iRODS®

Technology Update

Kory Draughn
Chief Technologist
iRODS Consortium

June 17-20, 2025
iRODS User Group Meeting 2025
Durham, NC

Contributors

iRODS Release	Issues Closed
4.3.3	89
4.3.4	134
5.0.0	178
5.0.1	5

Kory Draughn

Alan King

Markus Kitsinger

Derek Dong

Martin Jaime Flores Jr

Terrell Russell

Justin James

Daniel Moore

Nishant Dash

Violet White

Felix A. Croes

Phillip Davis

Awab Masroor

Peter Verraedt

June Releford

Leonardo Lenoci

4.3.3 Release Highlights

- Removed support for CentOS 7
- Support for Ubuntu 24.04
- Support for PostgreSQL 15
- Support for newer compilers and libstdc++
- Plugin package names include the server version which they were built against
- GenQuery2 improved support for SQL functions
  - SQL function syntax can be used in the WHERE and ORDER-BY clauses
  - Understands multi-argument SQL functions
- 10x performance improvement to timestamp generation for log messages

https://irods.org/2024/08/irods-4-3-3-is-released/

4.3.4 Release Highlights

Prepares the ground for iRODS 5.

- More deprecations
  - `ilocate`, `igetwild`, `imeta qu`, `imeta`'s interactive mode
  - Server monitoring tools, modification of ACL policy - i.e. Strict ACLs
- More bug fixes
- Cleaned up implementations
- `imiscsvrinfo` reports SSL/TLS certificate information
  - Made possible through enhancements to the `rcGetMiscSvrInfo` API endpoint
- Server removes ACLs following the removal of a user
- Expanded permission levels supported by the atomic ACL operations API endpoint

https://irods.org/2025/03/irods-4-3-4-is-released/

4.3 series EOL

Planning for 4.3.5 to be the final release of the 4.3 series.

Report those bugs!

5.0.0, 5.0.1 Releases

iRODS 5 represents a significant effort to deliver an enterprise grade data management platform.

Great focus has been put into making iRODS 5 easy to deploy and manage in a variety of environments.

https://irods.org/2025/05/irods-5-0-0-is-released/

https://irods.org/2025/05/irods-5-0-1-is-released/

iRODS 5 - Notable Mentions

- Released for Ubuntu 22, Ubuntu 24, Debian 12, and Enterprise Linux 9
- Semantic Versioning, decoupled versioning for plugins
- Migrated to OpenSSL 3
- Removed deprecated functionality
- Removed the ability to disable Strict ACLs
- Dedicated log categories for the database plugin, SQL, and GenQuery1

iRODS 5 - Revamped Startup/Shutdown Process

- Simplified implementation, deterministic behavior
- Server does not auto-reload changes to `server_config.json`
- Honors signals in the expected manner
  - `SIGINT`/`SIGTERM` for fast shutdown, `SIGQUIT` for graceful shutdown
  - `SIGHUP` for reloading configuration
    - `kill -HUP $(cat /var/run/irods/irods-server.pid)`
- Servers can be launched out of order
  - Provider waits for the database to become available
  - Consumers wait for the Provider to become available
- Separated upgrade step from server startup - `upgrade_irods.py`
- Server communicates its status to the service manager
- Removed the control plane, `irods-grid`, and `irodsctl`

iRODS 5 - Unified Server Configuration

All server configuration lives in `server_config.json`.

- Zone-wide options are stored in the catalog
- Server does not rely on service account's `irods_environment.json`

```
{
  "client_server_policy": "CS_NEG_REFUSE",
  "tls_client": {
    "ca_certificate_file": "<path>",
    "ca_certificate_path": "<path>",
    "verify_server": "hostname"
  },
  "tls_server": {
    "certificate_chain_file": "<path>",
    "certificate_key_file"" "<path>",
    "dh_params_file": "<path>"
  }
}
```

```
{
  "encryption": {
    "algorithm": "AES-256-CBC",
    "key_size": 32,
    "num_hash_rounds": 16,
    "salt_size": 8
  }
}
```

https://docs.irods.org/5.0.1/system_overview/configuration/

iRODS 5 - Access Time Tracking

The server has been updated to track access times for all replicas.

- Each replica of a data object maintains its own independent access time
- Access time is updated in the background via per-server memory queue
- Behavior is similar to Linux relatime

Access time tracking is controlled by four zone-wide configuration options.

Options can be modified by `iadmin set_grid_configuration`.

https://docs.irods.org/5.0.1/system_overview/data_objects/#access-time

iRODS 5 - GenQuery1

Replaced ad-hoc string parser with flex-bison parser.

- Supports escaping individual bytes using hexadecimal
- Supports escaping single quotes by using two adjacent single quotes

All known issues related to special characters, embedded single quotes, and reserved keywords have been resolved.

This improvement is limited to `iquest` and C/C++ clients only.

```
iquest "select DATA_NAME where DATA_NAME = 'that''s my data\x21'"
```

https://docs.irods.org/5.0.1/system_overview/genquery/#escaping-bytes-and-special-characters

iRODS 5 - GenQuery2

Updated so that users have full control over the placement of the `DISTINCT` keyword.

Motivated by a need for more control over queries which involve counting.

This is a major departure from GenQuery1. This change is NOT backward compatible with iRODS 4.3.4 and earlier.

```
# List all replicas
iquery "select COLL_NAME, DATA_NAME"

# List all data objects
iquery "select distinct COLL_NAME, DATA_NAME"

# Count data objects
iquery "select count(distinct DATA_NAME)"
```

iRODS 5 - Delay Rule Locking

Keeps multiple Delay Servers from executing the same delay rule simultaneously.

Delay rules are locked after being pulled from the in-memory queue, but before execution.

Locking a delay rule results in the Delay Server host information being attached to the delay rule entry in the Catalog.

- Delay Server's FQDN/Hostname
- Delay Server's PID
- Timestamp representing when the lock was acquired

Simplifies Delay Server Migration by allowing multiple Delay Servers to run in the same zone simultaneously.

Plugins and Clients

New Rule Engine Plugin Operations

Motivated by a need for rule engine plugins to manage resources external to the iRODS server.

- `setup`
  - For initialization of static configuration and/or external resources
  - Can be used to elide frequent reading of configuration files
  - Run during Agent Factory startup
- `teardown`
  - For cleaning up resources initialized by `setup`
  - Run during Agent Factory shutdown

Storage Tiering Capability Plugin

Plugin stability is significantly improved.

- Resolved all issues stemming from a lack of user permissions
- Data movement and metadata manipulation operations use the local service account
- Fixed various memory leaks
- Fixed issues with packaging for Enterprise Linux 8 and 9
- Access time metadata is updated on more API PEPs
  - i.e. `touch`, `get`, `replica_open`, `replica_close`
- Execution of rules provided by the plugin require `rodsadmin` level privileges
- Enabled use of the iRODS 4.3 logger

Automated Ingest Capability

Release	Issues Closed
0.5.0	11
0.6.0	47

- 0.5.0 introduced a breaking, but essential change to pre/post event handler methods
- Refactored implementation for clarity
- Added ability to delete data objects and collections from iRODS which DO NOT exist in the source being ingested *(**`rsync --delete`**-like functionality)*
- Added support for recursive removal in parallel
  - ~55% improvement over `irm -f`
- Expanded information reported about individual jobs

Python Rule Engine Plugin

- Fixed thread management for Python interpreter
- Removed Python 2 compatibility code
- Added support for GenQuery2 to Query module

PAM Interactive Authentication Plugin

A plugin that enables organizations to develop dynamic authentication flows which meet their needs.

Three releases since UGM 2024 ...

- **0.1.0** - 2024-08-26 *(initial release)*
- **0.1.1** - 2025-03-04
  - Fixed RPM packaging issues, no functional changes
- **0.1.2** - 2025-05-20
  - Compatible with iRODS 5

S3 Resource Plugin

- Bug fixes
  - Improved handling of number of threads and shared memory
  - Fixed crash due to mishandling of `std::string` objects
- Introduced the `checksum_read_buffer_size_in_bytes` configuration property
  - Allows admins to adjust the buffer size used for calculating checksums
- Cleaned up implementation
- Migrated to distro-provided packages

Globus Connector

- Fixed issue handling filenames containing apostrophes
- Updated framework used for testing
  - Added support for testing against unreleased versions of iRODS
  - Added support for additional Operating Systems
    - Ubuntu 22, Ubuntu 24, Debian 11, Debian 12
    - Enterprise Linux 8, Enterprise Linux 9

iRODS S3 API

- Implemented multipart enhancements
    - Conditionally write part directly to iRODS server
        - If the file offset can be determined, write part to iRODS server
        - Otherwise, write to local file, write to iRODS on `CompleteMultipartUpload`
    - Improved upload performance by ~40%
- Added support for more S3 API operations
    - e.g. `AbortMultipartUpload`

Big Picture

iRODS 5 delivers significant improvements over iRODS 4.

The Consortium is listening and is working hard to deliver the tools and functionality to solve enterprise-level data management challenges.

But, there's more work to be done.

Open Source Community Engagement

Get Involved

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

Consortium Membership

- Tell others
- Publish, cite, advocate, refer

Thank you!

Upcoming talks from the Core Development Team

- Python iRODS Client v3.1.1
  - Daniel Moore
- iRODS Build and Packaging: 2025 Update
  - Markus Kitsinger
- iRODS Roadmap 2025
  - Kory Draughn
  - Terrell Russell
- irods4j: A new Java client library designed for iRODS 4.3.2+
  - Kory Draughn
- Metalnx v3.1.0
  - Justin James
- iRODS HTTP API v0.5.0
  - Martin Flores