

Sustainability: The iRODS Consortium Economic Technological

Mary Whitton
Project Manager

Brand Fortner (by telecon)
iRODS Consortium Executive Director, RENCI



DFC October 2014 NSF Review Slide 1



Sustainability

- DFC produces open-source technology that facilitates federation of resources
 - Need to grow user and developer base
- DFC technology is built on and with iRODS technology
 - Technical dependence on iRODS for evolution
- The iRODS Consortium is dedicated to iRODS sustainability
 - Economic model for self-sufficiency
 - Engaged users to drive technical advances



DFC October 2014 NSF Review Slide 2



Updates from the Consortium

- iRODS 4.0 and Beyond
- An update on iRODS



DFC October 2014 NSF Review Slide 3



iRODS 4.0 and Beyond

Dan Bedard

iRODS Consortium, RENCI



DFC April 2013 NSF Review Slide 4



Motivation and Goals

iRODS is open source middleware for:

- Data Discovery,
- Workflow Automation,
- Secure Collaboration, and
- Data Virtualization

The iRODS Consortium wants iRODS to be *sustainable*.

That means it has to be:

- Reliable
- Supportable
- Usable
- Scalable



DFC October 2014 NSF Review Slide 5



Technical Framework for Sustainable iRODS

Components of a sustainable iRODS:

- Pluggable Architecture
- Configuration Management
- Grid Introspection
- Plugin Dependency Model

A feature set that remains **relevant**.



DFC October 2014 NSF Review Slide 6




iRODS 4.0: A First Step

Pluggable Architecture

- Simplifies development, testing, and support
- Broadens the developer community

Makes the long-term relevance of iRODS more likely.



DFC October 2014 NSF Review Slide 7



iRODS 4.0 Plugins, Presently

Resources:

- Compound
- Deferred
- Load Balanced
- MSO
- Non-Blocking
- Passthru
- Random
- Replication
- Round Robin
- Universal MSS
- Unix File System
- WOS
- HPSS
- S3

Authentication:

- Native
- PAM
- OSAuth
- GSI
- Kerberos

Network:

- TCP
- SSL

Database:

- PostgreSQL
- Oracle
- MySQL



DFC October 2014 NSF Review Slide 8



iRODS 4.0: Continuous Integration

Why Build and Test?

- Transparency (in both process and product)
- Use existing industry best practices
- Coverage -> Confidence in Refactoring
- Packaging -> Ease of installation and upgrade
- Test framework idempotency
- Test independence
- Topology awareness
- Automation, Automation, Automation

iRODS 4.0 transitioned legacy Perl-based test framework to a Python-based framework using unittest.

- We have increased code coverage from ~20% to ~57%
- Status always visible at <http://ci-dev.renci.org/hudson/view/iRODS>



DFC October 2014 NSF Review Slide 9



The iRODS Ecosystem: DFC Contributions

The DataNet Federation Consortium is creating national-scale research data federations.

Ongoing iRODS Development:

- Clients and Client Interfaces
 - Jargon and iDrop Web 3
 - Modeshape (plus WebDAV) plugin
- Messaging Interface
 - Integration with Elastic Search
- Metadata Templating and Ontology Discovery

This material is based upon work supported by the National Science Foundation under Cooperative Agreement OCI-0940841.



DFC October 2014 NSF Review Slide 10





www.datafed.org

www.irods.org



National Science Foundation Cooperative Agreement: OCI-0940841

Update on iRODS

Terrell Russell
iRODS Consortium, RENCi



DFC October 2014 NSF Review Slide 1



New iRODS releases

- 4.0.0 March 2014
- 4.0.1 June 2014
- 4.0.2 June 2014
- 4.0.3 August 2014



DFC October 2014 NSF Review Slide 2



Release History (since Feb 2014)

2014-03-28	4.0.0	Merged Codebase	4 th major release of iRODS and first merged open source release from RENC1.
2014-06-05	4.0.1	1st Point Release	Memory leak fixes, security fixes, --run-in-place, MacOSX support, schema update mechanism.
2014-06-17	4.0.2	2nd Point Release	Random and RoundRobin resource plugin fix, memory leak fixes, microservice fixes, security fixes, large collection recursive operations, and better server-server authentication setup.
2014-08-20	4.0.3	3rd Point Release	More flexible installation options (service account name/group), block storage operation fix, impostor resource, memory leak fixes, and security fixes.



DFC October 2014 NSF Review Slide 3



Highlights of iRODS 4.x

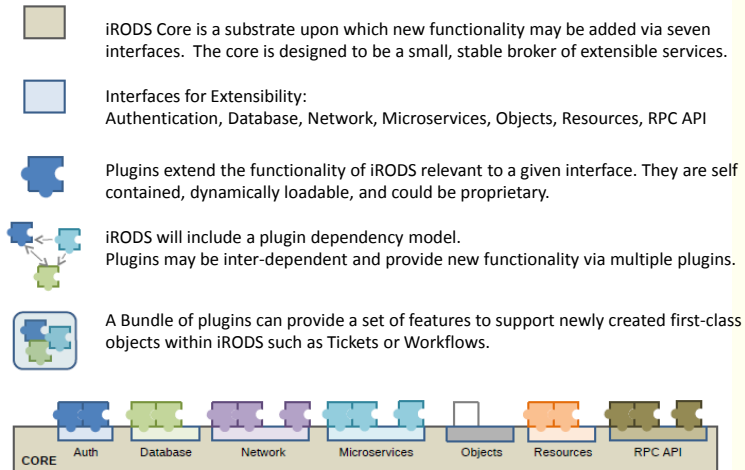
- Architecture Change
 - Pluggable Core
 - Composable Resources
- Build and Test
 - Binary Packages
 - Code Coverage
 - Static Analysis
 - Continuous Integration
 - Across six(+) supported Operating Systems



DFC October 2014 NSF Review Slide 4



Architecture




Plugins

- **29 plugins shipping**
 - 1 API
 - 3 authentication
 - 3 database
 - 6 microservice
 - 2 network
 - 14 resource
- **12 supported add-on plugins**
- **1 externally developed storage plugin**

Composable Resources

Resource Plugins

- Tree Metaphor
- Virtualization
- Extensible



Available

- 14 shipping
 - compound, mso, passthru, roundrobin, unixfilesystem, deferred, mssofile, random, structfile, load_balanced, mockarchive, nonblocking, replication, univmss
- 4 additional
 - WOS, S3, directaccess, HPSS



DFC October 2014 NSF Review Slide 7



Supporting DFC

Indexing Framework

- Rules and microservices
- Messaging based
- Event driven
- Continued hardening

Upgrading/Supporting DFC installations

- Migrations from 3.x to 4.x



DFC October 2014 NSF Review Slide 8



Future enhancements

- Release of new plugins
 - Don't require a core code release
 - User extensions welcome
 - May be added to testing and "certified"
 - May be tested, certified, and added to shipping code
 - May remain uncertified, distributed independently
- Release 4.1 – API compatible, refactoring
- Release 5.0 – Major new functionality



DFC October 2014 NSF Review Slide 9



www.datafed.org

www.irods.org



National Science Foundation Cooperative Agreement: OCI-0940841