



Author: Renaissance Computing Institute (RENCI)
Version: 3.0.1
Date: 2013-11-16

Table of Contents

1	Release Notes	5
2	License	5
3	Overview	6
4	Download	6
	4.1 Binaries	6
	4.2 Open Source	6
5	Installation	7
6	Quickstart	7
	6.1 Changing the administrator account password	8
	6.2 Changing the Zone name	8
	6.3 Add additional resource(s)	9
	6.4 Add additional user(s)	10
7	Upgrading	10
	7.1 RPM based systems	10
	7.2 DEB based systems	10
8	Migration from Community iRODS	10
9	Federation with Community iRODS	11
10	Backing Up	11
11	Assumptions	11
12	Architecture	11
	12.1 Dynamic Policy Enforcement Points	12
	12.1.1 Available Plugin Operations	12
	12.1.2 Available Values within Dynamic PEPs	13
13	Pluggable Microservices	15
14	Composable Resources	15
	14.1 Tree Metaphor	15
	14.2 Virtualization	16
	14.3 Coordinating Resources	17
	14.3.1 Compound	17
	14.3.2 Random	17
	14.3.3 Replication	18
	14.3.4 Round Robin	18
	14.3.5 Passthru	18
	14.3.6 Expected	18
	14.4 Storage Resources	18

14.4.1	Unix File System	18
14.4.2	Structured File Type (tar, zip, gzip, bzip)	18
14.4.3	Non-Blocking	19
14.4.4	Mock Archive	19
14.4.5	Universal Mass Storage Service	19
14.4.6	Expected	19
14.5	Managing Child Resources	19
14.6	Example Usage	20
14.6.1	Example 1	20
14.7	Rebalancing	20
15	Pluggable Authentication	20
16	Pluggable Network	21
17	Users & Permissions	21
18	Rule Engine	22
18.1	File Locking	22
18.2	Delay execution	22
19	Authentication	22
19.1	PAM	23
19.1.1	User Setup	23
19.1.2	Server Configuration	23
19.1.3	Server SSL Setup	23
19.1.3.1	Generate a new RSA key	23
19.1.3.2	Acquire a certificate for the server	24
19.1.3.3	Create the certificate chain file	24
19.1.3.4	Generate OpenSSL parameters	24
19.1.3.5	Place files within accessible area	25
19.1.3.6	Set SSL environment variables	25
19.1.3.7	Restart E-iRODS	25
19.1.4	Client SSL Setup	25
19.1.5	Environment Variables	26
20	Configuration	27
21	Troubleshooting	27
21.1	Common Errors	27
21.1.1	E-iRODS Server is down	27
21.1.2	No such file or directory	28
21.1.3	No rows found in the E-iRODS Catalog	28

Enterprise iRODS (E-iRODS) Manual

21.1.4	Access Control and Permissions	28
21.1.5	Credentials	28
22	Glossary	28
23	Known Issues	32
24	History of Releases	32

1 Release Notes

This is the second release of the Enterprise integrated Rule-Oriented Data System (E-iRODS).

E-iRODS is developed under the auspices of the iRODS Consortium. This release was prepared by the Renaissance Computing Institute (RENCI) and released under the New BSD (BSD-3) License.

Development of this release has focused on the following features and efforts:

- API Compatibility with Community iRODS for purposes of Federation
- Met all US Export Control requirements for distribution
- Introduced pluggable interface for networks
 - TCP
 - SSL
- Introduced pluggable interface for authentication
 - Native iRODS password
 - OSAuth
 - PAM
 - LDAP (via PAM)
- Additional coordinating resource plugins
 - Compound
 - Replication
- Migration support for composable resources
 - Rebalance mechanism
 - Adding/removing child resources with existing data
 - itrim from child resources of replication resources
- Tree view for visualizing resource hierarchies
- Full python-based testing framework
- Certified against full Jargon test suite
- Continuous integration testing via hudson
- Continuous static analysis via cppcheck
- Increased code coverage from ~50% to over 56%
- Continuously built and tested across 3 major linux distributions
- Packaged for 3 major linux distributions
- Support for package upgrade via package manager

2 License

Copyright (c) 2005-2013, Regents of the University of California, the University of North Carolina at Chapel Hill, and the Data Intensive Cyberinfrastructure Foundation All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- Neither the name of the University of California, San Diego (UCSD), the University of North Carolina at Chapel Hill nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

3 Overview

This manual attempts to provide standalone documentation for E-iRODS (<http://ei rods.org>) as packaged by the Renaissance Computing Institute (RENCI).

<http://ei rods.org>

Additional documentation is available on the iRODS wiki (<http://i rods.org>), the iRODS Doxygen site (<http://i rods.org/doxygen>), and in the two books published by the iRODS team:

(2010) iRODS Primer: integrated Rule-Oriented Data System (Synthesis Lectures on Information Concepts, Retrieval, and Services) <http://www.amazon.com/dp/1608453332>

(2011) The integrated Rule-Oriented Data System (iRODS 3.0) Micro-service Workbook <http://www.amazon.com/dp/1466469129>

4 Download

E-iRODS is released in both binary package format and with full source code.

4.1 Binaries

RPM and DEB formats are available for both iCAT-enabled servers and resource-only servers. There are variations available for combinations of platform and operating system.

More combinations will be made available as our testing matrix continues to mature and increase in scope.

The latest files can be downloaded from <http://ei rods.org/download>.

4.2 Open Source

Repositories, ticket trackers, and source code are available from the developers area of the E-iRODS website.

<http://ei rods.org/developers>

5 Installation

The available packages declare the dependencies necessary to run E-iRODS and if satisfied, they install a service account and group named 'eirops', the E-iRODS binaries, microservice documentation, and this manual.

Installation of the Postgres iCAT DEB:

```
$ (sudo) dpkg -i eirops-3.0.1-64bit-icat-postgres.deb
$ (sudo) apt-get -f install
```

Installation of the Resource RPM:

```
- Make sure to read ./packaging/RPM_INSTALLATION_HOWTO.txt before trying to install the RPM package.
$ (sudo) rpm -i eirops-3.0.1-64bit-resource.rpm
```

For the iCAT-enabled server packages, the E-iRODS server and EICAT database are started automatically with default values:

```
eirops@hostname:~/ $ ienv
NOTICE: Release Version = rods3.0, API Version = d
NOTICE: irodsHost=hostname
NOTICE: irodsPort=1247
NOTICE: irodsDefResource=demoResc
NOTICE: irodsHome=/tempZone/home/rods
NOTICE: irodsCwd=/tempZone/home/rods
NOTICE: irodsUserName=rods
NOTICE: irodsZone=tempZone
NOTICE: irodsClientServerNegotiation=request_server_negotiation
NOTICE: irodsClientServerPolicy=CS_NEG_REFUSE
NOTICE: irodsEncryptionKeySize=32
NOTICE: irodsEncryptionSaltSize=8
NOTICE: irodsEncryptionNumHashRounds=16
NOTICE: irodsEncryptionAlgorithm=AES-256-CBC
```

For the resource-only packages, the server is not started automatically. The administrator will need to run the `./packaging/setup_resource.sh` script and provide the following five pieces of information before E-iRODS can start and connect to its configured iCAT Zone:

1. Hostname or IP
2. iCAT Port
3. iCAT Zone
4. E-iRODS administrator username
5. E-iRODS administrator password

6 Quickstart

Successful installation will complete and result in a running E-iRODS server. If you installed an iCAT-enabled E-iRODS server, a database of your choice will also have been created and running. The iCommand `ils` will list your new iRODS administrator's empty home directory in the iRODS virtual filesystem:

```
eirods@hostname:~/ $ ils
/tempZone/home/rods:
```

When moving into production, you will probably want to cover the next few basic steps:

6.1 Changing the administrator account password

The default installation of E-iRODS comes with a single user account 'rods' that is also an admin account ('rodsadmin') with the password 'rods'. You should change the password before letting anyone else into the system:

```
eirods@hostname:~/ $ iadmin moduser rods password <newpassword>
```

To make sure everything succeeded, you will need to re-authenticate and check the new connection:

```
eirods@hostname:~/ $ iinit
Enter your current iRODS password:
eirods@hostname:~/ $ ils
/tempZone/home/rods:
```

If you see an authentication or other error message here, please try again. The password update only manipulates a single database value, and is independent of other changes to the system.

6.2 Changing the Zone name

The default installation of E-iRODS comes with a Zone named 'tempZone'. You probably want to change the Zone name to something more domain-specific:

```
eirods@hostname:~/ $ iadmin modzone tempZone name <newzonename>
If you modify the local zone name, you and other users will need to
change your .irodsEnv files to use it, you may need to update
irods.config and, if rules use the zone name, you'll need to update
core.re. This command will update various tables with the new name
and rename the top-level collection.
Do you really want to modify the local zone name? (enter y or yes to do so):y
OK, performing the local zone rename
```

Once the Zone has been renamed, you will need to update your .irodsEnv file to match (note the three places where the updated zone name is located):

```
eirods@hostname:~/ $ cat .irods/.irodsEnv
# iRODS server host name:
irodsHost '<hostname>'
# iRODS server port number:
irodsPort 1247
# Default storage resource name:
irodsDefResource 'demoResc'
# Home directory in iRODS:
irodsHome '/*<newzonename>*/home/rods'
# Current directory in iRODS:
irodsCwd '/*<newzonename>*/home/rods'
# Account name:
```



```
irodsUserName 'rods'
# Zone:
irodsZone '**<newzonename>*'
# Enable Advanced Client-Server negotiation:
irodsClientServerNegotiation 'request_server_negotiation'
# Client-Server connection policy:
irodsClientServerPolicy 'CS_NEG_REFUSE'
# Client-Server Encryption Key Size In Bytes:
irodsEncryptionKeySize '32'
# Client-Server Encryption Salt Size In Bytes:
irodsEncryptionSaltSize '8'
# Client-Server Encryption Number of Hash Rounds:
irodsEncryptionNumHashRounds '16'
# Client-Server Encryption Algorithm:
irodsEncryptionAlgorithm 'AES-256-CBC'
```

Now, the connection should be reset and you should be able to list your empty iRODS collection again:

```
ei rods@hostname:~/ $ iinit
Enter your current iRODS password:
ei rods@hostname:~/ $ ils
/<newzonename>/home/rods:
```

6.3 Add additional resource(s)

The default installation of E-iRODS comes with a single resource named 'demoResc' which stores its files in the `/var/lib/eirods/iRODS/Vault` directory. You will want to create additional resources at disk locations of your choosing as the 'demoResc' may not have sufficient disk space available for your intended usage scenarios. The following command will create a basic 'unix file system' resource at a designated host at the designated full path:

```
ei rods@hostname:~/ $ iadmin mkresc <newrescname> 'unix file system' <fully.qualified.domain.name>:</full/path/to/new/vault>
```

Additional information about creating resources can be found with:

```
ei rods@hostname:~/ $ iadmin help mkresc
mkresc Name Type [Host:Path] [ContextString] (make Resource)
Create (register) a new storage or database resource.

Name is the name of the new resource.
Type is the resource type.
Host is the DNS host name.
And Path is the defaultPath for the vault.
ContextString is any contextual information relevant to this resource.
(semi-colon separated key=value pairs e.g. "a=b;c=d")

A ContextString can be added to a coordinating resource (where there is
no hostname or vault path to be set) by explicitly setting the Host:Path
to an empty string ('').
```

Creating new resources does not make them default for any existing or new users. You will need to make sure that default resources are properly set for newly ingested files.

6.4 Add additional user(s)

The default installation of E-iRODS comes with a single user 'rods' which is a designated 'rodsadmin' type user account. You will want to create additional user accounts (of type 'rodsuser') and set their passwords before allowing connections to your new grid:

```

eirods@hostname:~/ $ iadmin mkuser <newusername> rodsuser

eirods@hostname:~/ $ iadmin lu
rods#tempZone
<newusername>#tempZone

eirods@hostname:~/ $ iadmin help mkuser
mkuser Name[#Zone] Type (make user)
Create a new iRODS user in the ICAT database

Name is the user name to create
Type is the user type (see 'lt user_type' for a list)
Zone is the user's zone (for remote-zone users)

Tip: Use moduser to set a password or other attributes,
use 'aua' to add a user auth name (GSI DN or Kerberos Principal name)

```

It is best to change your Zone name before adding new users as any existing users would need to be informed of the new connection information and changes that would need to be made to their local `.irodsEnv` files.

7 Upgrading

E-iRODS did not support upgrading until 3.0.1rc1. From that point forward, upgrading should be handled by the host Operating System via the package manager as expected.

7.1 RPM based systems

```
$ (sudo) rpm -U eirops-3.0.1-64bit-icat-postgres-suse.rpm
```

Upgrading from before 3.0.1rc1 will be detected and require the running of an additional script to get your installation in working order. The RPM based upgrade for those older installs moves aside the `/var/lib/eirops/` home directory to `/var/lib/eirops_new/`. To restore this directory and reinstall all the correct packaged files, run the following script with the new RPM file:

```
$ sudo /var/lib/eirops_new/packaging/post30upgrade.sh <newfile.rpm>
```

7.2 DEB based systems

```
$ (sudo) dpkg -i eirops-3.0.1-64bit-icat-postgres.deb
```

8 Migration from Community iRODS

Support for migrating from Community iRODS is planned, but automated scripts and documentation have not yet been completed.

This section will be updated when support is included and tested.

9 Federation with Community iRODS

Enterprise iRODS has made some additions to the database tables for the resources (`r_resc_main`) and Data Objects (`r_data_main`) for the purposes of tracking resource hierarchy, children, parents, and other relationships. These changes would have caused a cross-zone query to fail when the target zone is a community version of iRODS.

In order to support commands such as `ils` and `ilsresc` across an Enterprise to Community federation, E-iRODS will detect the cross zone query and subsequently strip out any requests for columns which do not exist in the community table structure in order to allow the query to succeed.

There are currently no known issues with Federation, but this has not yet been comprehensively tested.

10 Backing Up

Backing up E-iRODS involves: The data, the iRODS system and configuration files, and the iCAT database itself.

Configuration and maintenance of this type of backup system is out of scope for this document, but is included here as an indication of best practice.

1. The data itself can be handled by the iRODS system through replication and should not require any specific backup efforts worth noting here.
2. The iRODS system and configuration files can be copied into iRODS as a set of Data Objects by using the `msiServerBackup` microservice. When run on a regular schedule, the `msiServerBackup` microservice will gather and store all the necessary configuration information to help you reconstruct your iRODS setup during disaster recovery.
3. The iCAT database itself can be backed up in a variety of ways. A PostgreSQL database is contained on the local filesystem as a `data/` directory and can be copied like any other set of files. This is the most basic means to have backup copies. However, this will have stale information almost immediately. To cut into this problem of staleness, PostgreSQL 8.4+ includes a feature called "[Record-based Log Shipping](#)". This consists of sending a full transaction log to another copy of PostgreSQL where it could be "re-played". This would bring the copy up to date with the originating server. Log shipping would generally be handled with a cronjob. A faster, seamless version of log shipping called "[Streaming Replication](#)" was included in PostgreSQL 9.0+ and can keep two PostgreSQL servers in sync with sub-second delay.

11 Assumptions

E-iRODS enforces that the database in use (PostgreSQL) is configured for UTF-8 encoding. This is enforced at the database level and then the tables inherit this setting.

The iRODS setting 'StrictACL' is configured on by default in E-iRODS. This is different from the community version of iRODS and behaves more like standard Unix permissions. This setting can be found in the `server/config/reConfigs/core.re` file under `acAclPolicy{}`.

12 Architecture

E-iRODS represents a major effort to analyze, harden, and package iRODS for sustainability, modularization, security, and testability. This has led to a fairly significant refactorization of much of the underlying codebase. The following descriptions are included to help explain the architecture of E-iRODS.

The core is designed to be as immutable as possible and serve as a bus for handling the internal logic of the business of iRODS (data storage, policy enforcement, etc.). Six or seven major interfaces will be exposed by the core and will allow extensibility and separation of functionality into plugins. A few plugins are included by default in E-iRODS to provide a set of base functionality.

The planned plugin interfaces and their status are listed here:

Plugin Interface	Status	Since
Pluggable Microservices	Complete	3.0b2
Composable Resources	Complete	3.0b3
Pluggable Authentication	Complete	3.0.1b1
Pluggable Network	Complete	3.0.1b1
Pluggable Database	Planned	
Pluggable RPC API	Planned	
Pluggable Rule Engine	Requested	

12.1 Dynamic Policy Enforcement Points

E-iRODS has introduced the capability for dynamic policy enforcement points (PEP). For every operation that is called, two policy enforcement points are constructed (both a pre and post variety), and if it has been defined in `core.re` or any other loaded rulebase file they will be executed by the rule engine.

The PEP will be constructed of the form `"pep_PLUGINOPERATION_pre"` and `"pep_PLUGINOPERATION_post"`.

For example, for `"resource_create"`, the two PEPs that are dynamically evaluated are `pep_resource_create_pre(*OUT)` and `pep_resource_create_post(*OUT)`. If either or both have been defined in a loaded rulebase file (`core.re`), they will be executed as appropriate.

The flow of information from the pre PEP to the plugin operation to the post PEP works as follows:

- `pep_PLUGINOPERATION_pre(*OUT)` - Should produce an `*OUT` variable that will be passed to the calling plugin operation
- `PLUGINOPERATION` - Will receive any `*OUT` defined by `pep_PLUGINOPERATION_pre(*OUT)` above and will pass its own `*OUT` variable to `pep_PLUGINOPERATION_post()`
- `pep_PLUGINOPERATION_post()` - Will receive any `*OUT` from `PLUGINOPERATION`. If the `PLUGINOPERATION` itself failed, the `*OUT` variable will be populated with the string `"OPERATION_FAILED"`.

12.1.1 Available Plugin Operations

The following operations are available for dynamic PEP evaluation. At this time, only very few operations themselves consider the output (`*OUT`) of its associated pre PEP.

Plugin Type	Plugin Operation
Resource	resource_create resource_open resource_read resource_write resource_stagetocache resource_synctoarch resource_registered resource_unregistered resource_modified resource_resolve_hierarchy resource_rebalance

Authentication	auth_client_start auth_agent_start auth_establish_context auth_agent_client_request auth_agent_auth_request auth_agent_client_response auth_agent_auth_response auth_agent_auth_verify
Network	network_client_start network_client_stop network_agent_start network_agent_stop network_read_header network_read_body network_write_header network_write_body

12.1.2 Available Values within Dynamic PEPs

The following Key-Value Pairs are made available within the running context of each dynamic policy enforcement point (PEP) based both on the plugin type as well as the first class object of interest. They are available via the rule engine in the form of `$KVPairs.VARIABLE_NAME` and are originally defined in `iRODS/lib/core/include/rodsKeyWdDef.h`.

Plugin Type	First Class Object	Variable Name
-------------	--------------------	---------------

Enterprise iRODS (E-iRODS) Manual

Resource	Data Object	physical_path mode_kw flags_kw resc_hier
	File Object	logical_path file_descriptor l1_desc_idx file_size repl_requested in_pdmo
	Structured Object	host_addr zone_name port_num sub_file_path offset dataType oprType
	Special Collection	spec_coll_class spec_coll_type spec_coll_obj_path spec_coll_resource spec_coll_resc_hier spec_coll_phy_path spec_coll_cache_dir spec_coll_cache_dirty spec_coll_repl_num

Plugin Type	First Class Object	Variable Name
Authentication	Native Password OS Auth PAM	zone_name user_name digest

Plugin Type	First Class Object	Variable Name
Network	TCP	tcp_socket_handle
	SSL	ssl_host ssl_shared_secret ssl_key_size ssl_salt_size ssl_num_hash_rounds ssl_algorithm

For example, within a PEP, you could reference `$KVPairs.file_size` and get the size of the file currently in context. Likewise, `$KVPairs.ssl_host` would provide the current hostname involved in an SSL connection.

Also, `$pluginInstanceName` is an additional available session variable that gives the instance name of the plugin from which the call is made.

For example, when running `iput -R myOtherResc newfile.txt`, a `fileCreate()` operation is called on "myOtherResc". This delegates the call to the myOtherResc plugin instance which is a

"resource_create" operation. When the `pep_resource_create_pre()` rule is evaluated, the value of `$pluginInstanceName` will be "myOtherResc". This allows rule authors to make decisions at a per-resource basis for this type of operation.

13 Pluggable Microservices

E-iRODS is in the process of being modularized whereby existing community iRODS functionality will be replaced and provided by small, interoperable plugins. The first plugin functionality to be completed was pluggable microservices. Pluggable microservices allow users to add new microservices to an existing E-iRODS server without recompiling the server or even restarting any running processes. A microservice plugin contains a single compiled microservice shared object file to be found by the server. A separate development package, including an example, is available at <http://ei rods.org/download>, and explains how this works in more detail.

14 Composable Resources

The second area of modularity to be added to E-iRODS consists of composable resources. Composable resources replace the concept of resource groups from community iRODS. There are no resource groups in E-iRODS.

14.1 Tree Metaphor

In computer science, a tree is a data structure with a hierarchical representation of linked nodes. These nodes can be named based on where they are in the hierarchy. The node at the top of a tree is the root node. Parent nodes and child nodes are on opposite ends of a connecting link, or edge. Leaf nodes are at the bottom of the tree, and any node that is not a leaf node is a branch node. These positional descriptors are helpful when describing the structure of a tree. Composable resources are best represented using this tree metaphor.

An E-iRODS composite resource is a tree with one 'root' node. Nodes that are at the bottom of the tree are 'leaf' nodes. Nodes that are not leaf nodes are 'branch' nodes and have one or more 'child' nodes. A child node can have one and only one 'parent' node.

The terms root, leaf, branch, child, and parent represent locations and relationships within the structure of a particular tree. To represent the functionality of a particular resources within a particular tree, the terms 'coordinating' and 'storage' are used in E-iRODS. Coordinating resources coordinate the flow of data to and from other resources. Storage resources are typically 'leaf' nodes and handle the direct reading and writing of data through a POSIX-like interface.

Any resource node can be a coordinating resource and/or a storage resource. However, for clarity and reuse, it is generally best practice to separate the two so that a particular resource node is either a coordinating resource or a storage resource.

This powerful tree metaphor is best illustrated with an actual example. You can now use `ilsresc --tree` to visualize the tree structure of a grid.

```

eirods@hostname:~/ $ ilsresc --tree
demoResc
randy:random
├── pt1:passthru
│   └── ufs5
├── repl1:replication
│   ├── pt2:passthru
│   │   └── pt3:passthru
│   │       └── pt4:passthru
│   │           └── ufs9
│   ├── ufs10
│   └── ufs11
└── ufs1
robin:roundrobin
├── repl2:replication
│   └── repl3:replication
│       ├── ufs6
│       ├── ufs7
│       └── ufs8
├── ufs3
└── ufs4
ufs2
test
test1
test2
test3

```

14.2 Virtualization

In iRODS, files are stored as Data Objects on disk and have an associated physical path as well as a virtual path within the iRODS file system. iRODS collections, however, only exist in the iCAT database and do not have an associated physical path (allowing them to exist across all resources, virtually).

Composable resources, both coordinating and storage, introduce the same dichotomy between the virtual and physical. A coordinating resource has built-in logic that defines how it determines, or coordinates, the flow of data to and from its children. Coordinating resources exist solely in the iCAT and exist virtually across all E-iRODS servers in a particular Zone. A storage resource has a Vault (physical) path and knows how to speak to a specific type of storage medium (disk, tape, etc.). The encapsulation of resources into a plugin architecture allows E-iRODS to have a consistent interface to all resources, whether they represent coordination or storage.

This virtualization enables the coordinating resources to manage both the placement and the retrieval of Data Objects independent from the types of resources that are connected as children resources. When E-iRODS tries to retrieve data, each child resource will "vote", indicating whether it can provide the requested data. Coordinating resources will then decide which particular storage resource (e.g. physical location) the read should come from. The specific manner of this vote is specific to the logic of the coordinating resource. A coordinating resource may lean toward a particular vote based on the type of optimization it deems best. For instance, a coordinating resource could decide between child votes by opting for the child that will reduce the number of requests made against each storage resource within a particular time frame or opting for the child that reduces latency in expected data retrieval times. We expect a wide variety of useful optimizations to be developed by the community.

An intended side effect of the tree metaphor and the virtualization of coordinating resources is the deprecation of the concept of a resource group. Resource groups in community iRODS could not be put into other resource groups. A specific limiting example is a compound resource that, by definition, was a group and could not be placed into another group. This significantly limited its functionality as a management tool. Groups in E-iRODS now only refer to user groups.

Read more about Composable Resources at <http://eirods.org/release/e-irods-composable-resources/>:

- [Paper \(279kB, PDF\)](#)
- [Slides \(321kB, PDF\)](#)
- [Poster \(6.4MB, PDF\)](#)

14.3 Coordinating Resources

Coordinating resources contain the flow control logic which determines both how its child resources will be allocated copies of data as well as which copy is returned when a Data Object is requested. There are several types of coordinating resources: compound, random, replication, round robin, passthru, and some additional types that are expected in the future. Each is discussed in more detail below.

14.3.1 Compound

The compound resource is a continuation of the legacy compound resource type from Community iRODS.

A compound resource has two and only two children. One must be designated as the 'cache' resource and the other as the 'archive' resource. This designation is made in the "context string" of the `addchildtoresc` command.

An Example:

```
ei rods@hostname:~/ $ iadmin addchildtoresc parentResc newChildResc1 cache
ei rods@hostname:~/ $ iadmin addchildtoresc parentResc newChildResc2 archive
```

Putting files into the compound resource will first create a replica on the cache resource and then create a second replica on the archive resource.

Getting files from the compound resource will behave in a similar way as the Community iRODS version. By default, the replica from the cache resource will always be returned. If the cache resource does not have a copy, then a replica is created on the cache resource before being returned.

This compound resource staging policy can be controlled with the policy key-value pair whose keyword is "compound_resource_cache_refresh_policy" and whose values are either "when_necessary" (default), or "always".

From the example near the bottom of the `core.re` rulebase:

```
# =====
# policy controlling when a dataObject is staged to cache from archive in a compound coordinating resource
# - the default is to stage when cache is not present ("when_necessary")
# =====
# pep_resource_resolve_hierarchy_pre(*OUT){*OUT="compound_resource_cache_refresh_policy=when_necessary";} # default
# pep_resource_resolve_hierarchy_pre(*OUT){*OUT="compound_resource_cache_refresh_policy=always";} #
```

Replicas within a compound resource can be trimmed. There is no rebalance activity defined for a compound resource. When the cache fills up, the administrator will need to take action as they see fit. This may include physically moving files to other resources, commissioning new storage, or marking certain resources "down" in the iCAT.

The "--purgec" option for `iput`, `iget`, and `irepl` is honored and will always purge the first replica (usually with replica number 0) for that Data Object (regardless of whether it is held within this compound resource). This is not an optimal use of the compound resource as the behavior will become somewhat nondeterministic with complex resource compositions.

14.3.2 Random

The random resource provides logic to put a file onto one of its children on a random basis. A random resource can have one or more children.

If the selected target child resource of a put operation is currently marked "down" in the iCAT, the random resource will move on to another random child and try again. Currently, the random resource will try 10 times the number of its children, and if still not succeeding, throw an error.

14.3.3 Replication

The replication resource provides logic to automatically manage replicas to all its children.

[Rebalancing](#) of the replication node is made available via the "rebalance" subcommand of `iadmin`. For the replication resource, all Data Objects on all children will be replicated to all other children. The amount of work done in each iteration as the looping mechanism completes is controlled with the session variable `replication_rebalance_limit`. The default value is set at 500 Data Objects per loop.

Getting files from the replication resource will show a preference for locality. If the client is connected to one of the child resource servers, then that replica of the file will be returned, minimizing network traffic.

14.3.4 Round Robin

The round robin resource provides logic to put a file onto one of its children on a rotating basis. A round robin resource can have one or more children.

If the selected target child resource of a put operation is currently marked "down" in the iCAT, the round robin resource will move onto the next child and try again. If all the children are down, then the round robin resource will throw an error.

14.3.5 Passthru

The passthru resource was originally designed as a testing mechanism to exercise the new composable resource hierarchies.

A passthru resource can have one and only one child.

14.3.6 Expected

A few other coordinating resource types have been brainstormed but are not functional at this time:

- Load Balanced (expected)
- Storage Balanced (%-full) (expected)
- Storage Balanced (bytes) (expected)
- Tiered (expected)

14.4 Storage Resources

Storage resources represent storage interfaces and include the file driver information to talk with different types of storage.

14.4.1 Unix File System

The unix file system storage resource is the default resource type that can communicate with a device through the standard POSIX interface.

14.4.2 Structured File Type (*tar, zip, gzip, bzip*)

The structured file type storage resource is used to interface with files that have a known format. By default these are used "under the covers" and are not expected to be used directly by users (or administrators).

These are used mainly for mounted collections.

14.4.3 Non-Blocking

The non-blocking storage resource behaves exactly like the standard unix file system storage resource except that the "read" and "write" operations do not block (they return immediately while the read and write happen independently).

14.4.4 Mock Archive

The mock archive storage resource was created mainly for testing purposes to emulate the behavior of object stores (e.g. WOS). It creates a hash of the file path as the physical name of the Data Object.

14.4.5 Universal Mass Storage Service

The univMSS storage resource delegates `stage_to_cache` and `sync_to_arch` operations to an external script which is located in the `iRODS/server/bin/cmd` directory. It currently writes to the Vault path of that resource instance, treating it as a unix file system.

When creating a "univmss" resource, the context string provides the location of the Universal MSS script.

Example:

```
eirods@hostname:~$ iadmin mkresc myArchiveResc univmss HOSTNAME:/full/path/to/Vault univMSSInterface.sh
```

14.4.6 Expected

A few other storage resource types have been brainstormed but are not included at this time:

- HPSS (expected)
- S3 (expected)
- WOS (expected)

14.5 Managing Child Resources

There are two new `iadmin` subcommands introduced with this feature.

`addchildtoresc`:

```
eirods@hostname:~$ iadmin h addchildtoresc
  addchildtoresc Parent Child [ContextString] (add child to resource)
Add a child resource to a parent resource. This creates an 'edge'
between two nodes in a resource tree.

Parent is the name of the parent resource.
Child is the name of the child resource.
ContextString is any relevant information that the parent may need in order
to manage the child.
```

`rmchildfromresc`:

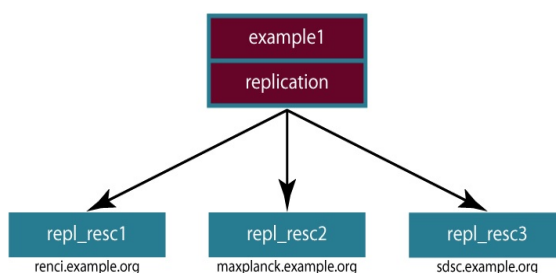
```
eirods@hostname:~$ iadmin h rmchildfromresc
  rmchildfromresc Parent Child (remove child from resource)
Remove a child resource from a parent resource. This removes an 'edge'
between two nodes in a resource tree.
```

Parent is the name of the parent resource.
Child is the name of the child resource.

14.6 Example Usage

Creating a composite resource consists of creating the individual nodes of the desired tree structure and then connecting the parent and children nodes.

14.6.1 Example 1



Example 1: Replicates Data Objects to three locations

A replicating coordinating resource with three unix file system storage resources as children would be composed with seven (7) iadmin commands:

```

ei rods@hostname:~/ $ iadmin mkresc example1 replication
ei rods@hostname:~/ $ iadmin mkresc repl_resc1 "unix file system" renci.example.org:/Vault
ei rods@hostname:~/ $ iadmin mkresc repl_resc2 "unix file system" maxplanck.example.org:/Vault
ei rods@hostname:~/ $ iadmin mkresc repl_resc3 "unix file system" sdsc.example.org:/Vault
ei rods@hostname:~/ $ iadmin addchildtoresc example1 repl_resc1
ei rods@hostname:~/ $ iadmin addchildtoresc example1 repl_resc2
ei rods@hostname:~/ $ iadmin addchildtoresc example1 repl_resc3
  
```

14.7 Rebalancing

A new subcommand for iadmin allows an administrator to rebalance a coordinating resource. The coordinating resource can be the root of a tree, or anywhere in the middle of a tree. The rebalance operation will rebalance for all decedents. For example, the iadmin command `iadmin modresc myReplResc rebalance` would fire the rebalance operation for the replication resource instance named myReplResc. Any Data Objects on myReplResc that did not exist on all its children would be replicated as expected.

For other coordinating resource types, rebalance can be defined as appropriate. For coordinating resources with no concept of "balanced", the rebalance operation is a "no op" and performs no work.

15 Pluggable Authentication

The authentication methods are now contained in plugins. By default, similar to iRODS 3.3 and prior, E-iRODS comes with native iRODS challenge/response (password) enabled. However, enabling an additional authentication mechanism is as simple as adding a file to the proper directory. The server does not need to be restarted.

Available authentication mechanisms include:

- Native iRODS password

- OSAuth
- GSI (Grid Security Infrastructure)
- PAM (Pluggable Authentication Module)
- Kerberos
- LDAP (via PAM)

16 Pluggable Network

E-iRODS now ships with both TCP and SSL network plugins enabled. The SSL mechanism is provided via OpenSSL and wraps the activity from the TCP plugin.

The SSL parameters are tunable via the following `.irodsEnv` variables:

```
# Enable Advanced Client-Server negotiation:
irodsClientServerNegotiation 'request_server_negotiation'
# Client-Server connection policy:
irodsClientServerPolicy 'CS_NEG_REFUSE'
# Client-Server Encryption Key Size In Bytes:
irodsEncryptionKeySize '32'
# Client-Server Encryption Salt Size In Bytes:
irodsEncryptionSaltSize '8'
# Client-Server Encryption Number of Hash Rounds:
irodsEncryptionNumHashRounds '16'
# Client-Server Encryption Algorithm:
irodsEncryptionAlgorithm 'AES-256-CBC'
```

The only valid value for `irodsClientServerNegotiation` at this time is `'request_server_negotiation'`. Anything else will not begin the negotiation stage and default to using a TCP connection.

The possible values for `irodsClientServerPolicy` include:

- `CS_NEG_REQUIRE`: This side of the connection requires an SSL connection
- `CS_NEG_DONT_CARE`: This side of the connection will connect either with or without SSL
- `CS_NEG_REFUSE`: (default) This side of the connection refuses to connect via SSL

In order for a connection to be made, the client and server have to agree on the type of connection they will share. When both sides choose `CS_NEG_DONT_CARE`, E-iRODS shows an affinity for security by connecting via SSL.

The remaining parameters are standard SSL parameters and made available through the EVP library included with OpenSSL. You can read more about these remaining parameters at <https://www.openssl.org/docs/crypto/evp.html>.

17 Users & Permissions

Users and permissions in E-iRODS are inspired by, but slightly different from, traditional UNIX filesystem permissions. Access to Data Objects and Collections can be modified using the `ichmod` iCommand.

Additionally, permissions can be managed via user groups in E-iRODS. A user can belong to more than one group at a time. The owner of a Data Object has full control of the file and can grant and remove access to other users and groups. The owner of a Data Object can also give ownership rights to other users, who in turn can grant or revoke access to users.

Inheritance is a collection-specific setting that determines the permission settings for new Data Objects and sub-Collections. Data Objects created within Collections with Inheritance set to Disabled do not inherit the parent Collection's permissions. By default, E-iRODS has Inheritance set to Disabled. More can be read from the help provided by `ichmod h`.

Inheritance is especially useful when working with shared projects such as a public Collection to which all users should have read access. With Inheritance set to Enabled, any sub-Collections created under the public Collection will inherit the properties of the public Collection. Therefore, a user with read access to the public Collection will also have read access to all Data Objects and Collections created in the public Collection.

18 Rule Engine

The Rule Engine, which keeps track of state and interprets both system-defined rules and user-defined rules, is a critical component of the E-iRODS system. Rules are definitions of actions that are to be performed by the server. These actions are defined in terms of microservices and other actions. The E-iRODS built-in Rule Engine interprets the rules and calls the appropriate microservices.

18.1 File Locking

A race condition occurs when two processes simultaneously try to change the same data. The outcome of a race condition is unpredictable since both threads are "racing" to update the data. To allow E-iRODS users to control such events, the iCommands `iput`, `iget`, and `irepl` each have both `--wlock` and `--rlock` options to lock the Data Objects during these operations. An `irodsServer` thread then purges unused locked files every 2 hours.

18.2 Delay execution

Rules can be run in two modes - immediate execution or delayed execution. Most of the actions and microservices executed by the rule engine are executed immediately, however, some actions are better suited to be placed in a queue and executed later. The actions and microservices which are to be executed in delay mode can be queued with the `delayExec` microservice. Typically, delayed actions and microservices are resource-heavy, time-intensive processes, better suited to being carried out without having the user wait for their completion. These delayed processes can also be used for cleanup and general maintenance of the E-iRODS system, like the `cron` in UNIX.

Monitoring the delayed queue is important once your workflows and maintenance scripts depends on the health of the system. The delayed queue can be managed with the following three iCommands:

1. `iqdel` - remove a delayed rule (owned by you) from the queue.
2. `iqmod` - modify certain values in existing delayed rules (owned by you).
3. `iqstat` - show the queue status of delayed rules.

19 Authentication

By default, E-iRODS uses a secure password system for user authentication. The user passwords are scrambled and stored in the iCAT database. Additionally, E-iRODS supports user authentication via PAM (Pluggable Authentication Modules), which in turn can be configured to support Kerberos and LDAP authentication systems. PAM and SSL have been configured 'on' out of the box with E-iRODS, but there is still some setup required to configure an installation to communicate with your local external authentication server of choice.

19.1 PAM

19.1.1 User Setup

PAM can be configured to support various authentication systems; however the E-iRODS administrator still needs to add the users to the E-iRODS database:

```
eirods@hostname:~/ $ iadmin mkuser newuser rodsuser
```

If the user's credentials will be exclusively authenticated with PAM, a password need not be assigned.

For PAM Authentication, the iRODS user selects the new iRODS PAM authentication choice (instead of password, or Kerberos) via their `.irodsEnv` file or by setting their environment variable:

```
eirods@hostname:~/ $ irodsAuthScheme=PAM
eirods@hostname:~/ $ export irodsAuthScheme
```

Then, the user runs 'iinit' and enters their system password. To protect the system password, SSL (via OpenSSL) is used to encrypt the 'iinit' session.

In order to use the E-iRODS PAM support, you also need to have SSL working between the iRODS client and server. The SSL communication between client and iRODS server needs some basic setup in order to function properly. Much of the setup concerns getting a proper X.509 certificate setup on the server side, and setting up the trust for the server certificate on the client side. You can use either a self-signed certificate (best for testing) or a certificate from a trusted CA.

19.1.2 Server Configuration

The following keywords are used to set values for PAM server configuration. These were previously defined as compile-time options. They are now configurable via the `iRODS/server/config/server.config` configuration file. The default values have been preserved.

- pam_password_length
- pam_no_extend
- pam_password_min_time
- pam_password_max_time

Descriptions of these values can be found in `iRODS/server/icat/src/icatHighLevelRoutines.c`.

19.1.3 Server SSL Setup

Here are the basic steps to configure the server:

19.1.3.1 Generate a new RSA key

Make sure it does not have a passphrase (i.e. do not use the `-des`, `-des3` or `-idea` options to `genrsa`):

```
eirods@hostname:~/ $ openssl genrsa -out server.key
```

19.1.3.2 Acquire a certificate for the server

The certificate can be either from a trusted CA (internal or external), or can be self-signed (common for development and testing). To request a certificate from a CA, create your certificate signing request, and then follow the instructions given by the CA. When running the 'openssl req' command, some questions will be asked about what to put in the certificate. The locality fields do not really matter from the point of view of verification, but you probably want to try to be accurate. What is important, especially since this is a certificate for a server host, is to make sure to use the FQDN of the server as the "common name" for the certificate (should be the same name that clients use as their irodsHost), and do not add an email address. If you are working with a CA, you can also put host aliases that users might use to access the host in the 'subjectAltName' X.509 extension field if the CA offers this capability.

To generate a Certificate Signing Request that can be sent to a CA, run the 'openssl req' command using the previously generated key:

```
eirods@hostname:~/ $ openssl req -new -key server.key -out server.csr
```

To generate a self-signed certificate, also run 'openssl req', but with slightly different parameters. In the openssl command, you can put as many days as you wish:

```
eirods@hostname:~/ $ openssl req -new -x509 -key server.key -out server.crt -days 365
```

19.1.3.3 Create the certificate chain file

If you are using a self-signed certificate, the chain file is just the same as the file with the certificate (server.crt). If you have received a certificate from a CA, this file contains all the certificates that together can be used to verify the certificate, from the host certificate through the chain of intermediate CAs to the ultimate root CA.

An example best illustrates how to create this file. A certificate for a host 'eirods.example.org' is requested from the proper domain registrar. Three files are received from the CA: eirods.crt, PositiveSSLCA2.crt and AddTrustExternalCARoot.crt. The certificates have the following 'subjects' and 'issuers':

```
openssl x509 -noout -subject -issuer -in eirods.crt
subject= /OU=Domain Control Validated/OU=PositiveSSL/CN=eirods.example.org
issuer= /C=GB/ST=Greater Manchester/L=Salford/O=COMODO CA Limited/CN=PositiveSSL CA 2
openssl x509 -noout -subject -issuer -in PositiveSSLCA2.crt
subject= /C=GB/ST=Greater Manchester/L=Salford/O=COMODO CA Limited/CN=PositiveSSL CA 2
issuer= /C=SE/O=AddTrust AB/OU=AddTrust External TTP Network/CN=AddTrust External CA Root
openssl x509 -noout -subject -issuer -in AddTrustExternalCARoot.crt
subject= /C=SE/O=AddTrust AB/OU=AddTrust External TTP Network/CN=AddTrust External CA Root
issuer= /C=SE/O=AddTrust AB/OU=AddTrust External TTP Network/CN=AddTrust External CA Root
```

The eirods.example.org cert was signed by the PositiveSSL CA 2, and that the PositiveSSL CA 2 cert was signed by the AddTrust External CA Root, and that the AddTrust External CA Root cert was self-signed, indicating that it is the root CA (and the end of the chain).

To create the chain file for eirods.example.org:

```
eirods@hostname:~/ $ cat eirods.crt PositiveSSLCA2.crt AddTrustExternalCARoot.crt > chain.pem
```

19.1.3.4 Generate OpenSSL parameters

Generate some Diffie-Hellman parameters for OpenSSL:

```
eirods@hostname:~/ $ openssl dhparam -2 -out dhparams.pem 2048
```


19.1.3.5 Place files within accessible area

Put the `dhparams.pem`, `server.key` and `chain.pem` files somewhere that the E-iRODS server can access them (e.g. in `iRODS/server/config`). Make sure that the `eirops` unix user can read the files (although you also want to make sure that the key file is only readable by the `eirops` user).

19.1.3.6 Set SSL environment variables

The server needs to read these variables on startup:

```
eirops@hostname:~/ $ irodsSSLCertificateChainFile=/var/lib/eirops/iRODS/server/config/chain.pem
eirops@hostname:~/ $ export irodsSSLCertificateChainFile
eirops@hostname:~/ $ irodsSSLCertificateKeyFile=/var/lib/eirops/iRODS/server/config/server.key
eirops@hostname:~/ $ export irodsSSLCertificateKeyFile
eirops@hostname:~/ $ irodsSSLDHParamsFile=/var/lib/eirops/iRODS/server/config/dhparams.pem
eirops@hostname:~/ $ export irodsSSLDHParamsFile
```

19.1.3.7 Restart E-iRODS

Restart the server:

```
eirops@hostname:~/ $ ./iRODS/irodsctl restart
```

19.1.4 Client SSL Setup

The client may or may not require configuration at the SSL level, but there are a few parameters that can be set via environment variables to customize the client SSL interaction if necessary. In many cases, if the server's certificate comes from a common CA, your system might already be configured to accept certificates from that CA, and you will not have to adjust the client configuration at all. For example, on an Ubuntu Precise system, the `/etc/ssl/certs` directory is used as a repository for system trusted certificates installed via an Ubuntu package. Many of the commercial certificate vendors such as VeriSign and AddTrust have their certificates already installed.

After setting up SSL on the server side, test SSL by using the PAM authentication (which requires an SSL connection) and running `iinit` with the log level set to `LOG_NOTICE`. If you see messages as follows, you need to set up trust for the server's certificate, or you need to turn off server verification.

Error from non-trusted self-signed certificate:

```
eirops@hostname:~/ $ irodsLogLevel=LOG_NOTICE iinit
NOTICE: environment variable set, irodsLogLevel(input)=LOG_NOTICE, value=5
NOTICE: created irodsHome=/dn/home/eirops
NOTICE: created irodsCwd=/dn/home/eirops
Enter your current PAM (system) password:
NOTICE: sslVerifyCallback: problem with certificate at depth: 0
NOTICE: sslVerifyCallback: issuer = /C=US/ST=North Carolina/L=Chapel Hill/O=RENCI/CN=eirops.example.org
NOTICE: sslVerifyCallback: subject = /C=US/ST=North Carolina/L=Chapel Hill/O=RENCI/CN=eirops.example.org
NOTICE: sslVerifyCallback: err 18:self signed certificate
ERROR: sslStart: error in SSL_connect. SSL error: error:14090086:SSL routines:SSL3_GET_SERVER_CERTIFICATE:certificate verify failed
sslStart failed with error -2103000 SSL_HANDSHAKE_ERROR
```

Error from untrusted CA that signed the server certificate:

```
eirops@hostname:~/ $ irodsLogLevel=LOG_NOTICE iinit
NOTICE: environment variable set, irodsLogLevel(input)=LOG_NOTICE, value=5
NOTICE: created irodsHome=/dn/home/eirops
NOTICE: created irodsCwd=/dn/home/eirops
Enter your current PAM (system) password:
NOTICE: sslVerifyCallback: problem with certificate at depth: 1
NOTICE: sslVerifyCallback: issuer = /C=US/ST=North Carolina/O=example.org/CN=eirops.example.org Certificate Authority
NOTICE: sslVerifyCallback: subject = /C=US/ST=North Carolina/O=example.org/CN=eirops.example.org Certificate Authority
NOTICE: sslVerifyCallback: err 19:self signed certificate in certificate chain
ERROR: sslStart: error in SSL_connect. SSL error: error:14090086:SSL routines:SSL3_GET_SERVER_CERTIFICATE:certificate verify failed
sslStart failed with error -2103000 SSL_HANDSHAKE_ERROR
```

Server verification can be turned off using the `irodsSSLVerifyServer` environment variable. If this variable is set to 'none', then any certificate (or none) is accepted by the client. This means that your connection

will be encrypted, but you cannot be sure to what server (i.e. there is no server authentication). For that reason, this mode is discouraged.

It is much better to set up trust for the server's certificate, even if it is a self-signed certificate. The easiest way is to use the `irodsSSLCACertificateFile` environment variable to contain all the certificates of either hosts or CAs that you trust. If you configured the server as described above, you could just set the following in your environment:

```
eirodu@hostname:~/ $ irodsSSLCACertificateFile=/var/lib/eirodu/iRODS/server/config/chain.pem
eirodu@hostname:~/ $ export irodsSSLCACertificateFile
```

Or this file could just contain the root CA certificate for a CA-signed server certificate. Another potential issue is that the server certificate does not contain the proper FQDN (in either the Common Name field or the subjectAltName field) to match the client's `irodsHost` variable. If this situation cannot be corrected on the server side, the client can set:

```
eirodu@hostname:~/ $ irodsSSLVerifyServer=cert
eirodu@hostname:~/ $ export irodsSSLVerifyServer
```

Then, the client library will only require certificate validation, but will not check that the hostname of the E-iRODS server matches the hostname(s) embedded within the certificate.

19.1.5 Environment Variables

All the environment variables used by the SSL support (both server and client side) are listed below:

irodsSSLCertificateChainFile (server)

The file containing the server's certificate chain. The certificates must be in PEM format and must be sorted starting with the subject's certificate (actual client or server certificate), followed by intermediate CA certificates if applicable, and ending at the highest level (root) CA.

irodsSSLCertificateKeyFile (server)

Private key corresponding to the server's certificate in the certificate chain file.

irodsSSLDHParamsFile (server)

The Diffie-Hellman parameter file location.

irodsSSLVerifyServer (client)

What level of server certificate based authentication to perform. 'none' means not to perform any authentication at all. 'cert' means to verify the certificate validity (i.e. that it was signed by a trusted CA). 'hostname' means to validate the certificate and to verify that the `irodsHost`'s FQDN matches either the common name or one of the subjectAltNames of the certificate. 'hostname' is the default setting.

irodsSSLCACertificateFile (client)

Location of a file of trusted CA certificates in PEM format. Note that the certificates in this file are used in conjunction with the system default trusted certificates.

irodsSSLCACertificatePath (client)

Location of a directory containing CA certificates in PEM format. The files each contain one CA certificate. The files are looked up by the CA subject name hash value, which must be available. If more than one CA certificate with the same name hash value exist, the extension must be different (e.g. 9d66eef0.0, 9d66eef0.1, etc.). The search is performed based on the ordering of the extension number, regardless of other properties of the certificates. Use the `'c_rehash'` utility to create the necessary links.

20 Configuration

There are a number of configuration files that control how an iRODS server behaves. The following is a listing of the configuration files in an E-iRODS installation.

This document is intended to explain how the various configuration files are connected, what their parameters are, and when to use them.

~/odbc.ini

This file, in the `irods` user's home directory, defines the `unixODBC` connection details needed for the `iCommands` to communicate with the `iCAT` database. This file was created by the installer package and probably should not be changed by the `sysadmin` unless they know what they are doing.

iRODS/config/irods.config

This file defines the main settings for the iRODS installation. It is created by the installer package and comes preconfigured with approved and tested settings. Changing this file will take effect after a restart of the iRODS server. It is recommended not to change this file.

iRODS/server/config/server.config

This file defines the behavior of the server Agent that answers individual requests coming into iRODS. It is recommended not to change this file.

~/irods/irodsA

This is the scrambled password file that is saved after an `init` is run. If this file does not exist, then each `iCommand` will prompt for a password before authenticating with the iRODS server. If this file does exist, then each `iCommand` will read this file and use the contents as a cached password token and skip the password prompt. This file can be deleted manually or can be removed by running `iexit full`.

~/irods/irodsEnv

This is the main iRODS configuration file defining the iRODS environment. Any changes are effective immediately since `iCommands` reload their environment on every execution.

21 Troubleshooting

21.1 Common Errors

Some of the commonly encountered E-iRODS errors along with troubleshooting steps are discussed below.

21.1.1 E-iRODS Server is down

Error Code: `USER_SOCKET_CONNECT_TIMEDOUT -347000`

Common areas to check for this error include:

ienv

- The `ienv` command displays the E-iRODS environment in use. This may help debug the obvious error of trying to connect to the wrong machine or to the wrong Zone name.

Networking issues

- Verify that a firewall is not blocking the connection on the E-iRODS port in use (default 1247).
- Check for network connectivity problems by pinging the server in question.

E-iRODS server logs

If the E-iRODS environment issues and networking issues have been ruled out, the E-iRODS server/client logs may provide additional information with regards to the specifics of the error in question.

21.1.2 *No such file or directory*

Common areas to check for this error include:

1. Permissions - Verify that the E-iRODS user has 'write' access to the directory in question
2. FUSE error
3. Zero byte files

21.1.3 *No rows found in the E-iRODS Catalog*

Error Code: CAT_NO_ROWS_FOUND -808000

This error occurs when there are no results for the database query that was executed. This usually happens when either:

1. the query itself is not well-formed (e.g. syntax error), or
2. the well-formed query produced no actual results (i.e. there is no data corresponding to the specified criteria).

21.1.4 *Access Control and Permissions*

Error Code: CAT_NO_ACCESS_PERMISSION -818000

This error can occur when an iRODS user tries to access an iRODS Data Object or Collection that belongs to another iRODS user without the owner having granted the appropriate permission (usually simply read or write).

With the more restrictive "StrictACL" policy being turned "on" by default in E-iRODS, this may occur more often than expected with Community iRODS. Check the permissions carefully and use `ils -AL` to help diagnose what permissions are set for the Data Objects and Collections of interest.

Modifying the "StrictACL" setting in the iRODS server's `core.re` file will apply the policy permanently; applying the policy via `irule` will have an impact only during the execution of that particular rule.

21.1.5 *Credentials*

Error Code: CAT_INVALID_USER -827000

This error can occur when the iRODS user is unknown or invalid in some way (for instance, no password has been defined for the user, or the user does not exist in that Zone). This error is most common while debugging configuration issues with Zone Federation.

22 Glossary

This glossary attempts to cover most of the terms you may encounter when first interacting with iRODS. More information can be found on the iRODS wiki at <http://irods.org>.

Action

An external (logical) name given to an iRODS Rule(s) that defines a set of macro-level tasks. These tasks are performed by a chain of microservices in accordance with external input parameters. This is analogous to head atom in a Prolog rule or trigger-name in a relational database.

Agent

A type of iRODS server process. Each time a client connects to a server, an agent is created and a network connection established between it and the client.

API

An Application Programming Interface (API) is a piece of software's set of defined programmatic interfaces to enable other software to communicate with it. iRODS defines a client API and expects that clients connect and communicate with iRODS servers in this controlled manner. iRODS has an API written in C, and another written in Java (Jargon).

Authentication Mechanisms

iRODS can employ various mechanisms to verify user identity and control access to Data Objects (iRODS files), Collections, etc. These currently include the default iRODS secure password mechanism (challenge-response), Grid Security Infrastructure (GSI), and Operating System authentication (OSAuth).

Audit Trail

List of all operations performed upon a Data Object, a Collection, a Resource, a User, or other iRODS entities. When Auditing is enabled, significant events in the iRODS system (affecting the iCAT) are recorded. Full activity reports can be compiled to verify important preservation and/or security policies have been enforced.

Client

A Client in the iRODS client-server architecture gives users an interface to manipulate Data Objects and other iRODS entities that may be stored on remote iRODS servers. iRODS clients include: iCommands unix-like command line interface, iDrop (ftp-like client java application), iDropWeb (web interface), etc.

Collection

All Data Objects stored in an iRODS system are stored in some Collection, which is a logical name for that set of Data Objects. A Collection can have sub-collections, and hence provides a hierarchical structure. An iRODS Collection is like a directory in a Unix file system (or Folder in Windows), but is not limited to a single device or partition. A Collection is logical so that the Data Objects can span separate and heterogeneous storage devices (i.e. is infrastructure and administrative domain independent). Each Data Object in a Collection must have a unique name in that Collection.

Data Grid

A grid computing system (a set of distributed, cooperating computers) that deals with the controlled sharing and management of large amounts of distributed data.

Data Object

A Data Object is a single "stream-of-bytes" entity that can be uniquely identified; a file stored in iRODS. It is given a Unique Internal Identifier in iRODS (allowing a global name space), and is associated with (situated in) a Collection.

Driver

A piece of software that interfaces to a particular type of resource as part of the iRODS server/agent process. The driver provides a common set of functions (open, read, write, close, etc.) which allow iRODS clients (iCommands and other programs using the client API) to access different devices via the common iRODS protocol.

Federation

Zone Federation occurs when two or more independent iRODS Zones are registered with one another. Users from one Zone can authenticate through their home iRODS server and have access rights on a remote Zone and its Data Objects, Collections, and Metadata.

Jargon

The Java API for iRODS. Read more at <https://www.irods.org/index.php/Jargon>.

iCAT

The iCAT, or iRODS Metadata Catalog, stores descriptive state metadata about the Data Objects in iRODS Collections in a DBMS database (e.g. PostgreSQL, MySQL, Oracle). The iCAT can keep track of both system-level metadata and user-defined metadata. There is one iCAT database per iRODS Zone.

IES (iCAT-Enabled Server)

A machine that runs both an iRODS server and the iCAT database for a particular Zone.

iCommands

iCommands are Unix utilities that give users a command-line interface to operate on data in the iRODS system. There are commands related to the logical hierarchical filesystem, metadata, data object information, administration, rules, and the rule engine. iCommands provide the most comprehensive set of client-side standard iRODS manipulation functions.

Inheritance

Collections in the iRODS logical name space have an attribute named Inheritance. When Collections have this attribute set to Enabled, new Data Objects and Collections added to the Collection inherit the access permissions (ACLs) of the Collection. Data Objects created within Collections with Inheritance set to Disabled do not inherit the parent Collection's ACL settings. `ichmod` can be used to manipulate this attribute on a per-Collection level. `ils -A` displays ACLs and the inheritance status of the current working iRODS directory.

Logical Name

The identifier used by iRODS to uniquely name a Data Object, Collection, Resource, or User. These identifiers enable global namespaces that are capable of spanning distributed storage and multiple administrative domains for shared Collections or a unified virtual Collection.

Management Policies

The specification of the controls on procedures applied to Data Objects in a Collection. Management policies may define that certain Metadata be required to be stored. Those policies could be implemented via a set of iRODS Rules that generate and verify the required Metadata. Audit Trails could be used to generate reports that show that Management Policies have been followed.

Metadata

Metadata is data about data. In iRODS, metadata can include system or user-defined attributes associated with a Data-Object, Collection, Resource, etc., stored in the iCAT database. The metadata stored in the iCAT database are in the form of AVUs (attribute-value-unit tuples).

Metadata Harvesting

The process of extraction of existing Metadata from a remote information resource and subsequent addition to the iRODS iCAT. The harvested Metadata could be related to certain Data Objects, Collections, or any other iRODS entity.

Micro-service

A set of operations performed on a Collection at a remote storage location.

Micro-services are small, well-defined procedures/functions that perform a certain server-side task and are compiled into the iRODS server code. Rules invoke Micro-services to implement Management Policies. Micro-services can be chained to implement larger macro-level functionality, called an Action. By having more than one chain of Micro-services for an Action, a system can have multiple ways of performing the Action. At runtime, using priorities and validation conditions, the system chooses the "best" micro-service chain to be executed.

Migration

The process of moving digital Collections to new hardware and/or software as technology evolves. Separately, Transformative Migration may be used to mean the process of manipulating a Data Object into a new format (e.g. gif to png) for preservation purposes.

Physical Resource

A storage system onto which Data Objects may be deposited. iRODS supports a wide range of disk, tape, and remote storage resources.

Resource

A resource, or storage resource, is a software/hardware system that stores digital data. iRODS clients can operate on local or remote data stored on different types of resources through a common interface.

Rules

Rules are a major innovation in iRODS that let users automate data management tasks, essential as data collections scale to petabytes across hundreds of millions of files. Rules allow users to automate enforcement of complex Management Policies (workflows), controlling the server-side execution (via Micro-services) of all data access and manipulation operations, with the capability of verifying these operations.

Rule Engine

The Rule Engine interprets Rules following the iRODS rule syntax. The Rule Engine, which runs on all iRODS servers, is invoked by server-side procedure calls and selects, prioritizes, and applies Rules and their corresponding Micro-services. The Rule Engine can apply recovery procedures if a Micro-service or Action fails.

Scalability

Scalability means that a computer system performs well, even when scaled up to very large sizes. In iRODS, this refers to its ability to manage Collections ranging from the data on a single disk to petabytes (millions of gigabytes) of data in hundreds of millions of files distributed across multiple locations and administrative domains.

Server

An iRODS server is software that interacts with the access protocol of a specific storage system. It enables storing and sharing data distributed geographically and across administrative domains.

Transformative Migration

The process of manipulating a Data Object from one encoding format to another. Usually the target format will be newer and more compatible with other systems. Sometimes this process is "lossy" and does not capture all of the information in the original format.

Trust Virtualization

The management of Authentication and authorization independently of the storage location.

Unique Internal Identifier

See Logical Name.

User Name

Unique identifier for each person or entity using iRODS; sometimes combined with the name of the home iRODS Zone (as username#Zonename) to provide a globally unique name when using Zone Federation.

Vault

An iRODS Vault is a data repository system that iRODS can maintain on any storage system which can be accessed by an iRODS server. For example, there can be an iRODS Vault on a Unix file system, an HPSS (High Performance Storage System), or an IBM DB2 database. A Data Object in an iRODS Vault is stored as an iRODS-written object, with access controlled through the iCAT catalog. This is distinct from legacy data objects that can be accessed by iRODS but are still owned by previous owners of the data. For file systems such as Unix and HPSS, a separate directory is used; for databases such as Oracle or DB2 a system-defined table with LOB-space (Large Object space) is used.

Zone

An iRODS Zone is an independent iRODS system consisting of an iCAT-Enabled Server (IES), optional additional distributed iRODS Servers (which can reach hundreds, worldwide) and clients. Each Zone has a unique name. When two iRODS Zones are configured to interoperate with each other securely, it is called (Zone) Federation.

23 Known Issues

```

Ticket Item: [#1260] unixODBC on OpenSuSE 12.x fails when iRODS resource name contains a hyphen (aka "hpss-sdsc")
-----
unixODBC on OpenSuSE 12.x fails when iRODS resource name contains a hyphen (aka "hpss-sdsc").

Also, 'moon landing' in rules3.0/rulewriteKeyValPairs.r.

Other Operating Systems and versions do not exhibit this behavior.

Mar 6 09:47:34 pid:21588 NOTICE: rsAuthCheck user rods#tempZone
Mar 6 09:47:34 pid:21588 NOTICE: rsAuthResponse set proxy authFlag to 5, client authFlag to 5, user:rods#tempZone
proxy:rods client:rods
Mar 6 09:47:34 pid:21588 NOTICE: bindVar[1]=RajaBase
Mar 6 09:47:34 pid:21588 NOTICE: bindVar[2]=acRegisterData
Mar 6 09:47:34 pid:21588 NOTICE: bindVar[3]=acRegisterData()
Mar 6 09:47:34 pid:21588 NOTICE: bindVar[4]=($objPath like "/home/raja#sdsc/myImportantFiles\*" && $dataSize > 1000000)
Mar 6 09:47:34 pid:21588 NOTICE: bindVar[5]={
  msiRegisterData()::recover_msiRegisterData();
  msiQueue("msiReplicateData('\hpss-sdsc') :: recover_msiReplicateData;");
}
Mar 6 09:47:34 pid:21588 NOTICE: bindVar[6]=@REL
Mar 6 09:47:34 pid:21588 NOTICE: cllExecSqlWithResult: SQLExecDirect error: -1, sql:select rule_id from R_RULE_MAIN
where rule_base_name = ? and rule_name = ? and rule_event = ? and rule_condition = ? and rule_body = ? and
rule_recovery = ?
Mar 6 09:47:34 pid:21588 NOTICE: SQLSTATE: 01000
Mar 6 09:47:34 pid:21588 NOTICE: SQLCODE: 4294967295
Mar 6 09:47:34 pid:21588 NOTICE: SQL Error message: [unixODBC]Error while executing the query (non-fatal);
ERROR: syntax error at or near "hpss" at character 344
Mar 6 09:47:34 pid:21588 NOTICE: chlInsRuleTable cmlGetIntegerValueFromSqlV3 find rule if any failure -806000
Mar 6 09:47:34 pid:21588 NOTICE: rsGeneralRowInsert: rcGeneralRowInsert failed
Mar 6 09:47:34 pid:21588 ERROR: executeRuleAction Failed for msiAdmInsertRulesFromStructIntoDB status = -806000
CAT_SQL_ERR

Ticket Item: [#1807] Fix OSAuth testing in CI
-----

Test for OSAuth yielded a rxAuthResponse error.

iexit needed to be called before a new iinit is issued with a new environment.

```

24 History of Releases

Date	Version	Description
2013-11-16	3.0.1	Second Release. This is the second open source release from RENC1. It includes Federation compliance with Community iRODS and signaling for dynamic post-PEPs to know whether their operation failed.
2013-11-14	3.0.1rc1	First Release Candidate of Second Release. This is the first release candidate of the second open source release from RENC1. It includes a new "--tree" view for <i>ilsresc</i> and a more powerful <i>irodsctl stop</i> . In addition, package managers should now be able to handle upgrades more gracefully.
2013-11-12	3.0.1b2	Second Beta of Second Release. This is the second beta of the second open source release from RENC1. It includes certification work with the Jargon library, more CI testing, and minor fixes.

2013-10-31	3.0.1b1	<p>First Beta of Second Release.</p> <p>This is the first beta of the second open source release from RENCI. It includes pluggable network and authentication support as well as a rebalance option and migration support for the composable resources.</p>
2013-06-05	3.0	<p>First Release.</p> <p>This is the first open source release from RENCI. It includes all the features mentioned below and has been both manually and continuously tested.</p>
2013-05-14	3.0rc1	<p>First Release Candidate.</p> <p>This is the first release candidate from RENCI. It includes PAM support, additional resources (compound, universalMSS, replication, random, and nonblocking), and additional documentation.</p>
2013-03-15	3.0b3	<p>Third Beta Release.</p> <p>This is the third release from RENCI. It includes a new package for CentOS 6+, support for composable resources, and additional documentation.</p>
2012-06-25	3.0b2	<p>Second Beta Release.</p> <p>This is the second release from RENCI. It includes packages for iCAT, Resource, iCommands, and development, in both DEB and RPM formats. Also includes more documentation.</p>
2012-03-01	3.0b1	<p>Initial Beta Release.</p> <p>This is the first release from RENCI, based on the iRODS 3.0 community codebase.</p>