

# Policy Based Data Management

## iRODS

Reagan W. Moore (DICE-UNC)

Arcot Rajasekar (DICE-UNC)

Mike Wan (DICE-UCSD)

Wayne Schroeder (DICE-UCSD)

Mike Conway (DICE-UNC)

Antoine de Torcy (DICE-UNC)

Hao Xu (UNC)

Sheau-Yen Chen (DICE-UCSD)

Jason Coposky (RENCI)

Lisa Stillwell (RENCI)

Leesa Brieger (RENCI)

<http://irods.diceresearch.org>



THE UNIVERSITY  
of NORTH CAROLINA  
at CHAPEL HILL



rencl

# Agenda

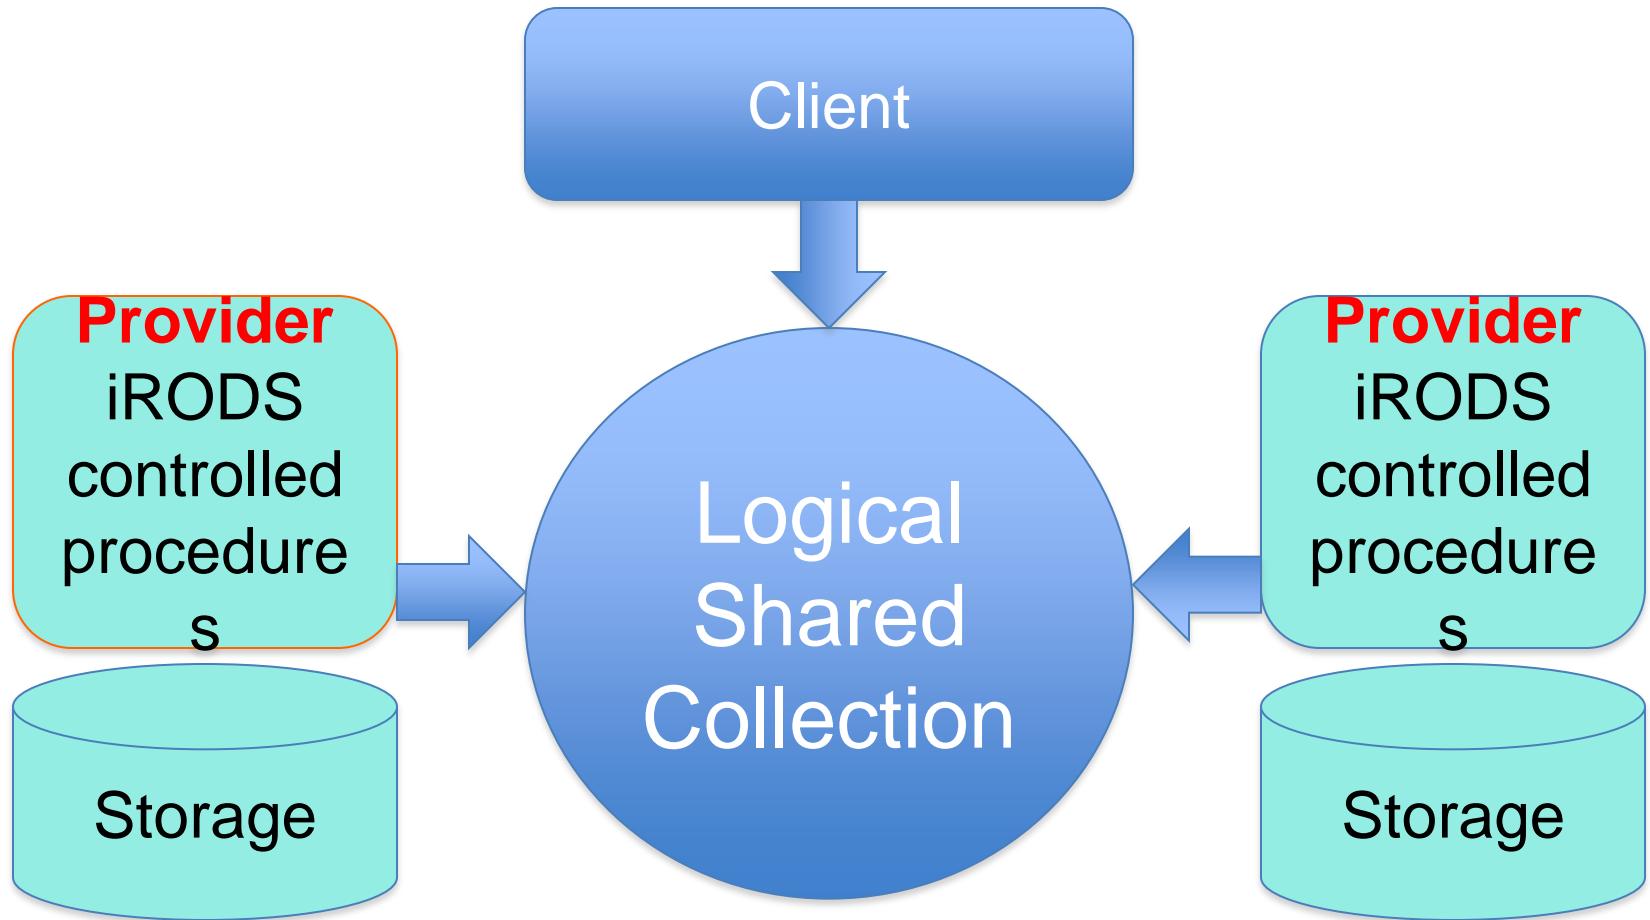
- 09:00 Introduction to iRODS technology – (Moore, UNC-CH)
- 10:00 Open discussion on data management needs – (Moore, UNC-CH)
- 10:30 Break
- 11:00 Installation of iRODS clients (Windows, Mac, Linux) – (Moore UNC-CH)
- 11:45 Initial demonstration of iRODS client access – (Moore, UNC-CH)
- 12:00 Lunch
- 14:00 Introduction to iRODS rules and micro-services – (Moore, UNC-CH)
- 14:45 Demonstration of federation of data grids - (Yutaka Kawai, KEK, Japan, Moore UNC-CH)
- 15:00 Simple rule examples including database queries (Moore, UNC-CH)
- 15:30 Break
- 16:00 Complex rule examples including scheduling (Moore, UNC-CH)
- 16:30 SRM development status– (Wei-Long UENG, ASGC)
- 17:00 Open discussion of iRODS applications in digital libraries, data grids, and preservation environments – (Moore, UNC-CH)
- 17:30 End of tutorial



THE UNIVERSITY  
of NORTH CAROLINA  
at CHAPEL HILL



# Policy-based Data Sharing



Consensus on Policies and Procedures  
controls the shared data



THE UNIVERSITY  
of NORTH CAROLINA  
at CHAPEL HILL



# Applications

- International projects
  - Cyber Square Kilometer Array (radio astronomy), Cinegrid (movies)
- National data grids
  - Australia, New Zealand, Portugal, UK, France
- Federal agency archives
  - NASA Center for Climate Simulation, National Optical Astronomy Observatories
- Institutional repositories
  - French National Library, Carolina digital repository, Broad Institute genomics data, Sanger Institute



THE UNIVERSITY  
of NORTH CAROLINA  
at CHAPEL HILL



# iRODS Version 3.0

- Released Sept. 30, 2011
- Next release scheduled for March, 2012
- New features
  - New rule engine
    - Strong parameter typing
    - Optimized performance (thousands of rules)
    - Expanded rule programming language
  - Rule versioning
  - Distributed rule base management
  - Soft links to external resources
  - New transport management for large files
  - Improved Java interface
  - Windows native port (C++ compilation)
  - Dropbox interface (iDrop)



THE UNIVERSITY  
of NORTH CAROLINA  
at CHAPEL HILL



# Under Development

- Tickets (3.1)
  - Ability to give access to specified file for specified time duration
- netCDF library support (3.1)
- Pluggable Authentication Modules (3.2)
  - Use external identity resource
- Registration of workflows (3.2)
  - Micro-service structured objects
  - Re-execute workflows for derived data products
  - Policy-encoded objects



THE UNIVERSITY  
of NORTH CAROLINA  
at CHAPEL HILL



# Policy-based Data Environments

- *Purpose* - reason a collection is assembled
- *Properties* - attributes needed to ensure the **purpose**
- *Policies* - controls for enforcing desired **properties**,  
**mapped to computer actionable rules**
- *Procedures* - functions that implement the **policies**  
**mapped to computer actionable workflows**
- *Persistent state information* - results of applying the **procedures**  
**mapped to system metadata**
- *Assessment criteria* - validation that **state information** conforms to the desired **purpose**  
**mapped to periodically executed policies**

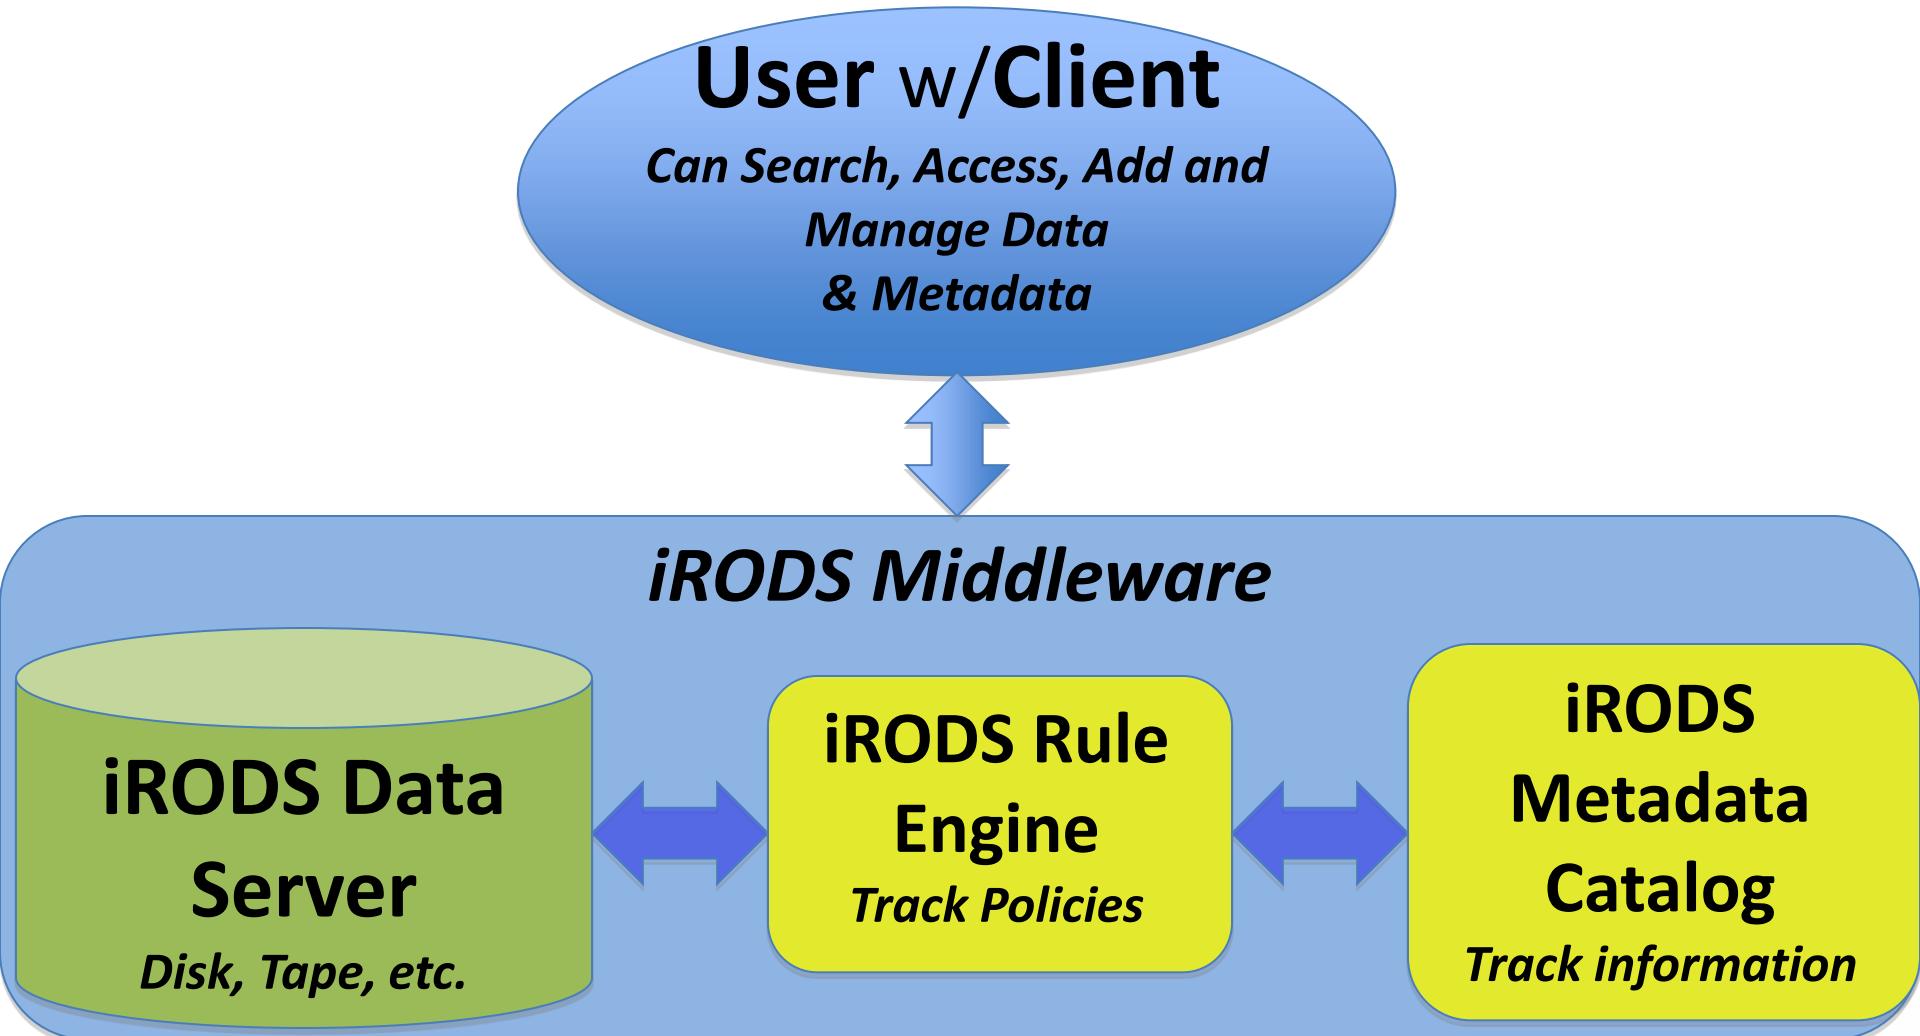


THE UNIVERSITY  
of NORTH CAROLINA  
at CHAPEL HILL



rencl

# Overview of iRODS Architecture



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

# Data Management Applications

## Generic Infrastructure

- Data grids – PB-size distributed collections
  - Astronomy – NOAO, CyberSKA, LSST
  - High Energy Physics – BaBar, KEK
  - Earth Systems – NASA (MODIS data set)
  - Australian Research Collaboration Service
  - Genomics – UNC-CH/RENCI
- Institutional repositories
  - Carolina Digital Repository
- Libraries
  - Texas Digital Libraries
  - Seismology - Southern California Earthquake Center
- Archives
  - Ocean Observatories Initiative



THE UNIVERSITY  
of NORTH CAROLINA  
at CHAPEL HILL



# Data Virtualization

Data Grid

**Access Interface**

Map from the actions requested by the client to multiple policy enforcement points.

**Policy Enforcement Points**

Map from policy to standard micro-services.

**Standard Micro-services**

Map from micro-services to standard Posix I/O operations.

**Standard I/O Operations**

**Storage Protocol**

Map standard I/O operations to the protocol supported by the storage system

**Storage System**



THE UNIVERSITY  
of NORTH CAROLINA  
at CHAPEL HILL



# Data Grid Clients (48)

API	Client	Developer	Language
<b>Browser</b>			
	DCAPE	UNC	
	iExplore	RENCI-Oleg	C++
	JUX	IN2P3	Jargon
	Peta Web browser	PetaShare	
	iDrop web browser	Mike Conway	Java
	Davis web interface	ARCS	
<b>Digital Library</b>	Rich web client	Lisa Stillwell - RENCI	
	Akubra/iRODS	DICE	Jargon
	Dspace	MIT	
	Fedora on Fuse	IN2P3	FUSE
	Fedora/iRODS module	DICE	Jargon
	Islandora	DICE	Jargon
	Curators Workbench	CDR-UNC-CH	Jargon
<b>File System</b>			
	Davis - Webdav	ARCS	Jargon
	Dropbox / iDrop	DICE-Mike Conway	Jargon
	FUSE	IN2P3, DICE,	FUSE
	FUSE optimization	PetaShare	FUSE
	OpenDAP	ARCS	
	PetaFS (Fuse)	Petashare - LSU	
	Petashell (Parrot)	PetaShare	



THE UNIVERSITY  
of NORTH CAROLINA  
at CHAPEL HILL



rencl

# iRODS Clients (Cont.)

<b>Grid</b>	GridFTP - Griffin	ARCS	
	Jsaga	IN2P3	Jargon
	Parrot	UND - Doug Thain	
	SRM	Academia Sinica	
	Saga	KEK	
<b>I/O Libraries</b>			
	PRODS - PHP	Renci - Lisa Stillwell	
	C API	DICE-Mike Wan	C
	C I/O library	DICE-Wayne Schroeder	C
	Fortran	Schroeder	C
	Eclipse file system	CDR - UNC-CH	Jargon
	Jargon	DICE-Mike Conway	Jargon
		SHAMAN-Jerome Fusillier	Python
<b>Portal</b>	Pyrods - Python		
	EnginFrame	NICE / RENCI	Jargon
	Petashare Portal	LSU	Jargon



THE UNIVERSITY  
of NORTH CAROLINA  
at CHAPEL HILL



# iRODS Clients (Cont.)

Tools			
Archive tools-NOAO	NOAO		
Big Board visualization	RENCI		
iFile	GA Tech		
i-commands	DICE		
Pcommands	PetaShare		
Resource Monitoring	IN2P3		
Sync-package	Academica Sinica		
URSpace	Teldap - Academica Sinica		
Web Service			
VOSpace	IVOA		
Shibboleth	King's College		
Workflows			
Kepler - actor	DICE	Jargon	
Stork - interoperability	LSU		
Workflow Virtualization	LSU		
Taverna - actor	RENCI		



THE UNIVERSITY  
of NORTH CAROLINA  
at CHAPEL HILL



rencl

# Policy Enforcement Points

- Currently have 74 locations within iRODS framework where policies are checked.
  - Each action may involve multiple policy enforcements points
- Policy enforcement points
  - Pre-action policy (selection of storage location)
  - Policy execution (file deletion control)
  - Post-action policy (derived data products)



THE UNIVERSITY  
of NORTH CAROLINA  
at CHAPEL HILL



rencl

icommands		none													
icp		x	x	x	x	x	x	x	x	x	x	x	x	x	x
icp -N 2		x	x	x	x	x	x	x	x	x	x	x	x	x	x
iphybun		x	x	x	x	x	x	x	x	x	x	x	x	x	x
irepl		x	x	x	x	x	x			x		x		x	
ibun -cD		x	x	x	x	x	x	x	x			x	x		
iput		x	x	x	x	x	x	x	x			x	x		
iphymv		x	x	x	x	x	x	x	x		x			x	
imv		x	x	x			x	x	x		x			x	x
irm		x	x	x			x	x	x		x			x	x
irm -r collection		x	x	x			x	x	x		x			x	x



THE UNIVERSITY  
of NORTH CAROLINA  
at CHAPEL HILL



rendi

# Policy Enforcement Points (74)

## ACTION

acCreateUser  
acDeleteUser  
acGetUserbyDN  
acTrashPolicy  
acAclPolicy  
acSetCreateConditions  
acDataDeletePolicy  
acRenameLocalZone  
acSetRescSchemeForCreate  
acRescQuotaPolicy  
acSetMultiReplPerResc  
acSetNumThreads  
acVacuum  
acSetResourceList  
acSetCopyNumber  
acVerifyChecksum  
acCreateUserZoneCollections  
acDeleteUserZoneCollections  
acPurgeFiles  
acRegisterData  
acGetIcatResults  
acSetPublicUserPolicy  
acCreateDefaultCollections  
acDeleteDefaultCollections

## PRE-ACTION POLICY

acPreProcForCreateUser  
acPreProcForDeleteUser  
acPreProcForModifyUser  
acPreProcForModifyUserGroup  
acChkHostAccessControl  
acPreProcForCollCreate  
acPreProcForRmColl  
acPreProcForModifyAVUMetadata  
acPreProcForModifyCollMeta  
acPreProcForModifyDataObjMeta  
acPreProcForModifyAccessControl  
acPreprocForDataObjOpen  
acPreProcForObjRename  
acPreProcForCreateResource  
acPreProcForDeleteResource  
acPreProcForModifyResource  
acPreProcForModifyResourceGroup  
acPreProcForCreateToken  
acPreProcForDeleteToken  
acNoChkFilePathPerm  
acPreProcForGenQuery  
acSetReServerNumProc  
acSetVaultPathPolicy

## POST-ACTION POLICY

acPostProcForCreateUser  
acPostProcForDeleteUser  
acPostProcForModifyUser  
acPostProcForModifyUserGroup  
acPostProcForDelete  
acPostProcForCollCreate  
acPostProcForRmColl  
acPostProcForModifyAVUMetadata  
acPostProcForModifyCollMeta  
acPostProcForModifyDataObjMeta  
acPostProcForModifyAccessControl  
acPostProcForOpen  
acPostProcForObjRename  
acPostProcForCreateResource  
acPostProcForDeleteResource  
acPostProcForModifyResource  
acPostProcForModifyResourceGroup  
acPostProcForCreateToken  
acPostProcForDeleteToken  
acPostProcForFilePathReg  
acPostProcForGenQuery  
acPostProcForPut  
acPostProcForCopy  
acPostProcForCreate



THE UNIVERSITY  
of NORTH CAROLINA  
at CHAPEL HILL



input .../src/irm.c

checks 10 policy hooks

**srbrick14:10900:ApplyRule#116:: acChkHostAccessControl**

srbrick14:10900:GotRule#117:: acChkHostAccessControl

**srbrick14:10900:ApplyRule#118:: acSetPublicUserPolicy**

srbrick14:10900:GotRule#119:: acSetPublicUserPolicy

**srbrick14:10900:ApplyRule#120:: acAclPolicy**

srbrick14:10900:GotRule#121:: acAclPolicy

**srbrick14:10900:ApplyRule#122:: acSetRescSchemeForCreate**

srbrick14:10900:GotRule#123:: acSetRescSchemeForCreate

srbrick14:10900:execMicroSrv#124:: msiSetDefaultResc(demoResc,null)

**srbrick14:10900:ApplyRule#125:: acRescQuotaPolicy**

srbrick14:10900:GotRule#126:: acRescQuotaPolicy

srbrick14:10900:execMicroSrv#127:: msiSetRescQuotaPolicy(off)

**srbrick14:10900:ApplyRule#128:: acSetVaultPathPolicy**

srbrick14:10900:GotRule#129:: acSetVaultPathPolicy

srbrick14:10900:execMicroSrv#130:: msiSetGraftPathScheme(no,1)

**srbrick14:10900:ApplyRule#131:: acPreProcForModifyDataObjMeta**

srbrick14:10900:GotRule#132:: acPreProcForModifyDataObjMeta

**srbrick14:10900:ApplyRule#133:: acPostProcForModifyDataObjMeta**

srbrick14:10900:GotRule#134:: acPostProcForModifyDataObjMeta

**srbrick14:10900:ApplyRule#135:: acPostProcForCreate**

srbrick14:10900:GotRule#136:: acPostProcForCreate

**srbrick14:10900:ApplyRule#137:: acPostProcForPut**

srbrick14:10900:GotRule#138:: acPostProcForPut

srbrick14:10900:GotRule#139:: acPostProcForPut

srbrick14:10900:GotRule#140:: acPostProcForPut

# Actionable Rule

- Name | Constraint | Workflow | Recovery
- Associate name with
  - Policy enforcement point
  - User defined rule
- Constraint is a test on any
  - Session variable
  - Persistent state attribute
- Workflow composed by chaining micro-services
- Recovery workflow composed from micro-services



THE UNIVERSITY  
of NORTH CAROLINA  
at CHAPEL HILL



# Hello World

```
myTestRule {  
# Rule to print out "Hello World"  
    writeLine("stdout", "Hello World");  
}  
INPUT null  
OUTPUT ruleExecOut
```



THE UNIVERSITY  
of NORTH CAROLINA  
at CHAPEL HILL



# Policies – actionable rules

- Retention, disposition, distribution, arrangement
- Authenticity, provenance, description
- Integrity, replication, synchronization
- Deletion, trash cans, versioning
- Archiving, staging, caching
- Authentication, authorization, redaction
- Access, approval, IRB, audit trails, report generation
- Assessment criteria, validation
- Derived data product generation, format parsing
- Federation of independent data grids



THE UNIVERSITY  
of NORTH CAROLINA  
at CHAPEL HILL



# Highly Controlled Environment

- All accesses are authenticated
  - GSI / Kerberos / Challenge-response / Shibboleth
- All operations are authorized
  - ACLs on files, storage
  - User groups, storage groups
- All policies evaluate a constraint
  - Constraints based on persistent state information and session information



THE UNIVERSITY  
of NORTH CAROLINA  
at CHAPEL HILL



# iRODS Extensible Infrastructure

- **Clients** – specific to discipline and life cycle state
- **Policies** – specific to discipline
- **Procedures** – specific to discipline
- Remaining infrastructure is generic
  - Network transport
  - Authentication / Authorization
  - Distributed storage access
  - Remote execution
  - Metadata management
  - Message passing / distributed debugging
  - Rule engine



THE UNIVERSITY  
of NORTH CAROLINA  
at CHAPEL HILL



# Extended Capabilities

- Replication
- Registration of files into the data grid
- Synchronization of remote directory
- Managed file transport (iDrop)
- Automated metadata extraction
- Queries on metadata, tags
- Server-side workflows (loop over result sets)
- Parallel I/O streams & RBUDP transport



THE UNIVERSITY  
of NORTH CAROLINA  
at CHAPEL HILL



# Cloud Integration

- iRODS drivers for
  - EC2
  - S3
- Applications include
  - <http://www.fiercegovernmentit.com/story/nasa-touts-nebula-faults-public-cloud/2011-08-22>
  - <http://nebula.nasa.gov/media/uploads/nasa-nebula-in-action.pdf>



THE UNIVERSITY  
of NORTH CAROLINA  
at CHAPEL HILL



# Integration of Data and Workflows

- Client-side workflows
  - Presentation / display
  - Client managed synchronization / transport
- Policy-driven workflows
  - Automation of administrative functions
  - Processing at the storage system
  - Data-intensive computations
- Traditional computer workflows
  - CPU-intensive computations



THE UNIVERSITY  
of NORTH CAROLINA  
at CHAPEL HILL



# Expectations

- Data collection sizes will increase
  - Now petabytes, soon exabytes
    - 1 PB/year = 32 MB/sec
    - 1 PB/day = 11.6 GB/sec
  - Indexing done in the storage system
- Integration of data manipulation with storage controllers - DDN SFA10KE
  - Analyses done within the storage system



THE UNIVERSITY  
of NORTH CAROLINA  
at CHAPEL HILL



rencl

# Future Applications

- Digital libraries
  - Continuous indexing of contents
- Scientific data collections
  - Extraction of features from data sets
  - Creation of derived data products
- Archives
  - Transformative migrations
  - Validation of assessment criteria



THE UNIVERSITY  
of NORTH CAROLINA  
at CHAPEL HILL



# iRODS Development

- Realized Objects
  - Micro-service Structured Object (MSSO)
  - Register into iRODS the workflow needed to create a derived data product and the input parameter / files
  - Accessing the link causes the derived data to be generated and registered as a replica
  - Implemented using compound resource concept
  - The realized object workflow can reference other realized objects, and dynamically re-create an object when any underlying input resource is altered.
  - Effectively manages provenance for workflows



THE UNIVERSITY  
of NORTH CAROLINA  
at CHAPEL HILL



# Automated Re-processing Design

Object A workflow      depends on

Object B workflow      depends on

Object C workflow      depends on

Object D

Change object D, and system automatically re-computes objects C, B, and A when object A is clicked



THE UNIVERSITY  
of NORTH CAROLINA  
at CHAPEL HILL



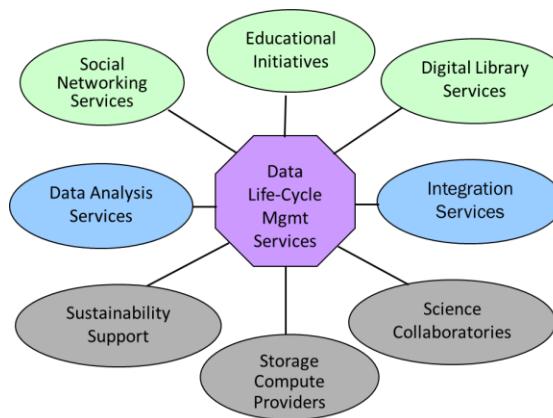
# DataNet Federation Consortium

## Data Driven Science

- Implement national data grid
  - Federate existing discipline-specific data management systems to enable national research collaborations
- Enable collaborative research on shared data collections
  - Manage collection life cycle as the user community broadens
- Integrate “live” research data into education initiatives
  - Enable student research participation through control policies

### Cyber-infrastructure Partners:

Univ. of North Carolina, Chapel Hill  
Univ. of California, San Diego  
Arizona State University  
Drexel University  
Duke University  
University of Arizona  
University of South Carolina



National Science Foundation Cooperative Agreement: OCI-0940841



THE UNIVERSITY  
of NORTH CAROLINA  
at CHAPEL HILL



Policy-based  
data management

renCI

Project

Shared Collection

Processing Pipeline

Digital Library

Reference Collection

Federation

**Collection Life Cycle**

### Science and Engineering Initiatives:

Ocean Observatories Initiative  
the iPlant Collaborative  
CUAHSI  
CIBER-U  
Odum Social Science Institute  
Temporal Dynamics of Learning Center

# Replication Validation Rule

- Loop over collection in sets of 256 files
- Restart capability for long-running session
  - Tracks identity of the files previously verified
  - Sets TEST\_DATA\_ID attribute on the collection
- Control execution rate
  - Slow down rate when executing too fast
- Audit trail of all changes to files
- Identify and replace corrupted files
- Create missing replicas
  - Distribute new replicas across storage vaults



THE UNIVERSITY  
of NORTH CAROLINA  
at CHAPEL HILL



rencl

# iRODS - Open Source Software

Reagan W. Moore

[rwmoores@renci.org](mailto:rwmoores@renci.org)

<http://irods.diceresearch.org>

***NSF OCI-0848296 “NARA Transcontinental Persistent Archives Prototype”***  
***NSF SDCI-0721400 “Data Grids for Community Driven Applications”***



THE UNIVERSITY  
of NORTH CAROLINA  
at CHAPEL HILL

