

DynaFed at RAL

Alastair Dewhurst



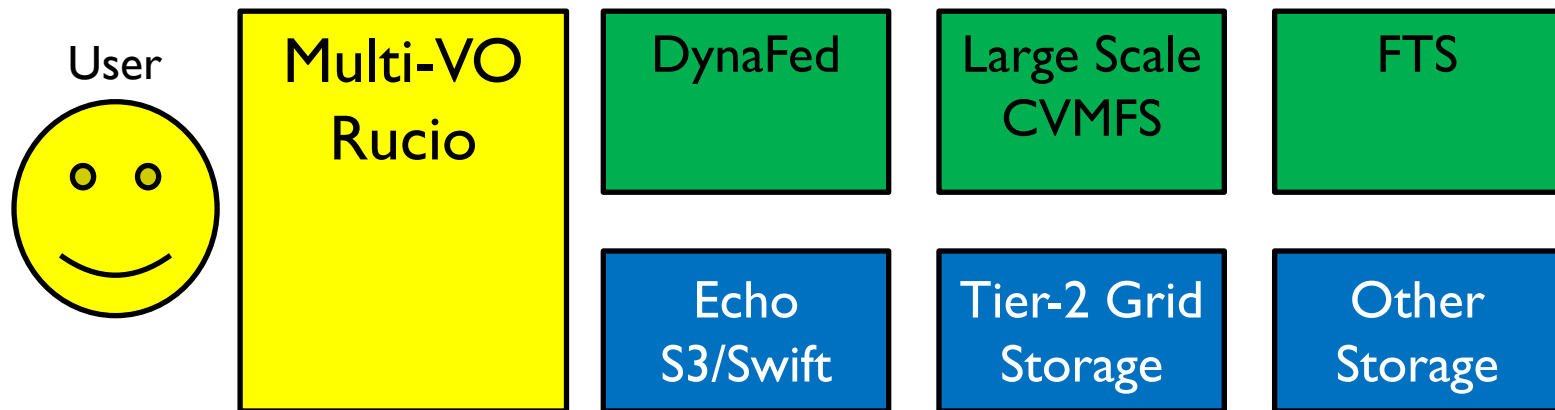
Introduction

- The RAL Tier-I supports all LHC VOs as well as a growing number of other experiments (many non-HEP).
- The RAL Tier-I Echo storage is backed by a very large Ceph cluster.
 - We wrote XRootD/GridFTP plugins.
 - We would prefer to just provide access via the RadosGW S3 / Swift APIs.
 - UK Tier-2s are looking at Ceph for future storage.
- DynaFed is the obvious choice to act as the intermediate layer.
 - For authentication and authorization we use a Grid-mapfile.



DDM as a service

- RAL is trying to provide a coherent set of data management tools to non WLCG experiments.
- We are developing a Multi-VO Rucio instance.
 - Currently being used by AENEAS for testing.
 - DynaFed is being used both by Rucio and individual local users.




AENEAS / SKA

- AENEAS aim is to produce a functional design for a distributed, federated European Science Data Centre for SKA.
 - Will finish at the end of 2019.
 - Trying to secure 6 month PhD student placement to further testing.
- Resources:
 - Rucio and FTS test instance at RAL
 - UK Grid Storage, Echo S3 via DynaFed, GridFTP in SA.

AENEAS transfer matrix.



Streaming through FTS instance
(not DynaFed) if necessary.



Diamond

The Diamond Light Source

- 20+ experiments at any one time.
- Experiments take a few weeks.
- Experiments take ~10-20TB data

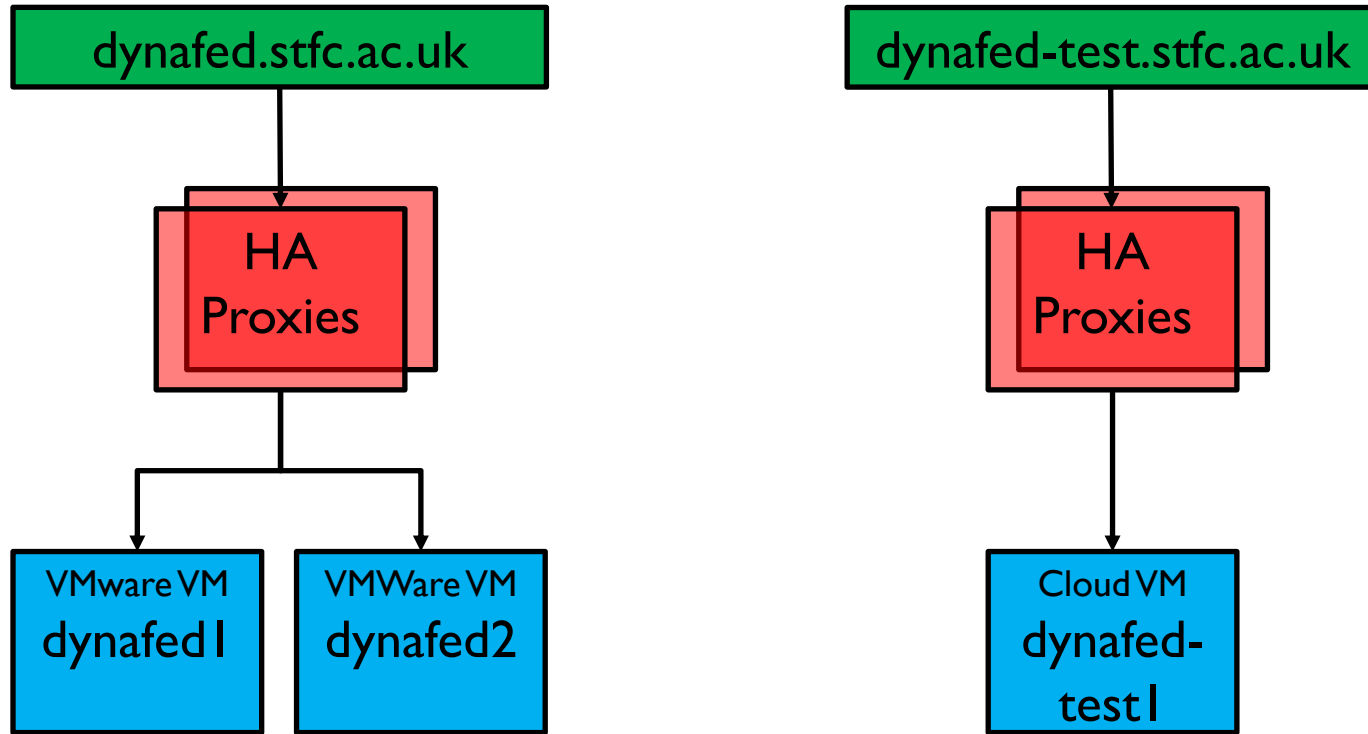
4PB per year data created and growing.



- Diamond have a Lustre file system for users to run analysis locally for first 60 days.
- RAL provides a tape archive for long term data preservation.
- We have a request to provide cloud storage to cover intermediate phase:
 - Allow users to access files once they have returned to their own institute.



RAL Setup



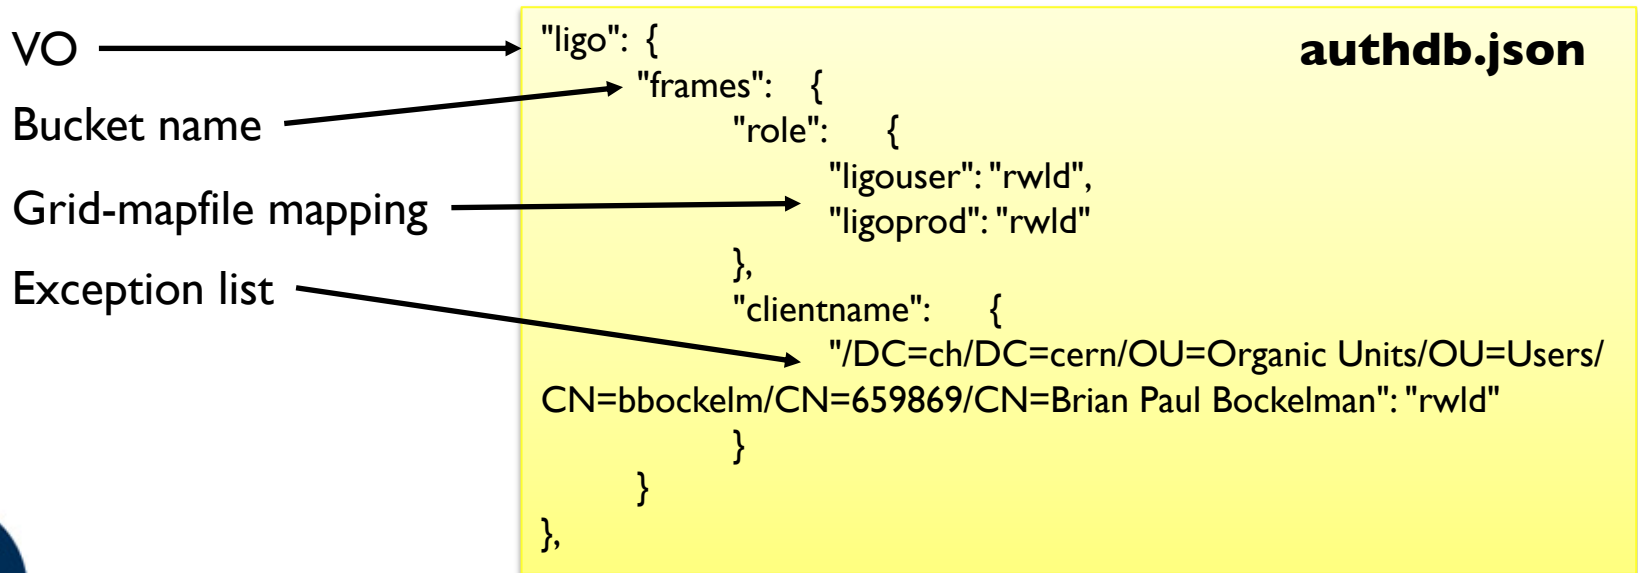
- Using version 1.4 currently.



Grid-mapfile Authorization

7

- Three files needed:
 - `/etc/ugr/conf.d/gridmap_auth.py`
 - `/etc/grid-security/grid-mapfile`
 - `/etc/grid-security/authdb.json`
- Followed similar logic to GridFTP/XRootD authdb.
- Additional feature allows authdb to contain exceptions.
- Requires buckets naming convention: `<VO>-<Name>`



Ldap Plugin

- For local users (e.g. Diamond), they would like access to their data via their STFC username/password.
- Graduate Louise Davies developed ldap plugging for DynaFed:
 - <https://github.com/stfc/dynafed-ldap-plugin>
- DynaFed allows uploads from an internal IP range to S3.
 - One bucket per Diamond experiment.
 - Ldap plugin provides read access for users to their experiment bucket.
 - **Last step is still highly experimental.**
- Appear to be having problems getting TPC to work on DynaFed running ldap plugin.



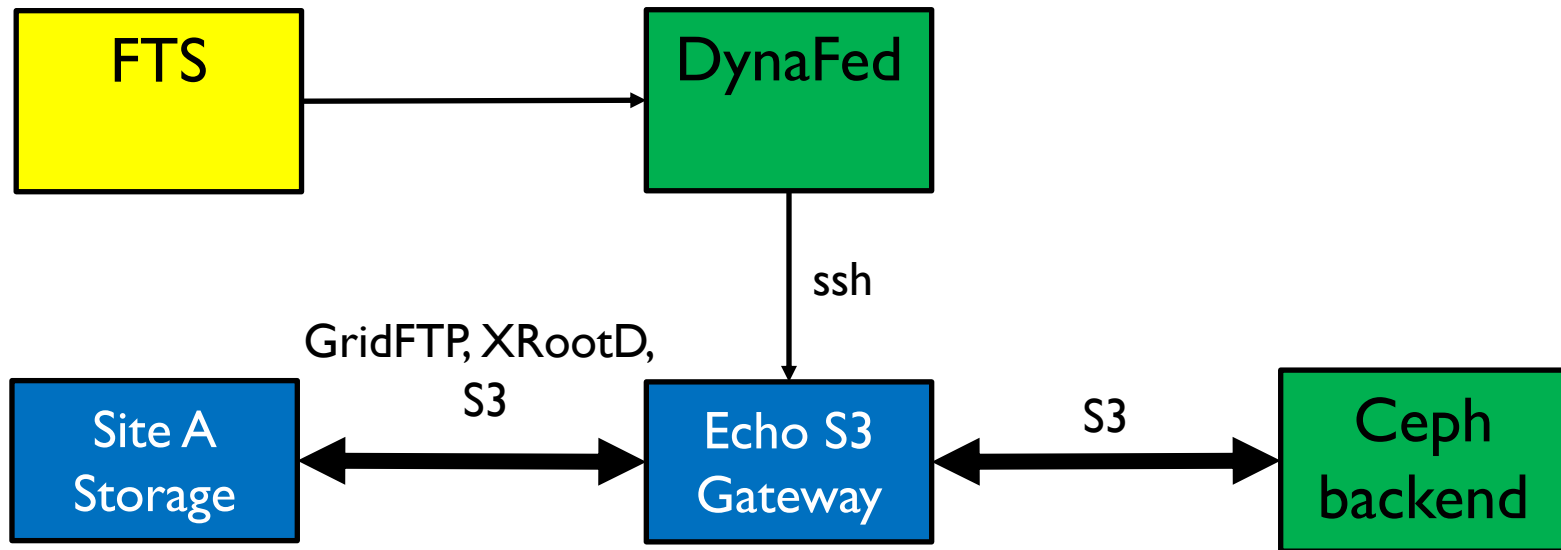
To do

- Monitoring
 - We have an Elastic Search Cluster so can copy Canadian setup.
 - **Is anything worth putting in InfluxDB instead?**
 - Would also like to implement quota monitoring.
- 4th Party Copies
- Checksumming



4th Party Copies

10



- If DynaFed is requested to mediate the transfers it will “stream” it by running a script.
- We can modify this to ssh into an Echo gateway and stream the transfer through that.
- This should offer “full TPC support” for all protocols!
- Significant DevOps work required.



Host Certificates

11

- RAL use QuoVadis commercial certificates on our S3 endpoints.
 - Web browsers will automatically trust it.
- Problems for Grid applications as the certificates are intermediate (not root).
 - Root certificate is IGTF trusted
 - [Some] Grid storage endpoints insist on a signing policy, which isn't present.
 - **Need to add this to the UK CA packages.**

