



PLUGIN_RUCIO.SO (WIP)

Mario Lassnig

European Organisation For Nuclear Research (CERN), Geneva, Switzerland

ph-adp-ddm-lab@cern.ch

mario.lassnig@cern.ch

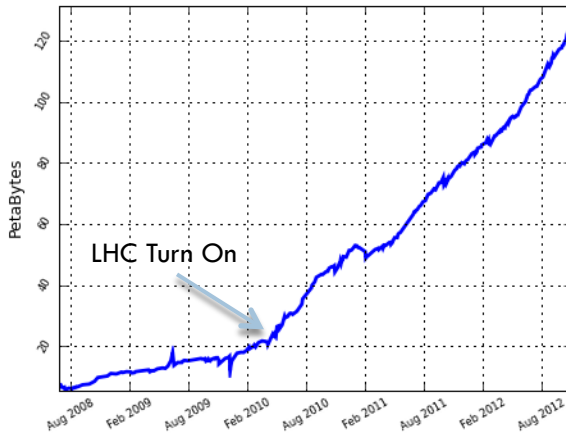
DPM Community Workshop, Paris, 2012

ATLAS Data Management

2

- ATLAS is currently accumulating $\sim 40\text{PB}$ of permanent data per year, and handling $\sim 20\text{PB}$ temporary data
 - ▣ Data is written to files, and files are aggregated logically in datasets
 - ▣ Files and datasets are replicated to data centres globally for distributed processing
 - ▣ Data management system (v1:DQ2, v2:Rucio) is responsible for the data workload
 - Interaction rate of events at $\sim 800\text{Hz}$ 24/7, with the occasional burst to $\sim 1500\text{Hz}$
 - ~ 15 million data management events per day, (e.g., *addFilesToDataset*, *scheduleFileForDeletion*, *download*, ...)
 - “Old” data does not become unused, due to the nature of high-energy physics analysis (reprocessing with new algorithms, ...)

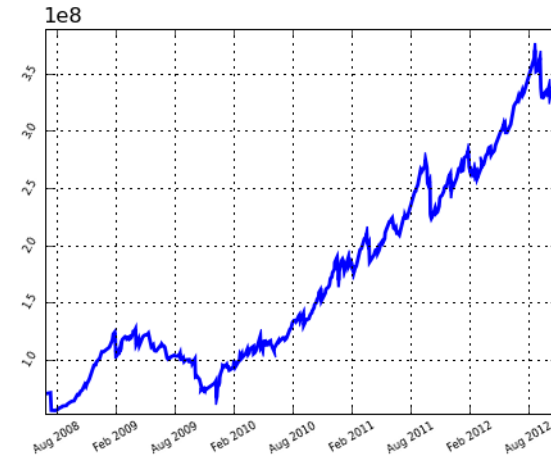
Total GRID space usage according to DQ2



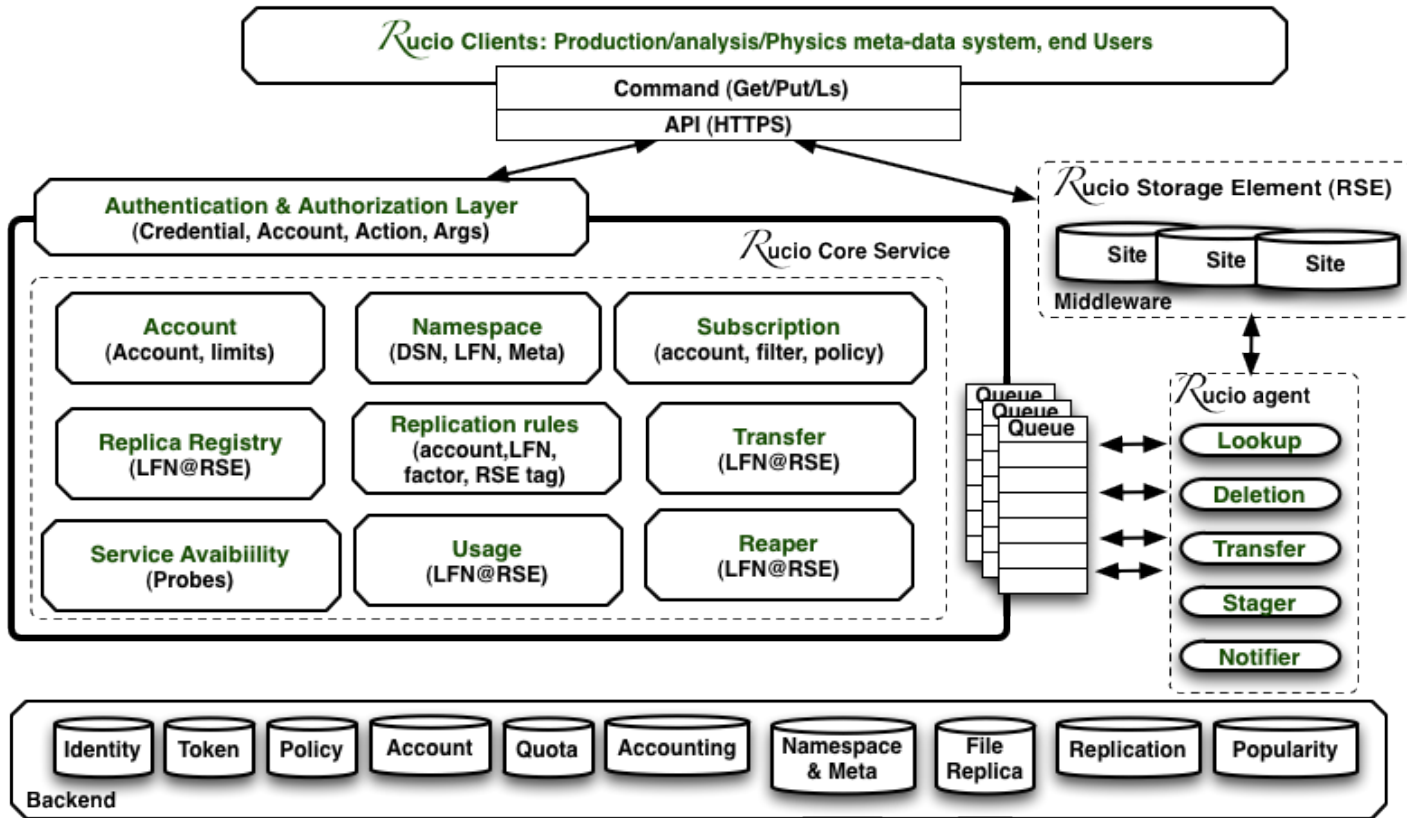
Total GRID datasets according to DQ2



Total GRID files according to DQ2




Complete system overview



Accounts and permissions

4

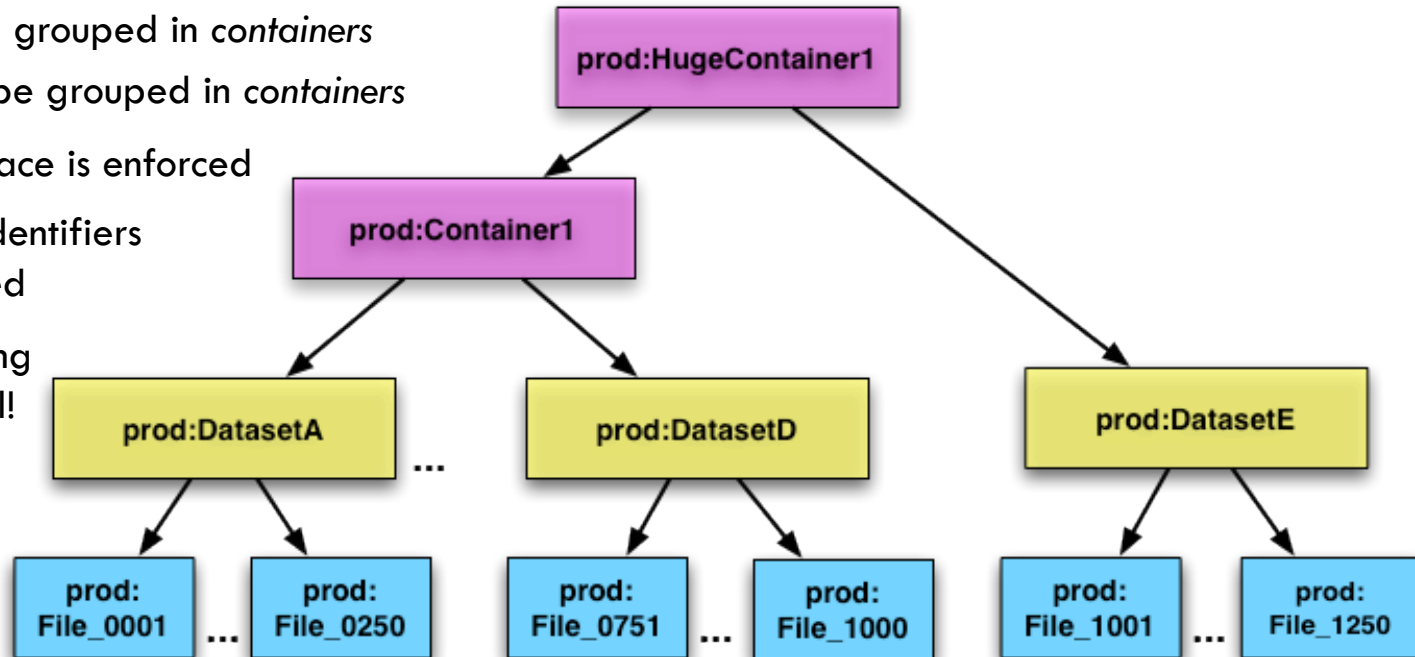
- There are many credentials that can identify a person
 - username and password
 - x509 certificates
 - Kerberos GSSAPI
 - OpenID
 - ...
- In many cases, these credentials are not issued or verified by the same entity
- Traditional groups and ACLs do not map very well to this
 - Many different credentials can map to many different roles
 - Micro-management nightmare
- Let credentials be verified by their issuing entity
- Permissions are handled on accounts, not credentials

 The image cannot be displayed. Your computer may not have enough memory to open the image, or the image may have been corrupted. Restart your computer, and then open the file again. If the red x still appears, you may have to delete the image and then insert it again.

Data identifiers (DID)

5

- Scoped namespace
- *Files must be grouped in datasets*
- *Datasets can be grouped in containers*
- *Containers can be grouped in containers*
- Unique namespace is enforced
- Deleted data identifiers cannot be reused
- Physical grouping on dataset level!



plugin_rucio.so

6

- First iteration
 - ▣ Allow browsing of ATLAS namespace through DMLite
- Second iteration
 - ▣ CRUD

- Main problem
 - ▣ Overlapping datasets and containers in a scoped namespace do not map well to a POSIX-like filesystem hierarchy

user_a: container_b



DID

plugin_rucio.so

7

- First iteration
 - Allow browsing of ATLAS namespace through DMLite
- Second iteration
 - CRUD

- Main problem
 - Overlapping datasets and containers in a scoped namespace do not map well to a POSIX-like filesystem hierarchy

`/user_a/container_b/user_a: container_c`



plugin_rucio.so

8

- First iteration
 - ▣ Allow browsing of ATLAS namespace through DMLite
- Second iteration
 - ▣ CRUD

- Main problem
 - ▣ Overlapping datasets and containers in a scoped namespace do not map well to a POSIX-like filesystem hierarchy

`/user_a/container_b/user_a/container_c/user_b: dataset_c`



plugin_rucio.so

9

- First iteration
 - Allow browsing of ATLAS namespace through DMLite
- Second iteration
 - CRUD
- Main problem
 - Overlapping datasets and containers in a scoped namespace do not map well to a POSIX-like filesystem hierarchy

`/user_a/container_b/user_a/container_c/user_b/dataset_c/user_c: file_1`



plugin_rucio.so

10

- First iteration
 - ▣ Allow browsing of ATLAS namespace through DMLite
- Second iteration
 - ▣ CRUD
- Main problem
 - ▣ Overlapping datasets and containers in a scoped namespace do not map well to a POSIX-like filesystem hierarchy


`/user_a/container_b/user_a/container_c/user_b/dataset_c/user_c/file_1`



plugin_rucio.so

11

- First iteration
 - ▣ Allow browsing of ATLAS namespace through DMLite
- Second iteration
 - ▣ CRUD
- Main problem
 - ▣ Overlapping datasets and containers in a scoped namespace do not map well to a POSIX-like filesystem hierarchy

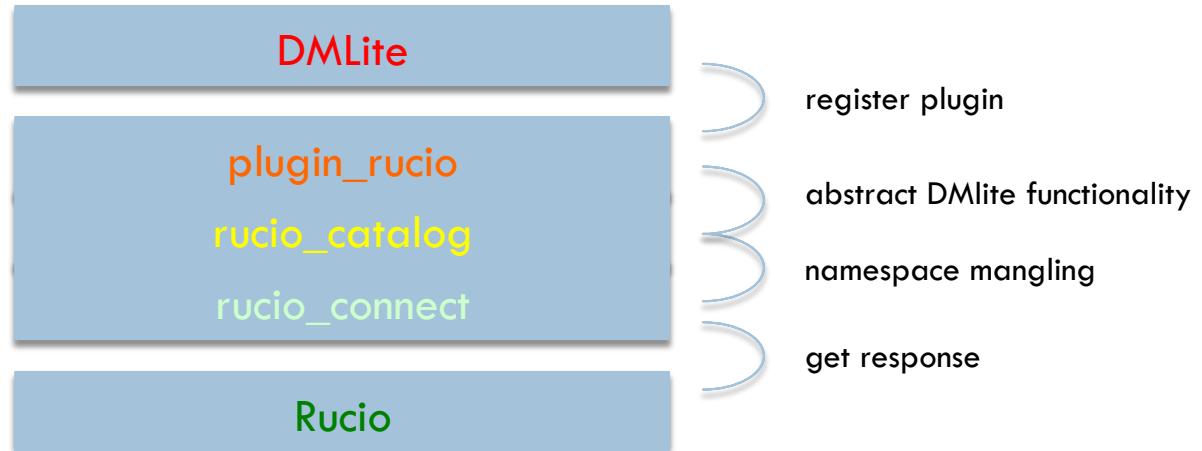


Simple solution

12

- Root becomes a list of all scopes
 - /scope_a, /scope_b, ...
- Listing a path then has to check if the last entry is either a scope or a data identifier
 - query Rucio appropriately and return the proper data

- Rucio API
 - HTTP/REST
 - libcurl
 - JSON
 - jansson
 - full support for JSON only in boost versions 1.42+
 - (SLC6 has 1.41)



DEMONSTRATION!



13

