iRODS®

Technology Update

Kory Draughn
Chief Technologist
iRODS Consortium

June 13-16, 2023
iRODS User Group Meeting 2023
Chapel Hill, NC

iRODS 4.2 Series

4.2.12 is the final release of the 4.2 series.

Limited to security fixes, bug fixes, and trivial enhancements.

Contributors

iRODS Release	Issues Closed
4.2.12	160

```
~/irods $ git shortlog --summary --numbered 4.2.11..4.2.12
    67  Alan King
    58  Kory Draughn
    13  Daniel Moore
     9  Justin James
     8  Markus Kitsinger (SwooshyCueb)
     6  Martin Jaime Flores Jr
     4  Felix A. Croes
     2  Alastair Smith
     1  Phillip Davis
     1  Terrell Russell
```

4.2.12 Core Server Improvements

- Microservices for read-only access to JSON objects
  - Useful in iRODS Rule Language (NREP) with JSON-based inputs/outputs
- Wider availability of admin keyword in various APIs and libraries
  - imeta
  - atomic ACLs/metadata endpoints
  - filesystem
  - msiDataObjChksum
- Improved user/group/password management
- Fixes and expansive tests for compound resource

Where's iRODS 4.3.1?

4.3.1 is coming together very well, but it needs a bit more work before it's ready.

We appreciate your patience.

For now, let's talk about what will be included once released ...

4.3.1 User Experience Updates

- Removed setup logic for rsyslog and logrotate
  - Multiple implementations of syslog
- Replaced **log_facility** with **server_zone** in log message output
- Exposed client connection information to acPreConnect()
- **ichmod** honors the permission model
- Unixfilesystem resource plugin supports detached mode

https://github.com/irods/irods/tree/7fc671b6

4.3.1 Core Server Enhancements

- **Added support for Address Sanitizer**
- New zone administration library for C++
- New ticket administration library for C++
- **New API plugin: rc_switch_user**
- Compile-time feature test macros
- **iRODS Project Templates for C++**
- **Improved documentation**

https://github.com/irods/irods/tree/7fc671b6

*** Bold items discussed in this talk ***

Address Sanitizer (ASan)

A very fast memory error detector for C/C++.

It detects several different issues such as memory leaks, use-after-free bugs, heap buffer overflows, etc.

Used to track down several memory leaks in iRODS 4.3.0.

Enabled via CMake by setting **IRODS_ENABLE_ADDRESS_SANITIZER** to **YES**.

For example:

```
    user@ugm2023:~ $ cmake ... -DIRODS_ENABLE_ADDRESS_SANITIZER=YES ...
```

New API Plugin - rc_switch_user

Allows the user associated with a connection to be switched to a different user.

Designed for client applications which act as servers (e.g. NFSRODS) and requires a proxied connection.

Benefits

- Avoids TCP connection setup and tear down
- Allows a single connection to be reused for multiple users
- Gets us closer to true connection pooling

New API Plugin - rc_switch_user (cont.)

Performance Testing Details

Setup

- Two custom client applications
- App A connects to a server *N* times as the same user
- App B makes one connection and calls **rc_switch_user** *N* times

Test results show a 98% performance improvement.

iRODS Project Templates for C++

Using the GitHub template repository feature, the iRODS Consortium now offers template repositories which allow C++ developers to jump directly into writing code for iRODS.

The Consortium supports three template repositories today.

- https://github.com/irods/irods_project_template_cpp_client
- https://github.com/irods/irods_project_template_cpp_microservice
- https://github.com/irods/irods_project_template_cpp_rule_engine_plugin

A template for building resource plugins is planned and will become available in the future.

Improved Documentation - Policy Cookbook

An online resource dedicated to providing best practices and the latest techniques to various policy-based situations encountered in the iRODS ecosystem.

The cookbook covers topics such as ...

- Synchronizing Delay Rules using Metadata
- Naming Schemes and Conventions
- Sharing data across PEPs
- Simulating User Quotas
- Implementing maintainable Policy through reusable rules

If you have suggestions on how to improve the cookbook, please reach out.

Improved Documentation - Data Objects

Information about data objects has been expanded.

Documentation for 4.3.1 includes details about ...

- The meaning of each replica status
    - intermediate, write-locked, etc.
- Logical Locking
- High-Level Operations
- R_DATA_MAIN
    - The database table which holds all replica information

We'll continue to expand on these topics as improvements to the server are made.

Core Development Team Talks

- Not in This Talk / Separate Talks
  - Terrell Russell and Violet White
    - iRODS S3 API: Presenting iRODS as S3
  - Kory Draughn
    - GenQuery2: A more standardized, powerful parser for the iRODS namespace
    - iRODS HTTP API
  - Derek Dong, Kory Draughn, and Terrell Russell
    - The iRODS CLI we deserve
  - Martin Flores
    - Authentication in iRODS 4.3: Investigating OAuth2 and OpenID Connect (OIDC)

- Included in This Talk
  - Justin James
    - S3 Resource Plugin
    - Globus Connector
  - Daniel Moore
    - Python iRODS Client
  - Markus Kitsinger
    - Audit AMQP Rule Engine Plugin
    - Python Rule Engine Plugin
    - Build and Packaging

S3 Resource Plugin Updates

S3 vendors certified as compatible with iRODS

- Fujifilm
- Oracle Cloud
- Wasabi S3

Bug Fixes and Enhancements

- Streaming of a file from local cache
  - Reorganized code so that each thread can handle more than one part
    - Allows better handling of part timeouts
    - Reduces the part size so the 2 minute timeout isn't triggered on large files
    - Allowed streaming of files from cache that are larger than 50 GiB
- Fixed a bug where the catalog consumer thought it was the provider when detached mode was enabled
- Fixed 4 MiB to 32 MiB file replication failures due to unexpected number of transfer threads

Globus Connector Updates

Two issues were fixed when testing the Globus Connector.

Issue #45

Update time not being updated due to the thread that closes the file not writing any data. Reorganized the code to make sure this thread takes part in the writing of data.

Issue #50

We were getting conflicts in some environments due to duplicate base64_encode and base64_decode methods being defined. These methods were moved into the Globus code with namespaces. The iRODS core code was also updated later to namespace these methods.

Python iRODS Client - from 1.1.4 to 1.1.8

Thanks to our contributors!

```
~/python-irodsclient $ git shortlog --summary --numbered v1.1.4..v1.1.8
    28  d-w-moore
     4  Terrell Russell
     1  Gwenael Leysour de Rohello
     1  John Constable
     1  Martin Jaime Flores Jr
     1  Paul Borgermans
     1  Sietse Snel
     1  jpmcfarland
```

Python iRODS Client - from 1.1.4 to 1.1.8

Major improvements

- Connection timeout fix
- ACLs interface (replaces permissions - more consistent)
- Correct path normalization
- Automatic SSL context handling / generation
- groupadmin capabilities
- Resource properties and methods to yield parent, hierarchy, etc

Minor improvements

- Windows compatibility; Unregister replicas; Guard password integrity; Session auto-close; GenQuery (NOT LIKE)

Python iRODS Client - Ongoing and Upcoming Work

Works in Progress

- Client-to-Resource redirect
- Auto-closing data objects

Yet to Come

- iRODS 4.3.0 authentication framework compatibility

Audit AMQP Rule Engine Plugin

- Modernization
  - Refactored to use nlohmann-json instead of jansson
  - Refactored to use qpid-proton's C++ API
  - Migrated to new logging framework
  - Miscellaneous other modernization
- Housekeeping
  - Repository reorganized and code reformatted
  - RPM package installation less fussy
  - Removed unused amqp_options configuration setting
  - Miscellaneous other housekeeping
- Removed JSON wrapper tokens
- Fixed JSON types for some fields
- More AMQP message metadata set
- Better handling of default configuration

Audit AMQP Rule Engine Plugin - ELK Stack

- Modernization
  - New Dockerfile syntax
  - Updated entire software stack
    - Container base image
    - Elasticsearch, Kibana, RabbitMQ
    - Temurin JDK
- Housekeeping
  - Reduced number and size of intermediate container images
  - Excluded more unneeded files from container image
- Updated for use with new version of the rule engine plugin
  - Workarounds for use with older/current versions of the plugin are togglable
- Replaced logstash with a Python daemon using qpid-proton's Python API
- Moved as much setup as possible to container build-time
- Added argument for specifying Java heap size

Python Rule Engine Plugin

- Repository reorganized and code reformatted
- Build-time memory usage reduced significantly
- Parallelized build time reduced by an order of magnitude
- Fixed package dependency declarations
    - No more pulling in unneeded development packages
- Miscellaneous housekeeping

Build and Packaging

We continue to move towards a more Normal and Boring approach to build and packaging.

- clang-tidy brings us closer to building against libstdc++
- Work has begun on unprivileged build and packaging in development environment containers

As work towards this goal continues, there will be significant changes to the pre-build configuration process. We plan to use CMake presets to smooth the transition.

We still do not have a time-table. See my talk at UGM 2022 for more details on our plans.

iRODS Internships - Summer 2023

Implement PUT_SYNC from S3 for iRODS Automated Ingest

The Automated Ingest tool is a Python application designed to keep the iRODS catalog up to date with changes in an existing filesystem or S3 bucket. The S3 scanner has implemented the REGISTER_SYNC operation for registering data in-place, but does not yet know how to make a copy of the scanned data into iRODS. Implementing PUT_SYNC will fill out this requested feature.

Add features to Zone Management Tool (ZMT)

The iRODS Zone Management Tool has become mostly feature complete in the last year. However, there are a few things that it still does not yet know how to manage. The open issues currently cover management of iRODS Tickets, the Delay Server, and a number of new health checks.

Implement new version of ZoneReport

An iRODS Zone can describe its own configuration with a ZoneReport. The schema that defines the ZoneReport is now a bit out of date as the server itself has changed for 4.3.0. We would like to refactor the zone_bundle.json schema and update the machinery that produces and depends on this format. Known issues include the naming of server roles, duplicate plugins, hierarchy information for resources, and self-aware versioning. Cleaning this up will affect the testing environment, downgrading service accounts, and clients such as the Zone Management Tool (ZMT).

iRODS Internships - Summer 2023 (cont.)

Document XML protocol

iRODS implements its own protocol. The protocol supports two encodings, binary and XML. To communicate with an iRODS server using either of these encodings means a developer must implement the protocol for the target programming language. However, this presents a real challenge because there is no formal documentation explaining the design or behavior of the protocol. This causes developers to look at other implementations for guidance, which normally results in differences between implementations.

Developers looking to implement the protocol will have a better experience if documentation is available. This will also help the Consortium understand what needs to be considered when designing and implementing a replacement.

https://irods.org/2023/01/irods-internship-summer-2023/

Big Picture

Core

- 4.3.x - Satisfy Roadmap (Cloud-friendliness, Replace PackStruct, etc)

Clients

- GUIs (Metalnx, ZMT, Kanki, et al.)
- Onboarding and Syncing (Automated Ingest)
- File System Integration (NFSRODS, SFTP)
- iRODS CLI
- iRODS HTTP API

Continue building out policy components (Capabilities).

We want installation and management of iRODS to become about policy design, composition, and configuration.

Please share your ...

- Use cases
- Pain points
- Hopes and dreams

Open Source Community Engagement

Get Involved

- Working Groups
- GitHub Issues
- Pull Requests
- Chat List
- Consortium Membership

Tell Others

- Publish, Cite, Advocate, Refer