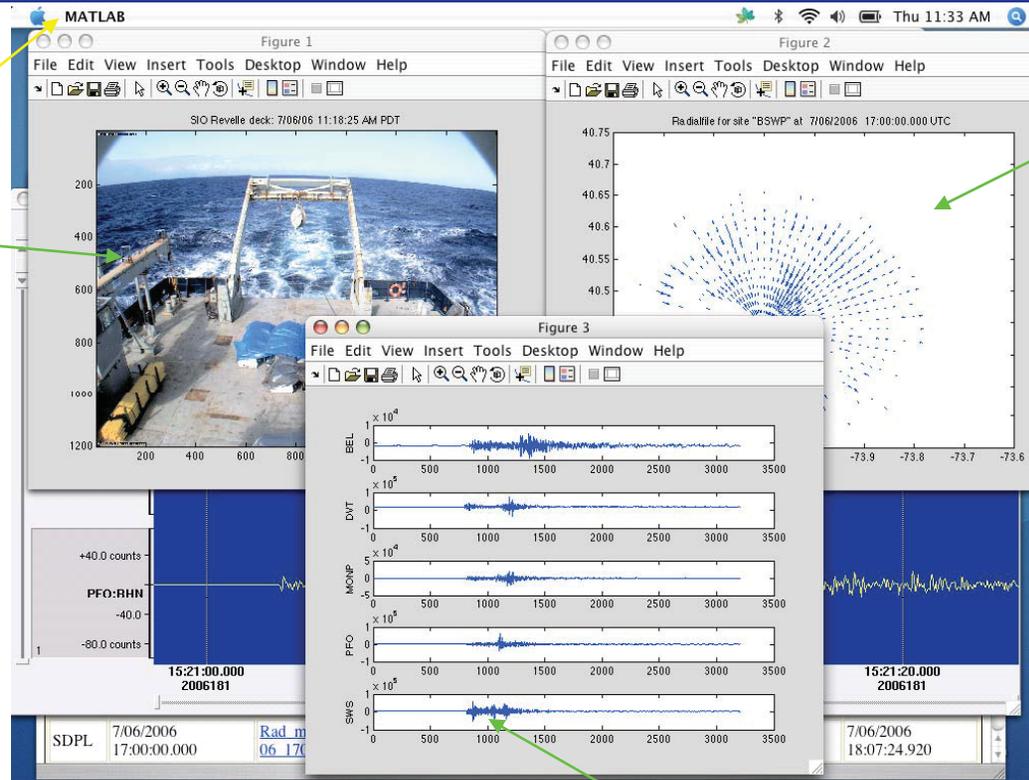
SRB

RPOP: Node in the
SRB Federation

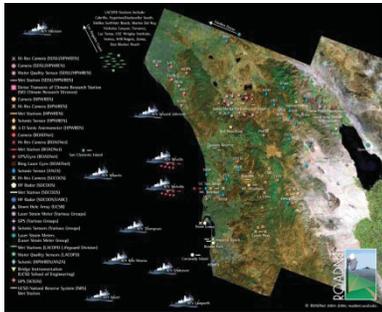
RPOP: Node in the underlying data grid

Tri-observatory Federation

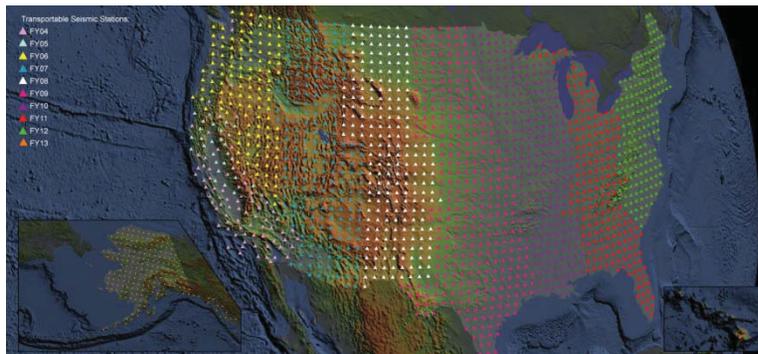
Southern California Coastal Ocean Observing System



ROADNet



EarthScope / USArray



- *Matlab tools*
- *Observatory-grade analysis tools*
- *Web access*

From NSF LOOKING Review 7/6/06, Calit2
University of California, San Diego

San Diego Supercomputer Center

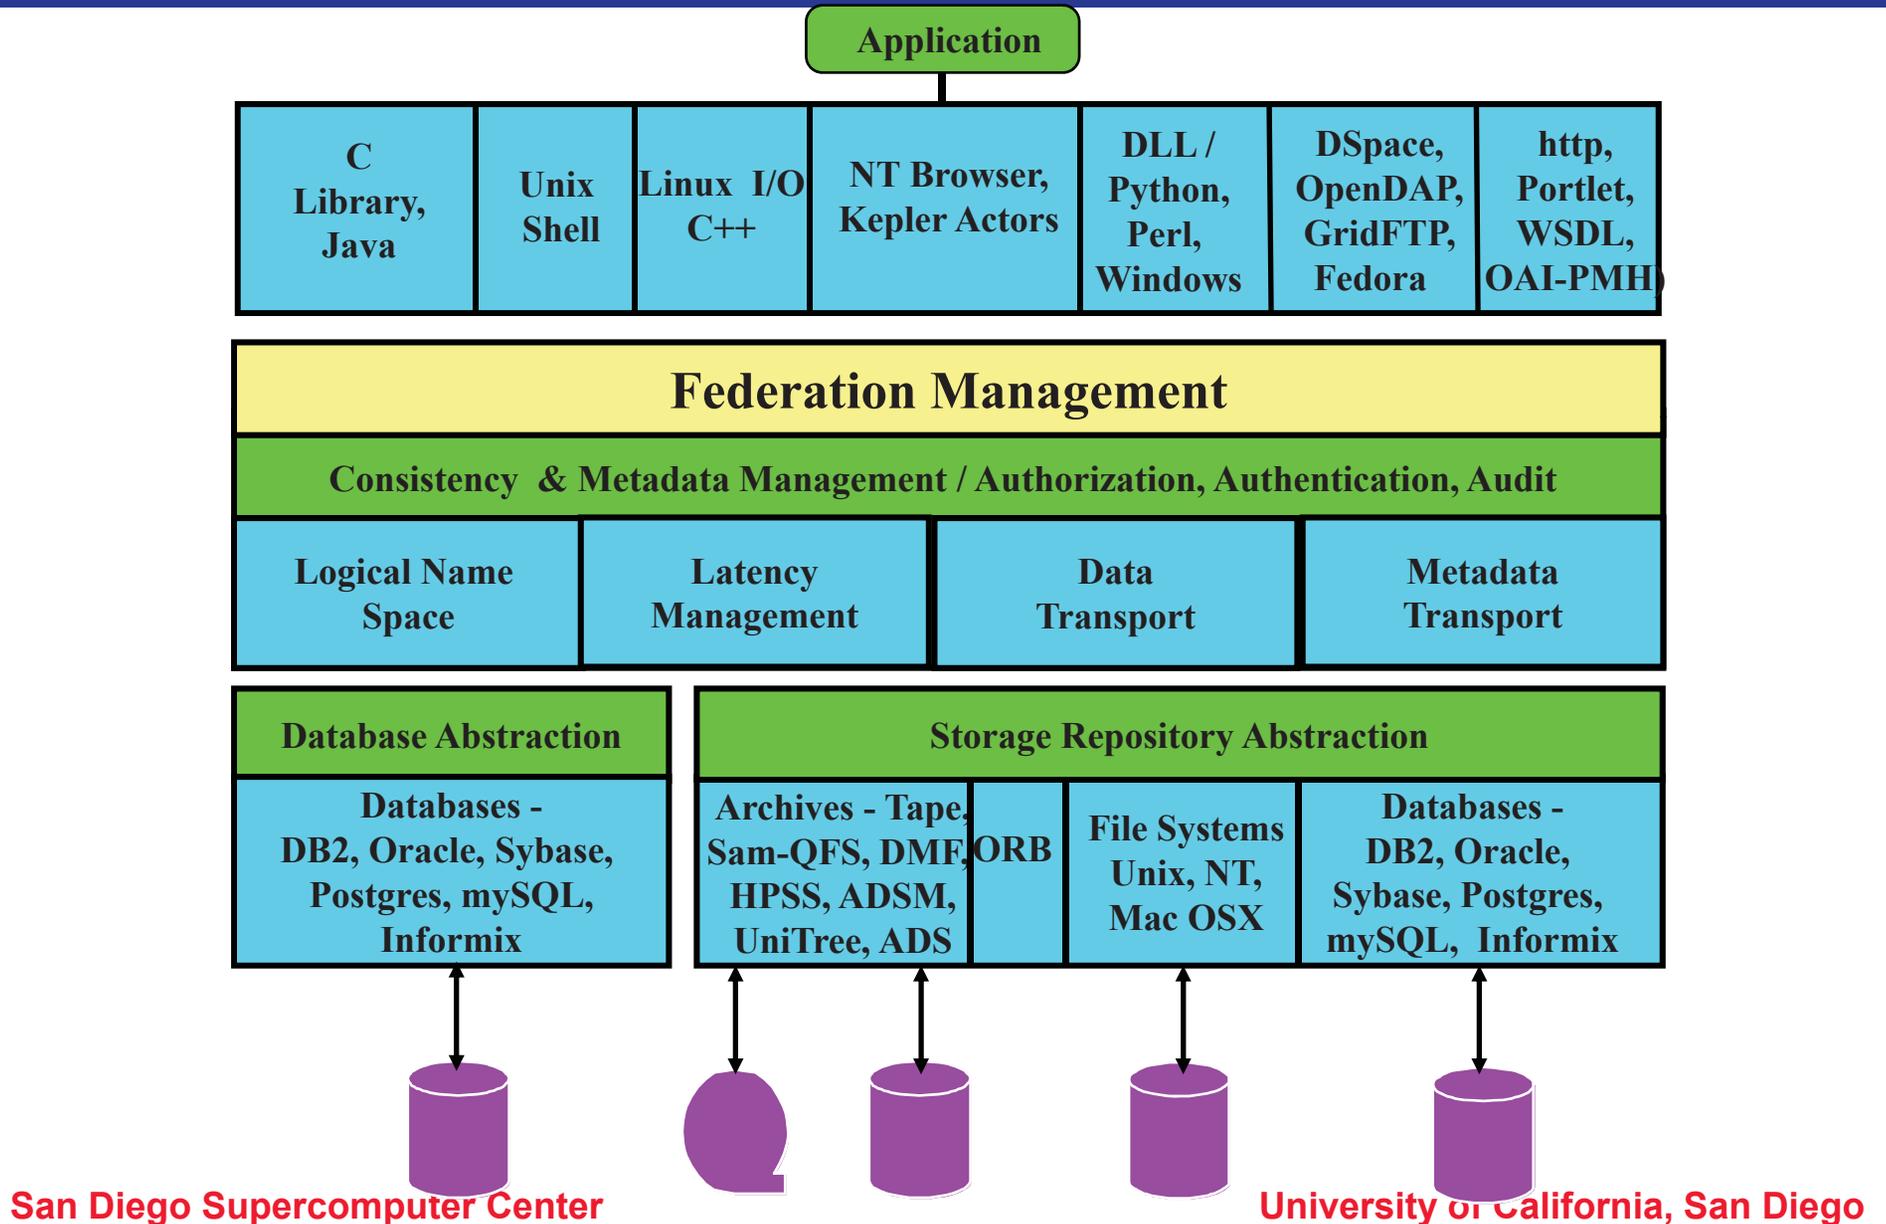
Cognitive Science Collaboratory

- **The NSF-funded Dynamic Learning Center**
 - Multi-institution group of scientists and educators
 - Investigate the role of time and timing in learning
- **Composed of four center initiatives**
 - Dynamics in the external world
 - Dynamics intrinsic to the brain
 - Dynamics of the muscles and body
 - Dynamics of learning
- **Data sharing facility**
 - Rules to validate enforcement of IRB policies
 - Shared collections
 - Publication of results
 - Archiving of data

Research Agenda

- **Require two levels of virtualization for managing operations**
 - Map from operations requested by client
 - To micro-services that are implemented by data grid
 - To operations executed on remote storage systems
- **Require two levels of virtualization for managing data**
 - Map from physical file naming used by storage system
 - To logical name space managed by the shared collection
 - To federated name space managed by federation of shared collections

Storage Resource Broker 3.4.2



Fundamental Data Management Concepts

- **Data virtualization**

- **Management of name spaces**

- Logical name space for users
- Logical name space for storage resources
- Logical name space for digital entities (files, URLs, SQL, tables, ...)
- Logical name space for metadata (user defined attributes)

- **Decoupling of access mechanisms from storage protocols**

- Standard operations for interacting with storage systems (80)
 - Posix I/O, bulk operations, latency management, registration, procedures, ...
- Standard client level operations for porting preferred interface (22)
 - C library calls, Unix commands, Java class library
 - Perl/Python/Windows load libraries, Perl/Python/Java/Windows web browsers, WSDL, Kepler workflow actors, DSpace and Fedora digital libraries, OAI-PMH, GridSphere portal, I/O redirection, GridFTP, OpenDAP, HDF5 library, Semplar MPI I/O, Cheshire

- **Management of state information resulting from standard operations**

Fundamental Data Management Concepts

• Trust virtualization

- Collection ownership of all deposited data
- Users authenticate to collection, collection authenticates to remote storage system
- Collection management of access controls
 - Roles for administration, read, write, execute, curate, audit, annotate
 - ACLs for each object
 - ACLs on metadata
 - ACLs on storage systems
 - Access controls remain invariant as data is moved within shared collection
- Audit trails
- End-to-end encryption

Research Objectives

- **What additional levels of virtualization are required to support advanced data management applications?**
- **Observe that each community imposes different management policies.**
 - Different criteria for data disposition, access control, data caching, replication
 - Assertions on collection integrity and authenticity
 - Assertions on guaranteed data transport
- **Need the ability to characterize the management policies and validate their application**

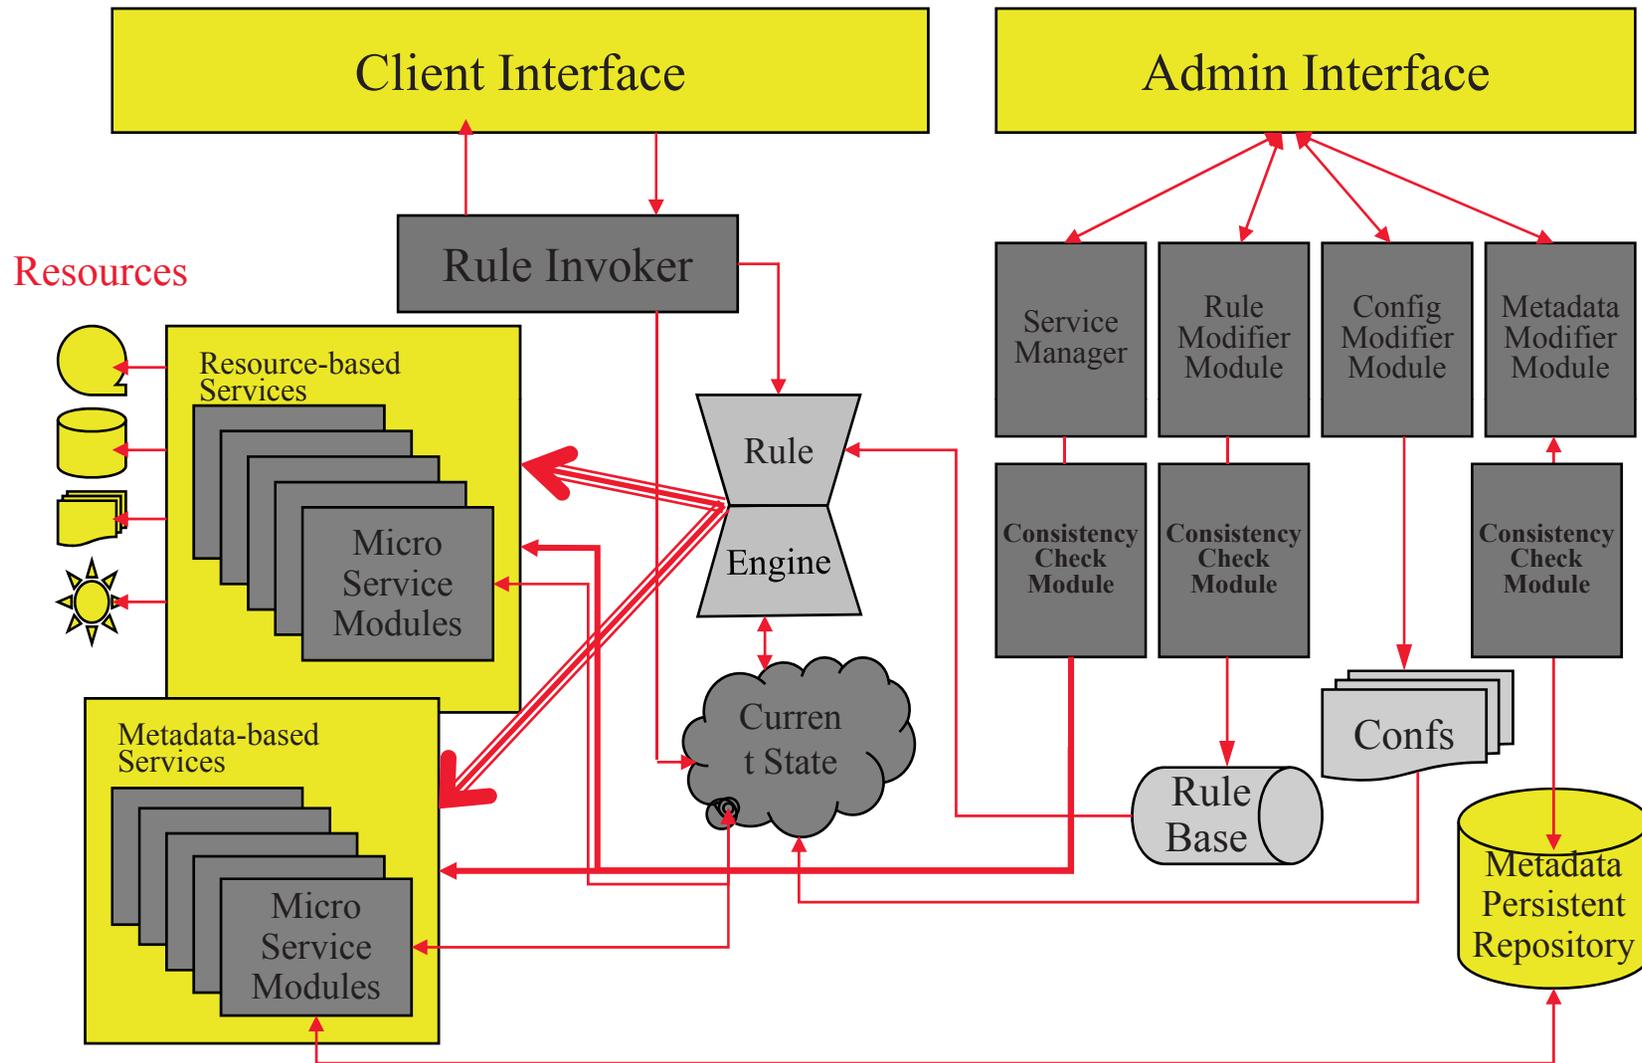
Levels of Virtualization

- **Require metadata (state information, descriptive metadata) for six name spaces**
 - Logical name space for users
 - Logical name space for digital entities (files, tables, URLs, SQL,...)
 - Logical name space for resources (storage systems, ORB, archives)
 - Logical name space for metadata (user defined metadata, extensible schema)
 - Logical name space for rules (assertions and constraints)
 - Logical name space for micro-services (data grid actions)
- **Associate state information and descriptive information with each name space**
- **Virtualization of management policies**

integrated Rule-Oriented Data System

- Integrate a rule engine with a data grid
- Map management policies to rules
- Express operations within the data grid as micro-services
- Support rule sets for each collection and user role
- On access to the system:
 - Select rule set (Collection : user role : desired operation)
 - Load required metadata (state information) into a temporary metadata cache
 - Evaluate rule input parameters and perform desired actions
 - Rules cast as Event:Condition:Action sets
 - Rules invoke both micro-services and rules
 - Provide recovery mechanism for each micro-service
 - On completion, load changed state information back into persistent metadata repository

iRODS - integrated Rule-Oriented Data System



Example Rules

```
0 ON    register_data
        IF          $objPath like /home/collections.nvo/2mass/fits-images/*
        DO          cut [nop]
        AND         check_data_type(fits image) [nop]
        AND         get_resource(nvo-image-resource) [nop]
        AND         registerData [recover_registerData]
        AND         addACLForDataToUser(2massusers.nvo,write) [recover_addACLForDataToUser]
        AND         extractMetadataForFitsImage [recover_extractMetadataForFitsImage]

1 ON    register_data
        IF          $objPath like /home/collections.nvo/2mass/*
        DO          get_resource(2mass-other-resource) [nop]
        AND         registerData [recover_registerData]
        AND         addACLForDataToUser(2massusers.nvo,write) [recover_addACLForDataToUser]

2 ON    register_data
        DO          get_resource(null) [nop]
        AND         registerData [recover_registerData]
```

Emerging Preservation Technology

- **NARA research prototype persistent archive demonstrated use of data grid technology to manage authenticity and integrity**
 - Federated data grids
- **Current challenge is the management of preservation policies**
 - Characterize policies as rules
 - Apply rules on each operation performed by the data grid
 - Manage state information describing the results of rule application
 - Validate that the preservation policies are being followed
- **Same challenge exists in grid services**
 - Characterize and apply rules that govern grid service application

ERA Capabilities

- **List of 854 required capabilities:**

- Management of disposition agreements describing how record retention and disposal actions
- Accession, the formal acceptance of records into the data management system
- Arrangement, the organization of the records to preserve a required structure (implemented as a collection/sub-collection hierarchy)
- Description, the management of descriptive metadata as well as text indexing
- Preservation, the generation of Archival Information Packages
- Access, the generation of Dissemination Information Packages
- Subscription, the specification of services that a user picks for execution
- Notification, the delivery of notices on service execution results
- Queuing of large scale tasks through interaction with workflow systems
- System performance and failure reports. Of particular interest is the identification of all failures within the data management system and the recovery procedures that were invoked.
- Transformative migration, the ability to convert specified data formats to new standards. In this case, each new encoding format is managed as a version of the original record.
- Display transformation, the ability to reformat a file for presentation.
- Automated client specification, the ability to pick the appropriate client for each user.

Summary of Mapping to Rules

- Multiple systems need to be integrated:
 - PAWN submission pipeline - 34 operations
 - Cheshire indexing system - 13 operations
 - Kepler workflow - 53 operations
 - iRODS data management - 597 operations
 - Operations facility - the remaining capabilities
- The 597 operations are executed by 174 generic rules
- The analysis identified five types of metadata attributes:
 - Collection metadata - 11 attributes
 - File metadata - 123 attributes
 - User metadata - 38 attributes
 - Resource metadata - 9 attributes
 - Rule metadata - 32 attributes

File Operations

- List files
- Display file (template)
- Set number of items per display page
- Format file
- Delete file
- Delete file authorized
- Delete file copies
- Delete file versions
- Erase file
- Replace file
- Set file version
- Create soft link
- Replicate file
- Synchronize replicas
- Physmove file
- Annotate file
- Access URL
- Regenerate system metadata
- Check vault
- Monitor space used
- Output file
- Register file
- Register collection hierarchy
- Delete collection
- Bulk move files (new hierarchy)
- Queue file for transfer
- Queue file for encrypted transfer
- Output file to media
- Modify file
- Redact file
- Edit file
- Replicate archives
- Monitor resources - hot page
- Track usage
- Set system parameter
- Predict resource requirements
- Inventory resources
- Log event
- Delete event log entry
- Identify data type
- Create access role
- Modify access control
- Generate notification
- Subscribe
- Delete subscription
- Modify subscription
- Suspend subscription
- Resume subscription
- Validate authenticity

Data Management Rules

- Execute rule
- Suspend rule
- Add rule
- Modify rule
- List rules
- List rule metadata
- Validate rule set
- Approve rule
- Queue rule
- List queued rules
- Set queued rule priority
- Adjust max run time
- Estimate service resources
- List metadata
- Get metadata
- Set metadata
- Bulk metadata load
- Delete metadata
- Define extensible schema
- Load extensible schema
- Export metadata
- Query metadata
- Save query
- Select saved query
- Run saved query
- Modify query
- Modify running query
- Save query result set
- Modify query result set
- Delete search results
- Annotate search result
- Sinit - set default workbench interface
- Register user
- Self-registration
- Delete user
- Suspend user
- Activate user
- Add resource
- Remove resource
- Set resource offline
- Set resource online
- Input file

Example Rules - Templates

- DIP format template
- Disposition agreement format template
- Disposition action format template
- Physical location report template
- Inventory report template
- Data movement summary report template
- Access report template
- File migration report template
- Document internal access control template
- AIP format template
- Transfer format template
- Access review determination rule template
- Access review determination report template
- Validate access classification rule template
- File transfer discrepancy report template
- Notification review report template
- Redaction rule template
- Search display template
- File display template (file type)
- Format conversion format template
- Workbench display template
- Request help format template
- System message format template
- Event log display template
- System report format template
- Monitor hot page format template
- Hot page report template
- Create DIP
- Modify DIP
- Application hot page report template
- COTS hot page report template
- Usage workflow report template
- System configuration display template
- Logistics report format template
- Inventory report format template
- Description extraction rule template
- Accounting report rule template
- Accounting report format template

Example Rules - Templates

- Identify template use
- Create template
- Modify template
- Delete template
- List templates
- Approve template
- Check template
- Assign template
- Template-based default setting
- Parse file
- Generate report
- Modify report
- Export record
- Export records
- Create disposition agreement
- Disposition record check
- Modify disposition agreement
- Compare disposition agreements
- Compare access review determinations
- Change review determination
- List review history
- Preservation assessment rule template
- Preservation assessment report format template
- Lifecycle parsing rules template
- Authenticity validation rule template
- Assess preservation
- Modify workbench
- Select workbench
- Create description
- Validate description
- Modify description
- Update description
- Approve description
- Create unique identifier
- Approve disposition agreement
- Validate transfer request
- Validate access classification
- Queue record for destruction
- Certify deletion of records
- Set disposition hold
- Unset disposition hold
- Record disposition action
- Register physical media location (URL)
- Verify transfer properties
- Preservation assessment

RLG/NARA TDR Assessment Criteria

- The assessment criteria can be mapped to management policies.
- The management policies can be mapped to a set of rules whose execution can be automated.
- The rules require definition of input parameters that define the assertion being implemented.
- The execution of the rules generates state information that can be evaluated to verify the assertion result
- The types of rules that are needed include:
 - Specification of assertions (setting rule parameters - flags and descriptive metadata)
 - Deferred consistency constraints that may be applied at any time
 - Periodic rules that execute defined procedures
 - Atomic rules applied on each operation (access controls, audit trails)
- The rules determine the metadata attributes that need to be managed

TDR - 174 Rules

#	Policy layers/types	TDR	Rule or procedure	State info (result of rule application)	Description
4.2	Format	Ž	Periodic rule - check consistency with required formats	List of supported formats and flag for SLA support level for each	Whether file format is accepted, preservation SLA for each accepted format; Also any requirements for quality within format (e.g. compliance with TIF 6.0 acceptance specs)
Ž	Ž	A5.1	Consistency rule - check that deposit agreement exists	Deposit agreement for storage of data specifying access, replicas, consistency checks	If repository manages, preserves, and/or provides access to digital materials on behalf of another organization, it has and maintains appropriate contracts or deposit agreements.
Ž	Ž	B2.1	Consistency rule that AIP definition exists	Statement of characteristics of each AIP	Repository has an identifiable, written definition for each AIP or class of information preserved by the repository
Ž	Ž	B2.2	Consistency rule - check allowed transformative migration is performed	Criteria for allowed transformative migrations	Repository has a definition of each AIP (or class) that is adequate to fit long-term preservation needs
Ž	Ž	B3.9	Set / Update descriptive metadata: Consistency check for changes to allowed transformative migrations	Procedure for updating transformative migration strategy: Audit trail of changes; Consistency check for changes to migration strategy	Repository has mechanisms to change its preservation plans as a result of its monitoring activities
Ž	Ž	B4.2	Consistency rule - check required metadata	Validation that minimum descriptive metadata is present	Repository captures or creates minimum descriptive metadata and ensures that it is associate with the AIP

iRODS Development

- **Open source software**

- 48,000 lines of “C” code
- Implemented 50 remote storage operations
- Implemented 13 client level operations
- Implemented client server model, with improved protocol

- **Standard build procedure**

- Built entire system on NMI testbed at University of Wisconsin

- **Rule engine**

- Nested Event-Condition-Action sets with recovery procedures for each action
- Named rule sets
- Logical name space for rules
- Logical name space for micro-services
- Logical name space for metadata

Rule Engine

- **Declarative Programming** - through a Rule-based Approach along with rule-consistency checks performed to verify rule execution for cycles and other consistency checks.
- **Transparent Processing & Agile Programming** - similar to Business Rules Logic.
- **Event Condition Action (ECA) Paradigm** - similar to active databases.
- **Transactional & Atomic Operations** - Similar to ACID properties of RDBMS. Each rule either succeeds completely or does not change the operational data (both transient and persistent metadata).
- **WorkFlow Paradigm** for defining a sequence of tasks.
- **Service oriented paradigm** based on micro-services and rules.
- **New Programming paradigms** - based on coding micro services and developing workflows (rules) and stitching the microservices at runtime to the requested operation.
- **Abstraction and logical naming at multiple levels:** data, collections, resources, users, metadata, methods, attributes, rules and micro-services
- **Novel management of version control in the execution architecture.** All versions can coexist. Users can apply their versions and rules at the same time to achieve their tasks.
- **Data grid paradigm** providing standard distributed data management functions:
- **Digital library paradigm** providing standard digital library functions:
- **Persistent archive paradigm** providing standard preservation functions:

iRODS Collaboration Areas

- Shibboleth-SRB/iRODS-Cheshire-uK eScience integration
- GSI support
- Time-limited sessions via the one-way hash authentication
- Python Client library
- Java Client library
- A GUI Browser (Java, or Python, or other)
- A driver for HPSS
- A driver for SAM-QFS
- Other drivers?
- Porting to many versions of Unix/Linux
- Porting to Windows
- Support for Oracle as the database
- Support for MySQL as the database
- A way for users to influence rules
- More extensive installation and test scripts
- AIP to aggregate small files
- MCAT to RCAT migration tools
- Extensible Metadata From the client level, User-defined metadata does not appear distinct from system or extensible metadata.
- Query condition/select clustering. Zones/Federation

Research Collaborations - UCSD

- **Creation of custom web interfaces to shared collections**
 - Yannis Katsis
 - Yannis Papakonstantinou
 - App2you collections and displays data
 - Template driven interface development
 - <https://app2you.org/video/tutorial.html>
- **Validation of rule set consistency**
 - Dayou Zhou
 - Alin Deutsch
 - Assert temporal properties of rule execution

More Information

moore@sdsc.edu

SRB:

<http://www.sdsc.edu/srb>

iRODS:

http://www.sdsc.edu/srb/future/index.php/Main_Page