Transcontinental Persistent Archive Prototype

Policy-Driven Data Preservation

Reagan W. Moore
University of North Carolina at Chapel Hill
rwmoore@renci.org
http://irods.diceresearch.org

NSF OCI-0848296 "NARA Transcontinental Persistent Archives Prototype" (2008-2012)
NSF SDCI 0721400 "Data Grids for Community Driven Applications" (2007-2010)

D·I·C·E i·R·O·D·S THE UNIVERSITY of NORTH CAROLINA at CHAPEL HILL UCSD NSF THE NATIONAL ARCHIVES ARCHIVES.GOV SRB

Topics

- Transcontinental Persistent Archive Prototype
  - Use data grid technology to build a preservation environment
  - Conduct research on preservation concepts
    - Infrastructure independence
    - Enforcement of preservation properties
    - Validation of assessment criteria
    - Automation of administrative processes
  - Demonstrate preservation on selected NARA digital holdings
    - Demonstrate generic infrastructure

Why Data Grids?

- Organize distributed data into shared collections
  - Virtualize collection properties
    - Manage retention, disposition, distribution, replication, integrity, authenticity, chain of custody, access controls, provenance, representation information, descriptive information, logical arrangement
  - Infrastructure independence: provide uniform interface to multiple storage systems
    - Manage interactions with Unix, Linux, Mac, Windows based storage systems
    - Enable use of multiple client interfaces across all storage systems
    - Provide scalability mechanisms such as optimized data transport (parallel I/O, single message small file transfer)

Observation # 1

- Data grids support virtualization of collections
  - Preservation is the extraction of records from the environment in which they were created, and the import of the records into a persistent archive
  - The archivist is in control:
    - Manages the construction of an archival form
    - Selects the properties of the collection that will be preserved
    - Selects the preservation assessment criteria

Observation # 2

- Preservation is communication with the future
  - We know that technology in the future will be more sophisticated than technology today
- Data grids manage technology evolution
  - At the point in time when new technology becomes available, both the old and new systems can be accessed simultaneously
  - Automate migration from the old technology to the new technology

Observation #3

- Preservation is the management of communication from the past
  - To make assertions about the preservation policies and procedures that were applied in the past, the persistent archive must manage and enforce consistent policies and procedures
  - Long-term preservation requires periodic validation of assessment criteria to ensure continued trustworthiness

Preservation is an Integral Part of the Data Life Cycle

- Organize project data into a shared collection
  - **Describe record context**
- Publish data in a digital library for use by other researchers
- Preserve reference collection for use by future research initiatives
- Compare new reference collection against prior state-of-the-art data
- Enable data-driven analyses that dynamically optimize research
  - **Use record context to enable analysis**

D·I·C·E i·R·O·D·S THE UNIVERSITY of NORTH CAROLINA at CHAPEL HILL UCSD NSF THE NATIONAL ARCHIVES ARCHIVES.GOV SRB

Overview of iRODS Data System

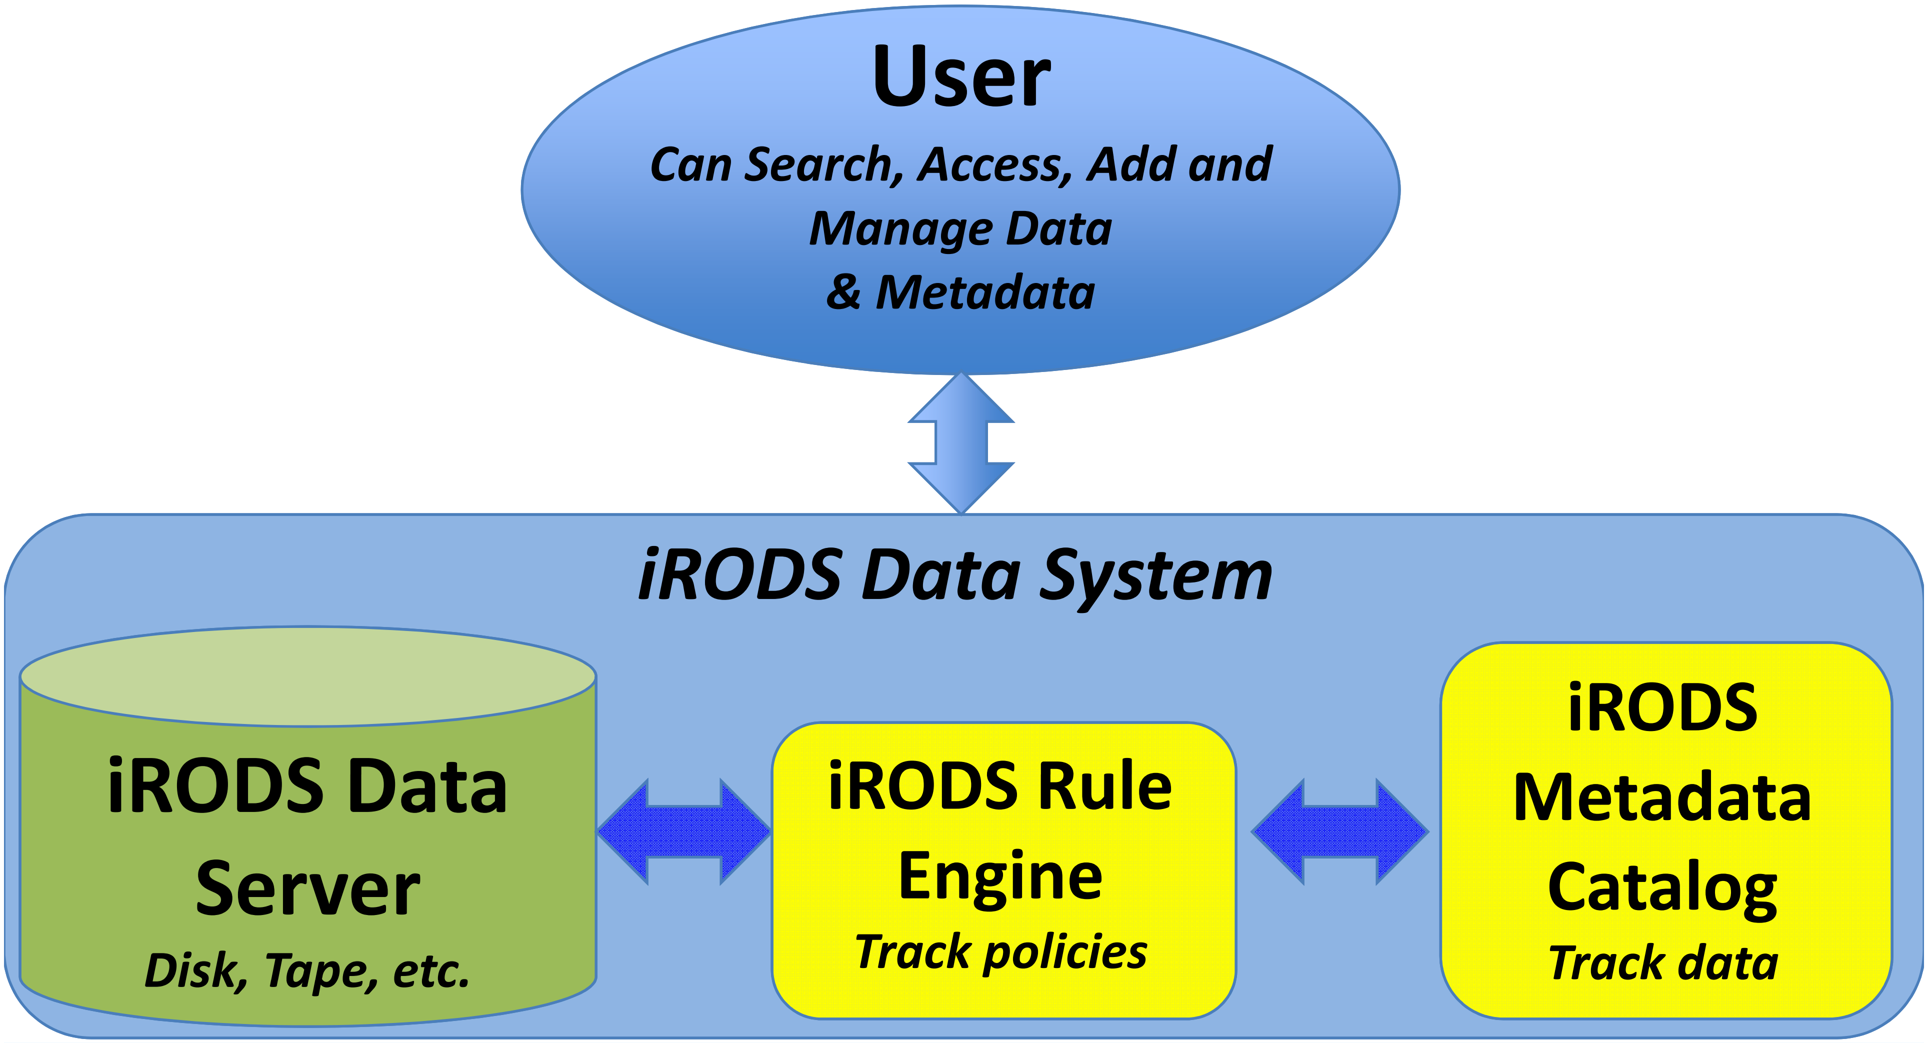

*Access data with Web-based Browser or iRODS GUI or Command Line clients.

Policy-based Data Management

- Turn management policies into computer actionable rules
  - Support dynamic rule base updates
- Turn management processes into remotely executable computer procedures (micro-services)
  - Apply procedural workflow at the storage system to filter, subset, manipulate data
  - Minimize the amount of data pulled over the network
  - Automate administrative tasks
- Validate assessment criteria
  - Automate validation of collection properties
  - ISO MOIMS-rac

Generic Data Management Systems

iRODS - integrated Rule-Oriented Data System

DataManagement Environment	**Conserved Properties**	**Control Mechanisms**	**Remote Operations**
Management Functions	**Assessment Criteria**	**Management Policies**	**Management Procedures**
	Data grid Š Management virtualization		
DataManagement Infrastructure	**State Information**	**Rules**	**Micro-services**
	Data grid Š Data and trust virtualization		
Physical Infrastructure	**Database**	**Rule Engine**	**Storage System**

THE UNIVERSITY of NORTH CAROLINA at CHAPEL HILL

Federation of Data Grids

- Federation policies govern interactions between data grids
  - Remote data grid forwards request to the federated data grid
  - Local policies always enforced
  - Multiple types of federation
    - Master-slave data grids
    - Central archive data grids
    - Chained data grids
    - Peer-to-peer data grids

National Archives and Records Administration Transcontinental Persistent Archive Prototype

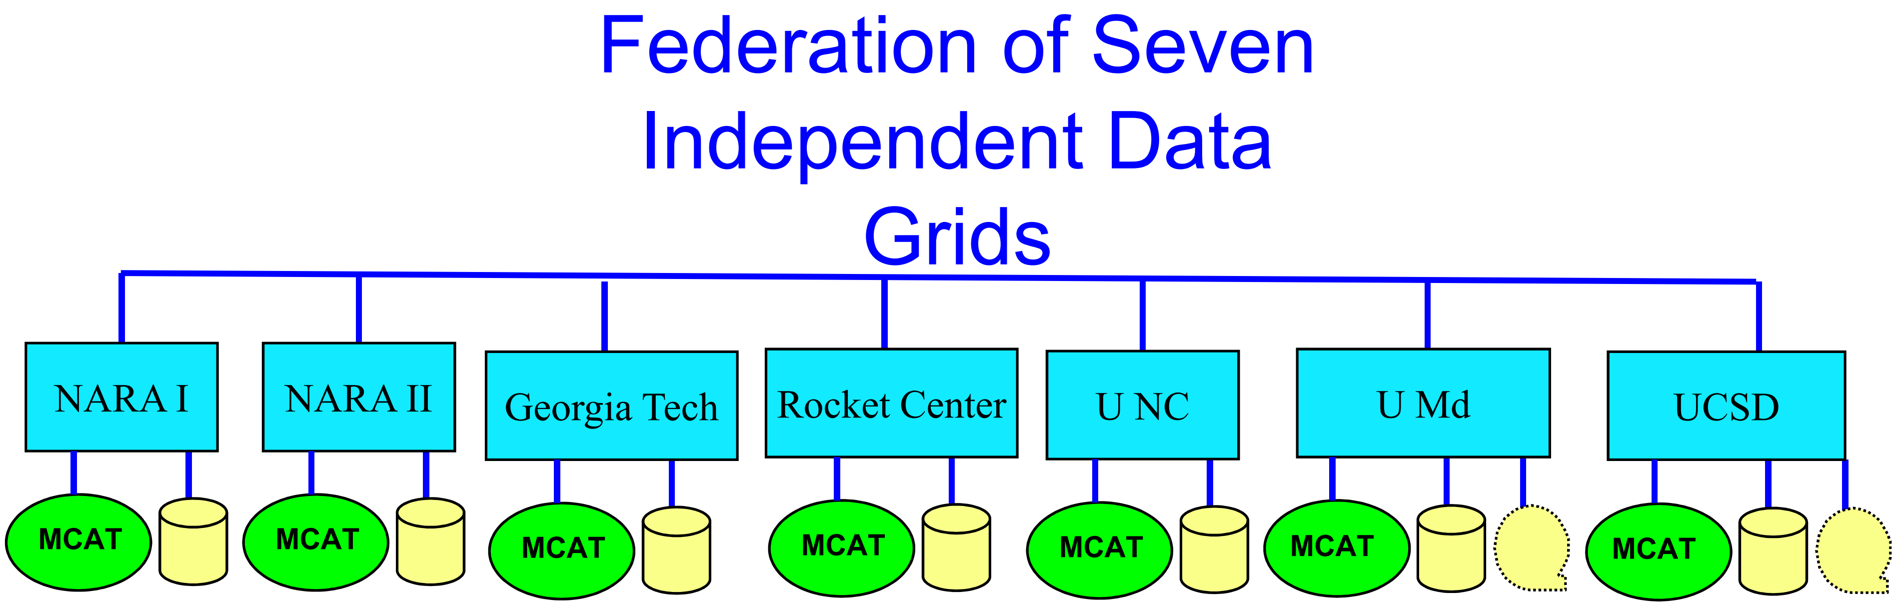

Extensible Environment, can federate with additional research and education sites. Each data grid uses different vendor products.

Policies

- Administrative
  - Retention, disposition, distribution, replication, deletion, registration, synchronization, integrity checks, IRB approval flags, addition of users, addition of resources
- Ingestion / Access
  - Metadata extraction, logical organization, derived data product generation, redaction, time-dependent access controls
- Validation
  - Authenticity verification, chain of custody, repository trustworthiness, audit trails

iRODS is a "coordinated NSF/OCI-Nat'l Archives research activity" under the auspices of the President's NITRD Program and is identified as among the priorities underlying the President's 2009 Budget Supplement in the area of Human and Computer Interaction Information Management technology research

Reagan W. Moore

rwmoore@renci.org

Current iRODS release 2.0.1

http://irods.diceresearch.org

Additional Slides

Building a Shared Collection

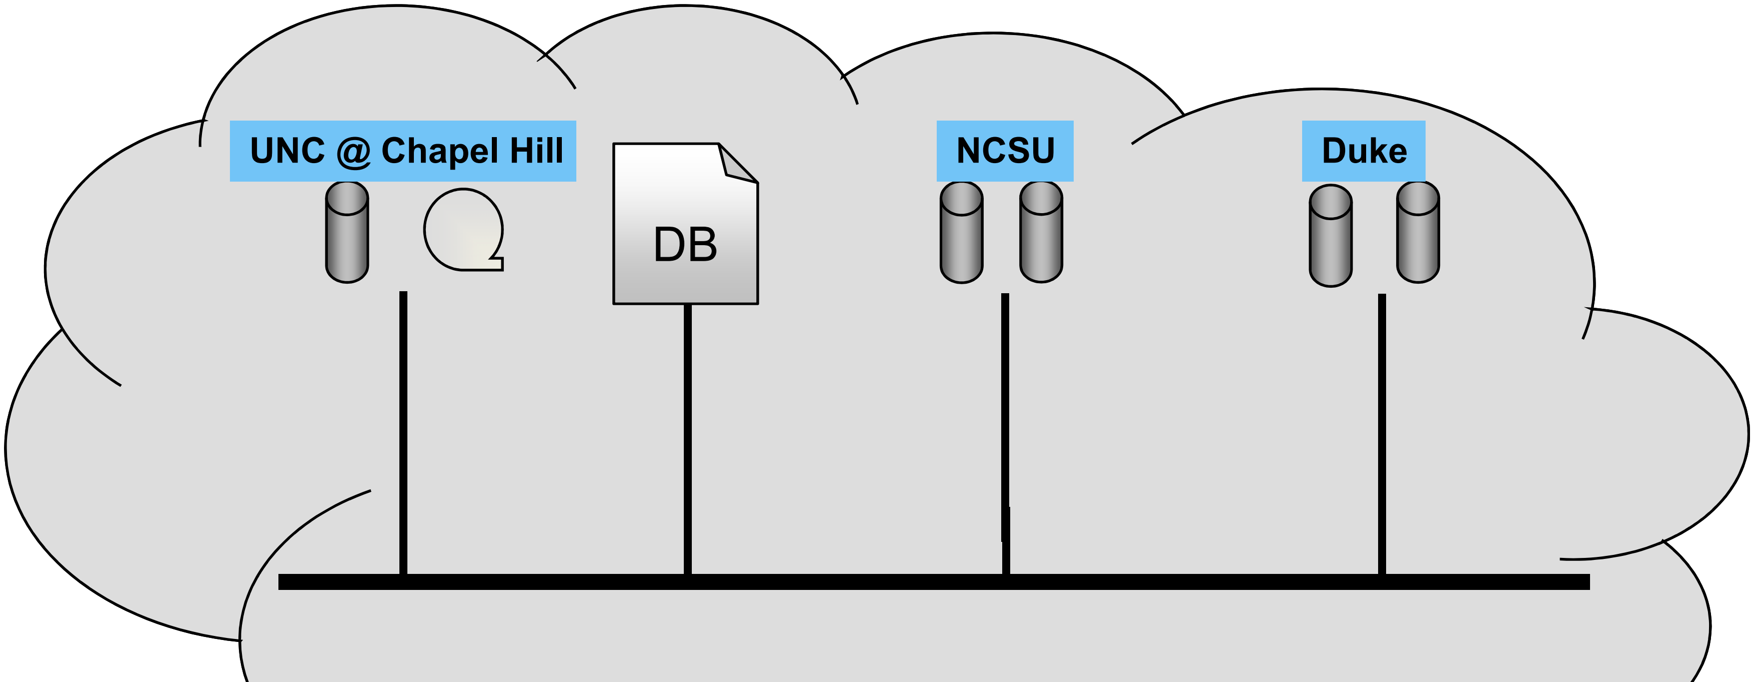

Have collaborators at multiple sites, each with different administration policies, different types of storage systems, different naming conventions.

Assemble a self-consistent, persistent distributed shared collection to support a specific purpose.

iRODS Shows Unified "Virtual Collection"

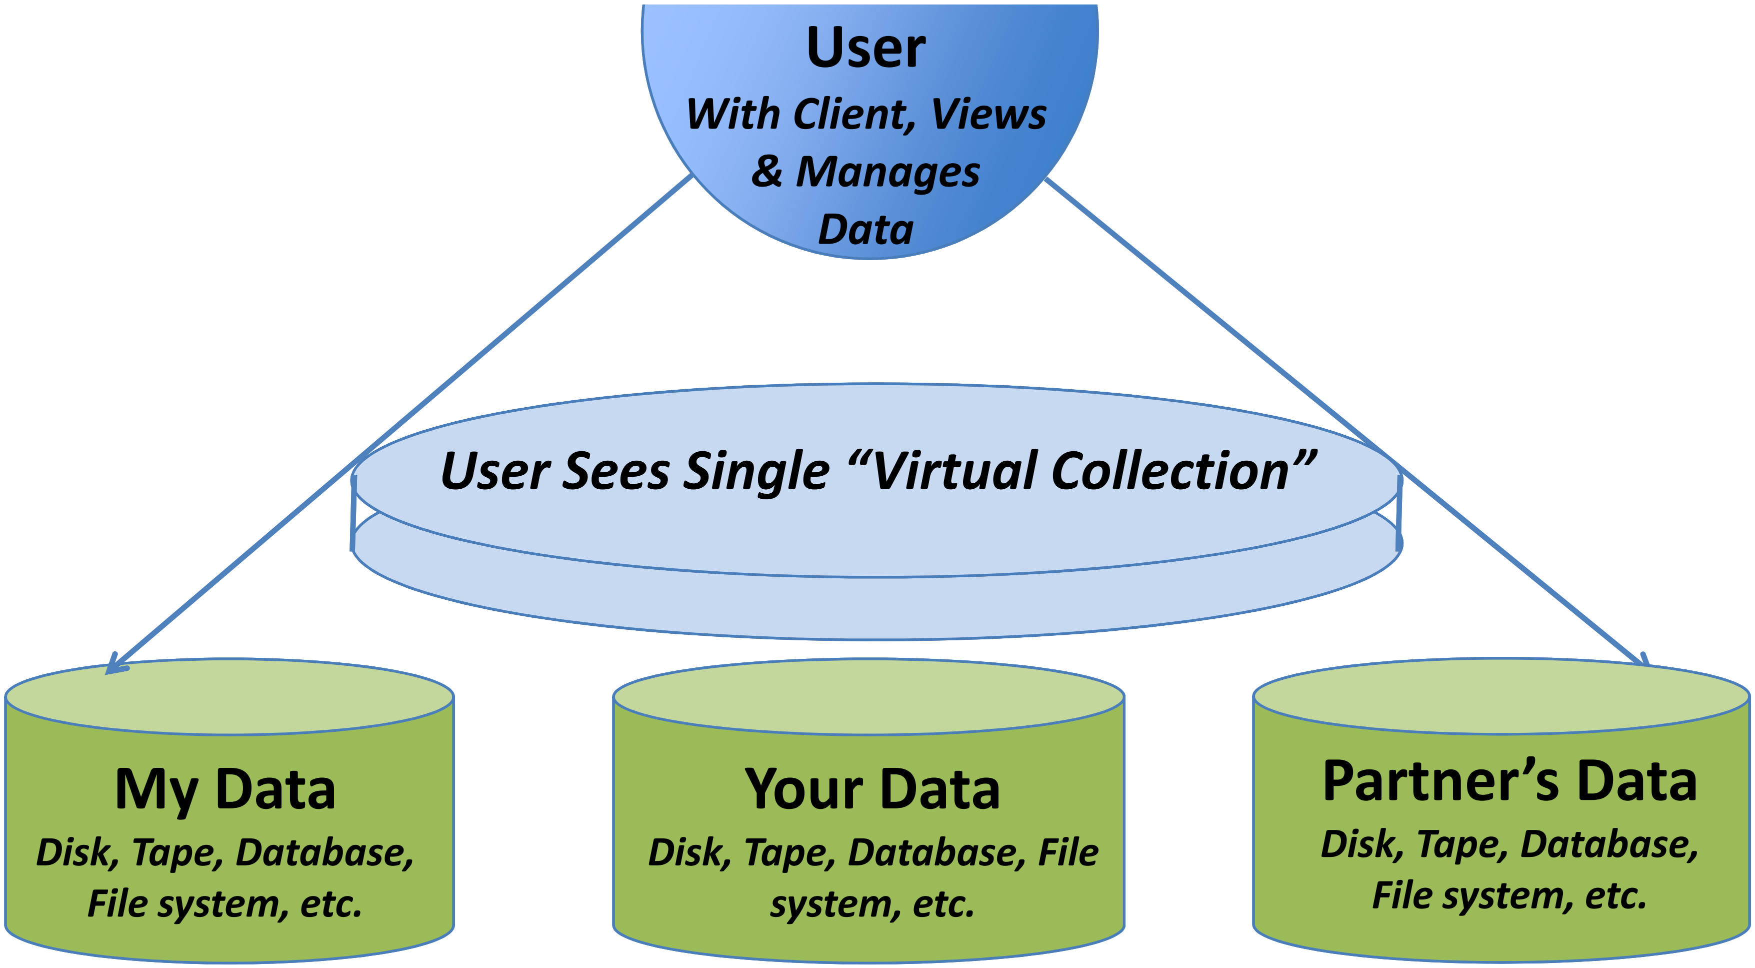

The iRODS Data Grid installs in a "layer" over existing or new data, letting you view, manage, and share part or all of diverse data in a unified Collection.

iRODS - Integrated Rule Oriented Data System

1. Shared collection assembled from data distributed across remote storage locations
2. Server-side workflow environment in which procedures are executed at remote storage locations
3. Policy enforcement engine, with computer actionable rules applied at the remote storage locations
4. Validation environment for assessment criteria
5. Consensus building system for establishing a collaboration (policies, data formats, semantics, shared collection)

Using a Data Grid - *Details*

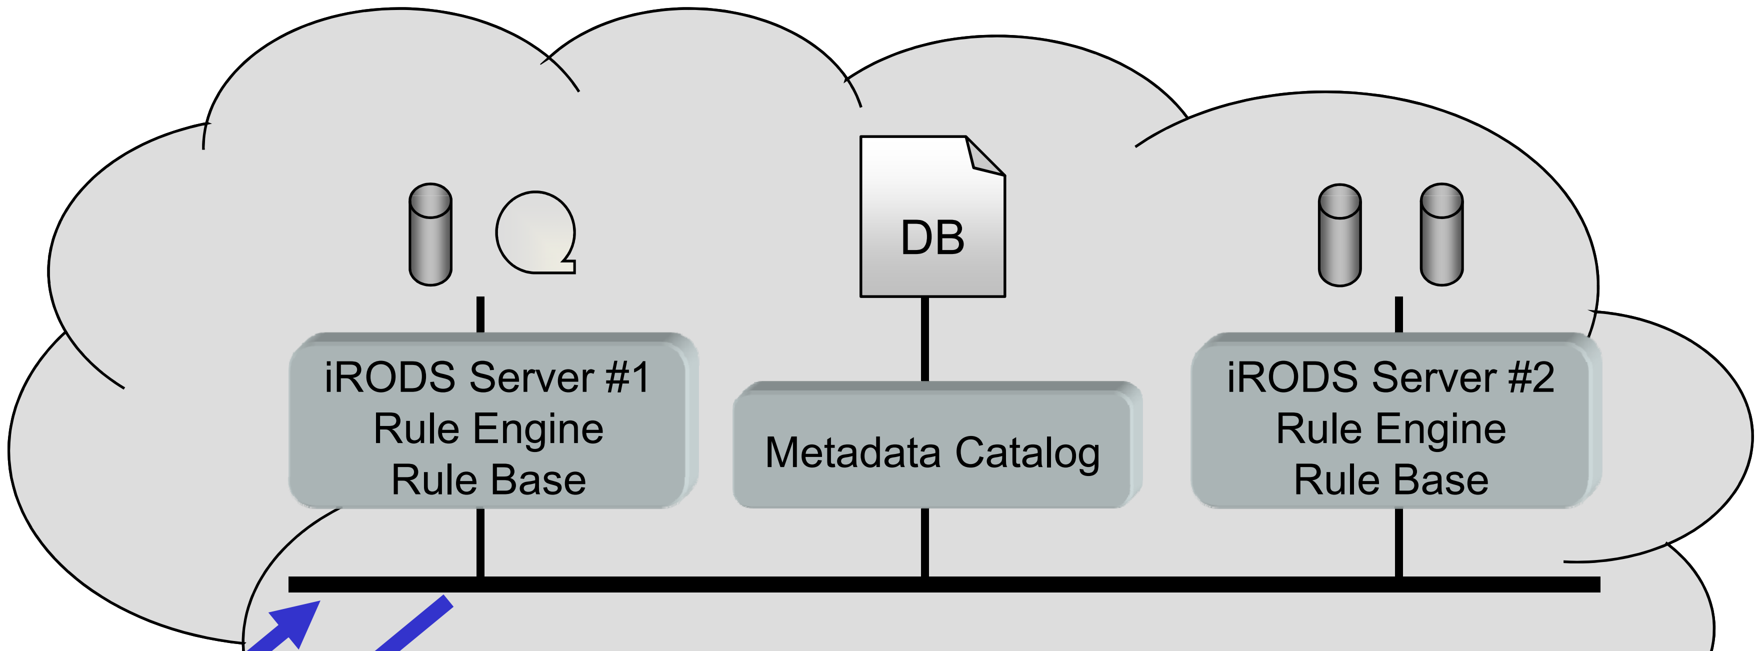

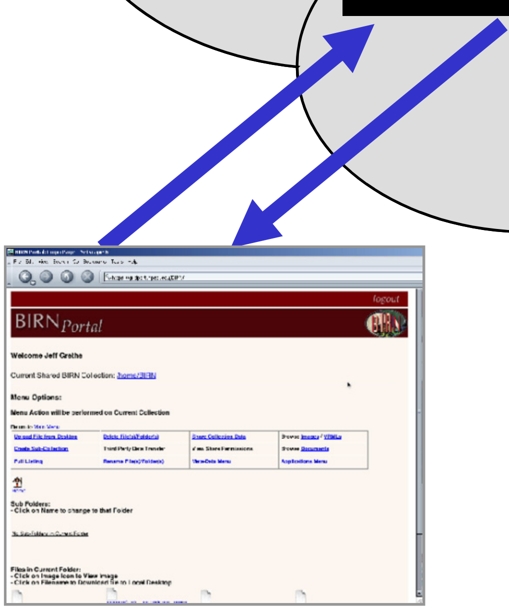

- User asks for data
- Data request goes to iRODS Server #1
- Server looks up information in catalog
- Catalog tells which iRODS server has data
- 1st server asks 2nd for data
- The 2nd iRODS server applies rules

Architecture

- Highly extensible, modular architecture
  - Peer-to-peer servers interact to form a data grid
- Layered architecture
  - Clients
  - Rules
  - Micro-services
  - Storage drivers
  - Structured information resource drivers

Data Virtualization

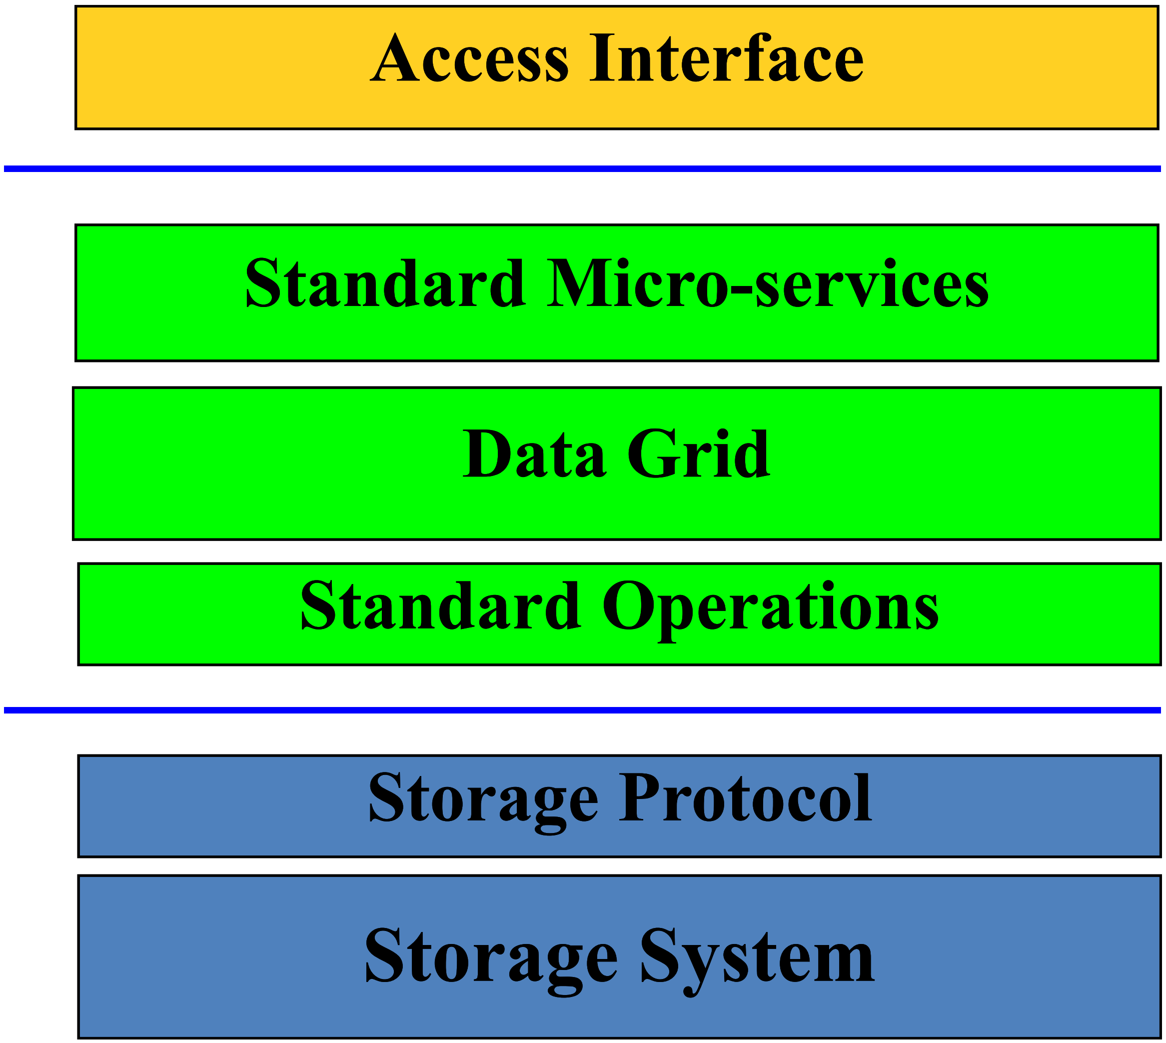

Map from the actions requested by the access method to a standard set of micro-services. The standard micro-services are mapped to the operations supported by the storage system

User Interfaces

- C library calls - Application level
- Unix shell commands - Scripting languages
- Java I/O class library (JARGON) - Web services
- SAGA - Grid API
- Web browser (Java-python) - Web interface
- Windows browser - Windows interface
- WebDAV - iPhone interface
- Fedora digital library middleware - Digital library middleware
- Dspace digital library - Digital library services
- Parrot - Unification interface
- Kepler workflow - Grid workflow
- Fuse user-level file system - Unix file system

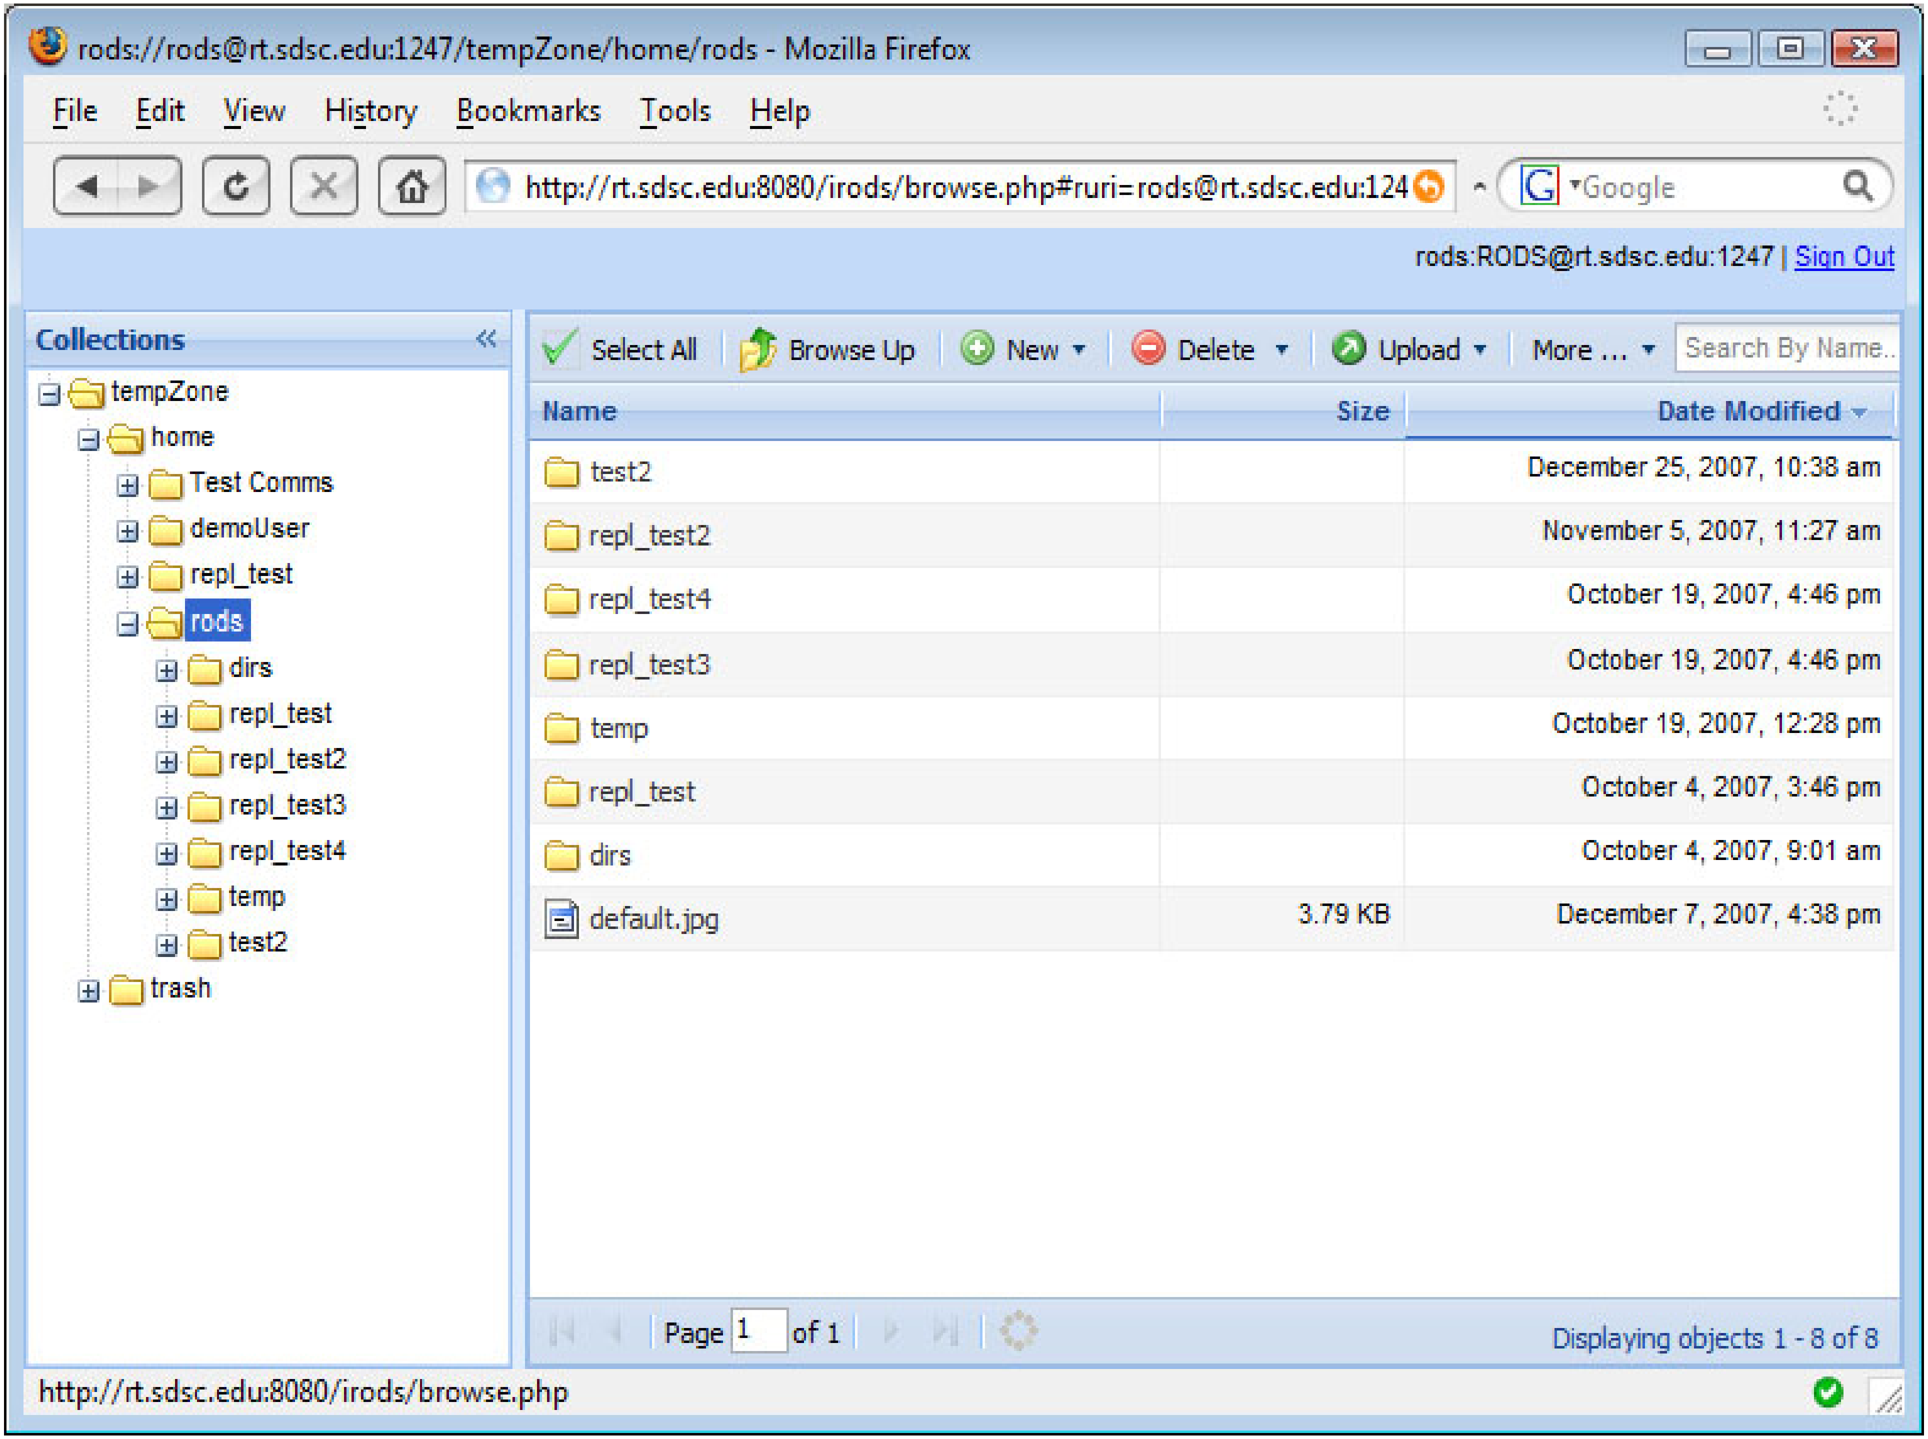

rods://rods@rt.sdsc.edu:1247/tempZone/home/rods - Mozilla Firefox
File Edit View History Bookmarks Tools Help
http://rt.sdsc.edu:8080/irods/browse.php#ruri=rods@rt.sdsc.edu:124
Google
rods:RODS@rt.sdsc.edu:1247 | Sign Out
Collections
tempZone
home
Test Comms
demoUser
repl_test
rods
dirs
repl_test
repl_test2
repl_test3
repl_test4
temp
test2
trash
Select All
Browse Up
New
Delete
Upload
More ...
Search By Name...
Name
Size
Date Modified
test2 December 25, 2007, 10:38 am
repl_test2 November 5, 2007, 11:27 am
repl_test4 October 19, 2007, 4:46 pm
repl_test3 October 19, 2007, 4:46 pm
temp October 19, 2007, 12:28 pm
repl_test October 4, 2007, 3:46 pm
dirs October 4, 2007, 9:01 am
default.jpg 3.79 KB December 7, 2007, 4:38 pm
Page 1 of 1
Displaying objects 1 - 8 of 8
http://rt.sdsc.edu:8080/irods/browse.php

iRODS Rules

- Server-side workflows

    Action | condition | workflow chain | recovery chain

- Condition - test on any attribute:
    - Collection, file name, storage system, file type, user group, elapsed time, IRB approval flag, descriptive metadata
- Workflow chain:
    - Micro-services / rules that are executed at the storage system
- Recovery chain:
    - Micro-services / rules that are used to recover from errors

ISO MOIMS-repository assessment criteria

- Are developing 150 rules that implement the ISO assessment criteria

90	***Verify descriptive metadata and source against SIP template and set SIP compliance flag***
91	***Verify descriptive metadata against semantic term list***
92	***Verify status of metadata catalog backup (create a snapshot of metadata catalog)***
93	***Verify consistency of preservation metadata after hardware change or error***

integrated Rule-Oriented Data System

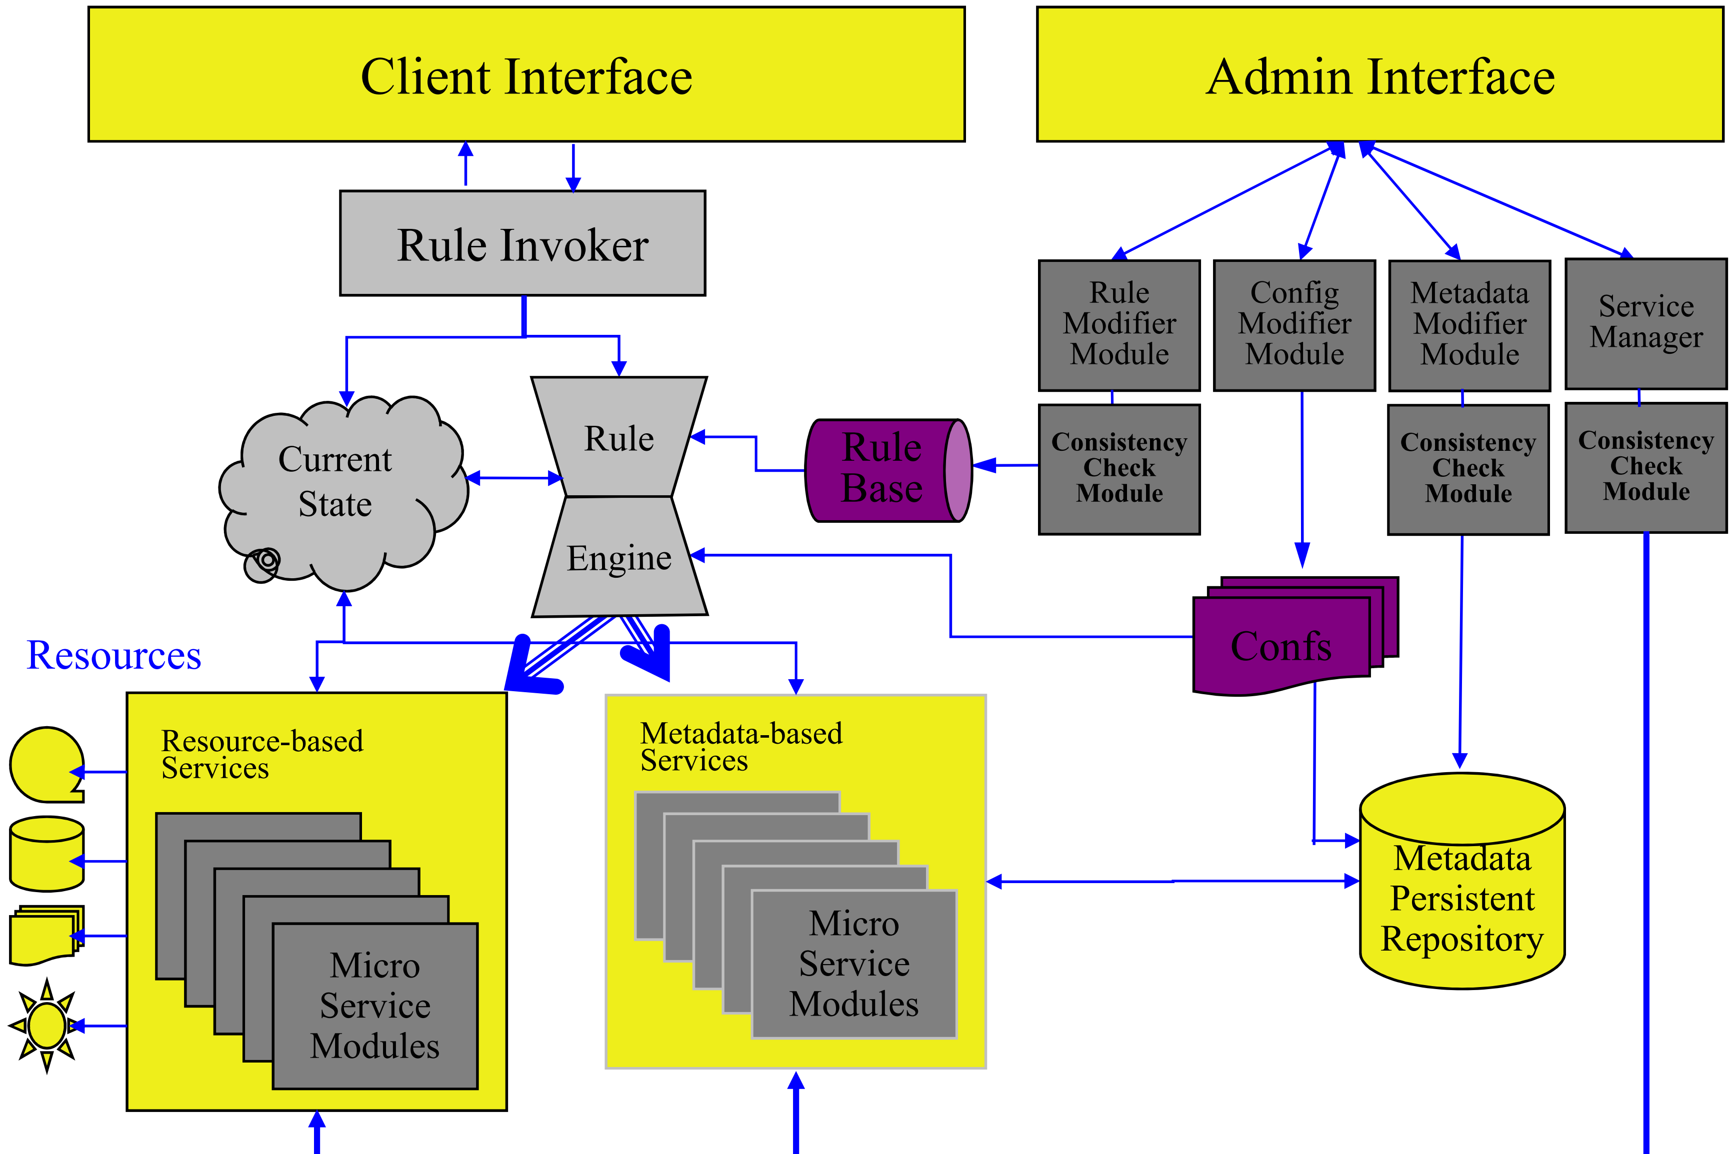

D·I·C·E i·R·O·D·S THE UNIVERSITY of NORTH CAROLINA at CHAPEL HILL UCSD NSF THE NATIONAL ARCHIVES ARCHIVES.GOV

Generic Infrastructure

- Data grids - sharing data
- Digital libraries - publishing data
- Persistent archives - preserving data
- Processing pipelines - analyzing data
- Real-time data management - federation of streams
- Integrated workflows - server and client side

- Switch applications by switching management policies
  - Build reference policy sets for each type of application

Scale

- Tens of millions to hundreds of millions of files
- Hundreds of terabytes to petabytes of data
- Hundreds of metadata attributes
- Hundreds of collaborators
- Tens to hundreds of policies
- Distributed internationally
- Federations of tens of data grids
- Thousands to tens of thousands of users

	As of 12/11//2006			As of 2/25/2008		
	Data_size (in GB)	*Count (files)*	*Curators*	*Data_size (in GB)*	*Count (files)*	*Curators*
Data Grid						
NSF/NVO	110,615.00	16,381,466	100	88,216.00	14,550,030	100
NSF/NPACI	35,909.00	7,458,960	380	43,684.00	7,643,389	380
PZONE	24,755.00	14,208,012	68	29,851.00	19,506,972	68
NSF/LDAS-SALK	163,706.00	176,897	67	211,542.00	173,806	67
NSF/SLAC-JCSG	18,494.00	1,945,302	55	26,100.00	2,675,426	55
NSF/TeraGrid	269,332.00	7,300,999	3,267	286,390.00	7,289,445	3,267
NCAR	2.00	8	2	76,255.00	435,597	2
LCA	1,834.00	39,611	2	4,544.00	78,289	2
NIH/BIRN	18,921.00	18,499,588	385	20,400.00	40,747,060	445
Others	8,013.00	161	227	8,013.00	161	227
Digital Library						
NSF/LTER	257.00	41,152	36	260.00	42,080	36
NSF/Portal	2,620.00	53,048	460	2,620.00	53,048	460
NIH/AFCS	733.00	94,686	21	733.00	94,686	21
NSF/SIO Explorer	2,681.00	1,201,719	27	3,053.00	1,220,303	27
NSF/SCEC	168,931.00	3,545,070	73	168,933.00	3,545,122	73
LLNL	8,176.00	335,540	5	18,934.00	2,338,384	5
CHRON	932.00	830,354	5	13,278.00	6,496,025	5
Persistent Archive						
NARA	4,713.00	5,992,817	58	5,036.00	6,409,726	58
NSF/NSDL	5,699.00	50,446,490	136	8,618.00	85,004,112	136
UCSD Libraries	5,080.00	1,077,202	29	5,210.00	1,720,463	29
NHPRC/PAT	3,756.00	527,695	28	2,575.00	1,050,795	28
RoadNet	2,057.00	712,534	30	3,886.00	1,792,185	30
UCTV	7,111.00	2,045	5	7,140.00	2,081	5
LOC	9,921.00	252,046	8	6,644.00	192,517	8
EarthSci	3,306.00	499,137	5	6,317.00	661,894	5
Total	877 TB	131 million	5479	1.04 PB	203 million	5539

Applications

- Institutional repositories
  - Carolina Digital Repository at University of North Carolina
  - Duke Medical Archive
- Regional data grids
  - RENCI data grid linking 7 engagement centers in North Carolina
  - HASTAC data grid linking humanities collections across 9 UC campuses
- National data grids
  - NARA Transcontinental Persistent Archive Prototype
  - NSF Temporal Dynamics of Learning Center data grid
  - NSF Ocean Observatories Initiative data grid
  - NASA Center for Computational Sciences archive
  - JPL Planetary Data System data grid
- International data grids
  - Australian Research Collaboration Service - ARCS
  - French National Library

Challenges - Social Consensus

- Building a consensus on management policies for the shared collection
- Translating service level agreements for shared use of resources into computer actionable rules
- Translating assessment criteria into computer executable procedures
- Defining federation policies for sharing data between data grids / institutions

D·I·C·E i·R·O·D·S THE UNIVERSITY of NORTH CAROLINA at CHAPEL HILL UCSD NSF THE NATIONAL ARCHIVES ARCHIVES.GOV SRB

Development Team

- DICE team
  - Arcot Rajasekar - iRODS development lead
  - Mike Wan - iRODS chief architect
  - Wayne Schroeder - iRODS developer
  - Bing Zhu - Fedora, Windows
  - Lucas Gilbert - Java (Jargon), DSpace
  - Paul Tooby - documentation, foundation
  - Sheau-Yen Chen - data grid administration
- Preservation
  - Richard Marciano - Preservation development lead
  - Chien-Yi Hou - preservation micro-services
  - Antoine de Torcy - preservation micro-services

Foundation

- Data Intensive Cyber-environments
  - Non-profit open source software development
  - Promote use of iRODS technology
  - Support standards efforts
  - Coordinate international development efforts
    - IN2P3 - quota and monitoring system
    - King's College London - Shibboleth
    - Australian Research Collaboration Services - WebDAV
    - Academia Sinica - SRM interface

Prioritize Development

- Generic infrastructure
  - Turn specific requests into generic framework
- Assign importance
  - Bug fixes
  - Funded development
  - Multiple requests
  - Critical need to meet major demonstration
- Incorporate community supplied mods
  - Generic infrastructure
  - Compliance with iRODS modular design

Features in Next Release

- Support for mySQL as the iCAT metadata catalog
- Support for Kerberos authentication
- Support for resource monitoring system
- Multi-tasking the batch server (irodsReServer) for more robust job execution.
- A new resource class - Compound Resource for a class of resources that support only put/get type functions
  - (e.g., ftp, HPSS parallel I/O, etc)
- Better support for writing micro-services - consolidation of data structures used by micro-services, more helper routines.
- Better Java interface for iRODS - parallel I/O, metadata support, etc.
- Multi-threading put/get of small files (if it can be done in time for the release)
- Better support for restricted listing of collections (ACLs).