

dCache and iRODs

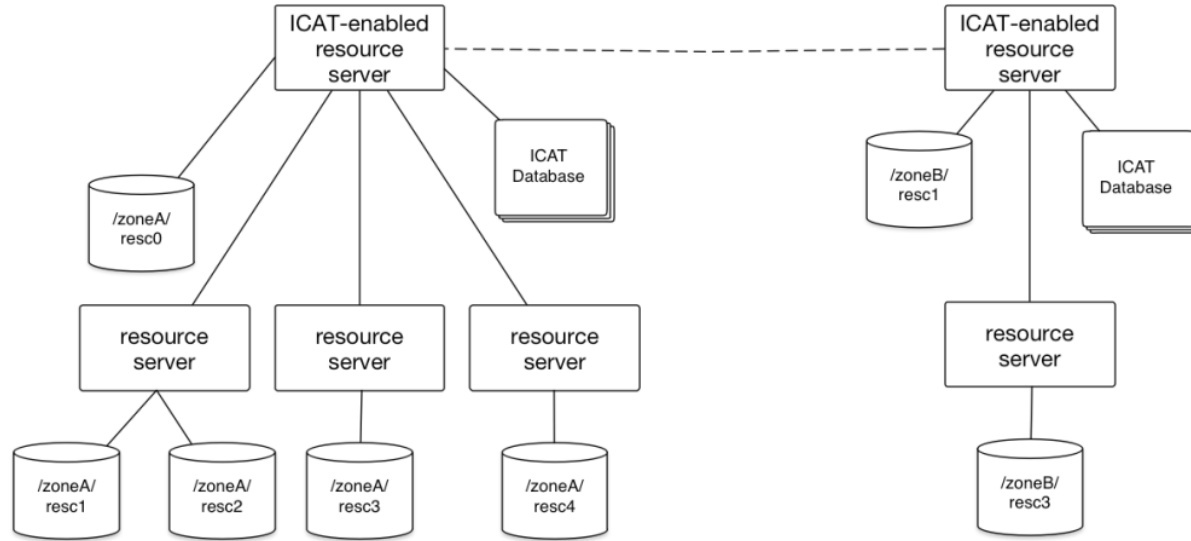


iRODs

- Open source data management middleware
- Metadata
- Actions triggered based on rule engine
- Federation
 - Sharing data
- Storage are “resources”
 - Composable resources
- Clients
 - Commandline (iCommands)
 - API (libraries C++, Python, Golang, PHP,....)
- WebDAV
- More info on irods.org

iRODs

iRODS Grid Topology



iRODs

- IRODs in use in a number of places
 - EUDAT's B2SAFE
 - SweStore
 -
- @SURF
 - Scale-out service
 - Hosting
 - YODA (GUI)

iRODs and dCache

- Make iRODs sit on top of a dCache nfs4.1 mounted filesystem
- IRODs has various plugins
- libunixfilesystem.cpp
- Make it work with WORM storage

iRODs and dCache

- Needed to modify `libunixfilesystem.cpp` at a few places
 - At places where a file is opened for (over)writing, throw it away first
- iCommands seem to work (`imkdir`, `icp`, `irm`, `imv`,...)
- There is one thing though....or actually two

iRODs and dCache

- In `libunixfilesystem.cpp` there is a `unix_file_open_plugin`
- This function is used when files are overwritten
- Opens files with `O_RDWR|O_CREAT|O_TRUNC` when files are overwritten
- That's what I filter for when deciding to throw away a file
- But are `O_RDWR|O_CREAT|O_TRUNC` used in this function for purposes other than overwriting files?
 - Needs to be checked out
 - The `O_RDWR` is what worries me

iRODs and dCache

- Then there is the client API
 - rcDataObjOpen, rcDataObjLseek, rcDataObjWrite
- This will not work on WORM storages

iRODs and dCache

- What needs to be done?
 - Make sure that throwing away existing files first will not create havoc
 - Fix API in some way

iRODs and dCache

- How?
 - Make a clear distinction between opening files for reading and (overwriting)
 - Maybe a separate WORM iRODs plugin?
 - Make the dangerous API calls return a proper error message
 - Put the file that is to be modified in some noWORM scratch space and copy it back when you are done
 - Or.....a NoWORM dCache???

