

Connecting iRODS to SUSE Enterprise Storage

Architecture and Implementation for S3 and RADOS

Guide

www.suse.com

Implementation Guide



Connecting iRODS to SUSE Enterprise Storage

Architecture and Implementation for S3 and RADOS.

Target Audience

This guide provides step-by-step instructions on how to connect iRODS software to SUSE Enterprise Storage via an S3 Plugin or RADOS Plugin.

Overview of the Solution

The integrated Rule-Oriented Data System (iRODS) is open source data management software used by research organizations and government agencies worldwide. It virtualizes data storage resources so that users can take control of their data, regardless of where and on what device the data is stored. As data volumes grow and data services become more complex, iRODS is serving an increasingly important role in data management.

SUSE Enterprise Storage (SES) provides the backend storage for an iRODS via an S3 plugin. SES is an intelligent software-defined storage solution powered by Ceph technology. It provides highly redundant, automated, self-healing, and unlimited scalability—all by using commodity off-the-shelf hardware.

Prerequisites

- A deployed and running iRODS server. Please see the iRODS server installation guides at <https://docs.irods.org>. Contact iRODS for additional details on deploying the iRODS server.
- A deployed and running SES cluster. Please see the SUSE Enterprise Storage deployment guide: <https://documentation.suse.com/ses/6/single-html/ses-deployment/index.html>.

What Was Tested?

- iRODS tested two plugins for SES connectivity: S3 and RADOS.
- The S3 plugin is fully supported by iRODS. The RADOS plugin was tested and used by many clients as well, but is not yet officially released for support (iRODS is working on this).

S3 Plugin Connectivity

Prerequisites

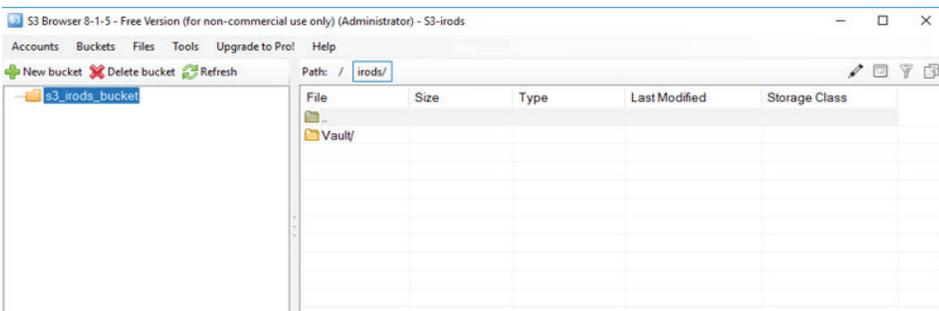
Install the S3 plugin on iRODS. Please follow the iRODS documentation for S3 plugin installation:
https://github.com/irods/irods_resource_plugin_s3.

Connecting to SES Cluster

- Set up DNS and verify that all names are fully resolvable by name and IP.
- Create a new user and a bucket from SES cluster.

The screenshot shows two forms in the iRODS web interface. The first form is titled "Edit user: s3_irods_user" and contains the following fields: "Username" (s3_irods_user), "Full name" (s3_irods_user), "Email address" (Email Address...), and "Max. buckets" (1000). The second form is titled "Edit bucket: s3_irods_bucket" and contains the following fields: "Name" (s3_irods_bucket), "Id" (5fa98368-5b7f-4f17-b427-643b25a7d596.5164107.1), and "Owner" (s3_irods_user). At the bottom right of the bucket form are "Submit" and "Back" buttons.

- Verify that you can connect to your bucket.
 - In the example below, you can use another test Windows VM and the S3 browser to verify the connection. This can be downloaded from: <http://s3browser.com/>.



Configuring the iRODS S3 Plugin

You can find the example configuration steps at: https://github.com/irods/irods_resource_plugin_s3.

In the following commands, we have two environment variables:

- For `S3_DEFAULT_HOSTNAME`, use your S3 gateway host.
- For `S3_AUTH_FILE`, use the full path for the key pair location generated during a bucket creation on the cluster. The AWS/S3 key pair file `S3_AUTH_FILE` should have exactly two values: Access Key ID and Secret Access Key.

Run the following with your environment variables:

```
irods@hostname $ iadmin mkresc compResc compound
```

```
irods@hostname $ iadmin mkresc cacheResc unixfilesystem <hostname>:</full/path/to/Vault>
```

```
irods@hostname $ iadmin mkresc archiveResc s3 <hostname>:</s3BucketName>/irods/Vault "S3_DEFAULT_HOSTNAME=s3.  
amazonaws.com;S3_AUTH_FILE=</full/path/to/AWS.keypair>;S3_RETRY_COUNT=<num reconn tries>;S3_WAIT_TIME_  
SEC=<wait between retries>;S3_PROTO=<HTTP|HTTPS>"
```

```
irods@hostname $ iadmin addchildtoresc compResc cacheResc cache
```

```
irods@hostname $ iadmin addchildtoresc compResc archiveResc archiveirods@hostname $ iput -R compResc foo.txt
```

```
irods@hostname $ ireg -R archiveResc </s3BucketName>/full/path/in/bucket /full/logical/path/to/dataObject
```

Once completed, verify that you can put/pull files from the storage bucket.

Additional Options

S3 Cacheless Mode was also verified for support.

Please follow the iRODS documentation for the proper installation steps:

https://github.com/irods/irods_resource_plugin_s3#using-the-s3-plugin-in-cacheless-mode.

RADOS Plugin Connectivity

Prerequisite

Install the RADOS plugin on iRODS. Please follow the iRODS documentation for RADOS plugin installation instructions: https://github.com/irods/irods_resource_plugin_rados/blob/master/README.md.

SES Cluster Setup Steps (2 options)

You can create a pool for iRODS to the SES cluster via the Open Attic dashboard or via the SES CLI. Below are two setup methods to achieve the same result:

1. *Using the Open Attic dashboard:*

- *Create an iRODS pool on Ceph.*
- *Create a pool from Open Attic.*

Edit Ceph pool: irods

Name	<input type="text" value="irods"/>
Pool type	<input type="text" value="Replicated pool"/>
Crush ruleset	<input type="text" value="replicated_rule"/>
	Size Steps
Replicated size	<input type="text" value="3"/>
Placement groups *	<input type="text" value="128"/>
	<small>The pg size of a pool can only be increased in edit mode.</small>
	Calculation help

Compression

Mode	<input type="text" value="none"/>
------	-----------------------------------

Applications

Add applications	<input type="text" value="- Select an application to use -"/>	<input type="button" value="+"/>
Applications *	<input type="text" value="rgw"/>	<input type="button" value="🗑"/>

2. If you prefer to use the SES CLI terminal on the admin node:

- Create the iRODS OSD pool:

```
SES:~ # ceph osd pool create irods 128
```

- Set permissions:

```
SES:~ # ceph auth get-or-create client.irods osd 'allow rw pool=irods' mon 'allow r' > /etc/ceph/client.irods.keyring
```

- Copy the key from the newly created keyring and create the Ceph config files on the resource iRODS server:

```
SES:~ # touch /etc/irods/irados.config && chown irods: /etc/irods/irados.config && chmod 600 /etc/irods/irados.config
```

3. Once complete, test uploading and downloading files.

Additional contact information and office locations:
www.suse.com

www.suse.com

