

IBM COS iRODS plugin

The New Object Store plugin for iRODS

Anu Khera
Sr. Offering Manager
IBM Cloud Object Storage

Introduction

- IBM Cloud Object Storage (IBM COS) is a massively scalable object storage platform for hybrid cloud environments that provides standard S3.
- IBM are committed to delivering end-to-end software-defined storage solutions around IBM Cloud Object Storage and are working with communities including iRODS.
- IBM are investing into a long term strategy around iRODS that focuses on iRODS interoperability with object storage platforms and the optimization of iRODS for IBM Cloud Object Storage.
- We are pleased to announce the imminent release of the IBM COS iRODS plugin.

The Problem

- iRODS users need a way to interface with object stores – something fast, direct, and optimized for use with stored objects.
- Most iRODS use cases require a random-access workflow and operate on increasingly large (>100GB) objects
- The current iRODS resource plugin for S3 does not provide *native* support for random-access workflows including offset reads and writes
- The iRODS S3 plugin must consequently be deployed as compound resource using a POSIX file system cache which is unfeasible for large object processes
- IBM have researched this problem and have developed the **IBM COS iRODS plugin**

The IBM COS iRODS Plugin

- **The IBM COS iRODS Plugin is a native iRODS resource plugin optimized for IBM Cloud Object Storage**
- Unlike the iRODS S3 plugin, the **IBM COS iRODS plugin** does away with the inefficient cache layer and provides native support for efficient random I/O operations expected of most iRODS applications

Current iRODS S3 workflow

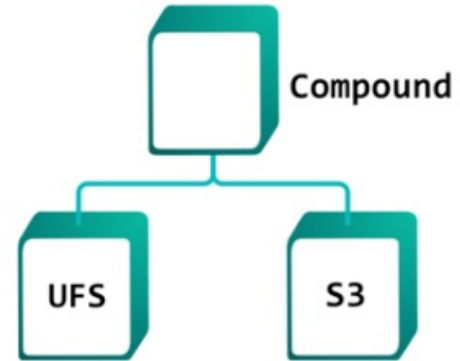
- **iRODS Compound Resources provide a POSIX interface to S3**
 - Pair a Unix File System resource with an S3 resource
 - The UFS resource adds POSIX operations to the immutable S3 archive

- **Writing objects (PUT)**

- Data objects are delegated to the cache first and registered
- Then a copy is sent to the S3 object store
- Offset writes result in a new immutable object

- **Reading objects (GET)**

- Replica is staged to cache if required
- Read always happens from cache
- Offset reads of a few bytes require full object to be read



- **This configuration results in a large and unpredictable file cache**

iRODS IBM COS Workflow

- **iRODS IBM COS Resources provide a native POSIX interface to S3 and IBM COS**
 - No local file system cache is required
 - POSIX operations are supported using a ‘smart plugin’
- **Writing objects (PUT)**
 - Data objects broken into chunks and individual chunks are written to the object store
 - Offset writes result only update the impacted chunks
- **Reading objects (GET)**
 - Only the required chunks are read from the object store
 - Offset reads of a few bytes read only the relevant few chunks
- **This eliminates the need for an inefficient and unpredictable file cache**

Comparing the IBM COS iRODS Plugin

	S3 Plugin	Compound Resource (S3 & UFS)	IBM COS plugin
S3 Compatible	✓	✓	✓
IBM COS Compatible	✗	✗	✓
Open Source	✓	✓	✓
POSIX compliant	✗	✓	✓
No expensive cache layer		✗	✓
Optimized I/O	✗	✗	✓
IBM COS Optimized			✓

iRODS IBM COS Plugin Specs

– Features

- Native iRODS Resource Plugin
- Full POSIX Compliance
- Optimized workflows
- Multiple Authentication Modes

– Supported iRODS versions

- iRODS 4.03 - 4.1.10
- *iRODS 4.2 planned*
- All Operating Systems planned

– Object Stores

- IBM Cloud Object Storage Public
- IBM Cloud Object Storage On Prem
- IBM Cloud Object Storage Hosted
- IBM Cloud Object Storage Hybrid