

DynaFed at RAL

Ian Johnson, Alastair Dewhurst
Scientific Computing Department
STFC Rutherford Appleton Laboratory



Context

2

- The RAL Tier-1 uses DynaFed mainly to increase the usability of ECHO S3 storage
- DynaFed provides convenient naming of data locations for various VO users, and can simplify their workflows
- Sources can be renamed but interface to user remains constant
- The service has near-production status; there aren't (yet) any users who would find its absence too detrimental



ATLAS Support

3

- DynaFed configured to view ATLAS Event Service Buckets.
 - Not used by production transfers, but helpful when debugging problems.
- DynaFed was used during testing of Microsoft Azure.
 - Aim of testing was to create complete site in the cloud.
- DynaFed provided frontend / authorization layer.
- Had to also deploy an FTS instance in Azure.
 - Transfers went via protocol translation through this instance.
- If the protocol translation feature could be added to DynaFed, then installing just a DynaFed service in a cloud would allow VOs (well ATLAS at least) to fully integrate cloud storage.



LHCb support

4

- Andrew Lahiff wrote ARC CE plugin to upload jobs logs to S3 buckets.
 - Users can then access the logs via DynaFed.
 - Object Lifetime on bucket means files get automatically cleaned up after a certain period.
- Used by LHCb to store job logs from RAL farm for extended periods of time.
- There is also an HTCondor plugin for S3, this allows all job output to be uploaded to a bucket.
- (Very) Small VO's find this a useful way of browsing job outputs.

Upload directly to S3, access via DynaFed



LIGO Support

5

- RAL are working with LIGO to provide CPU and S3 storage to assist UK exploitation of the data.
- A relatively small amount of data (~10TB) has been uploaded via DynaFed (using Davix).
- LIGO Data Replication (LDR) transfers data via GridFTP.
- Users primarily access their data via CVMFS, which in turn pulls data directly from S3 (or an alternate location if data not in ECHO yet).



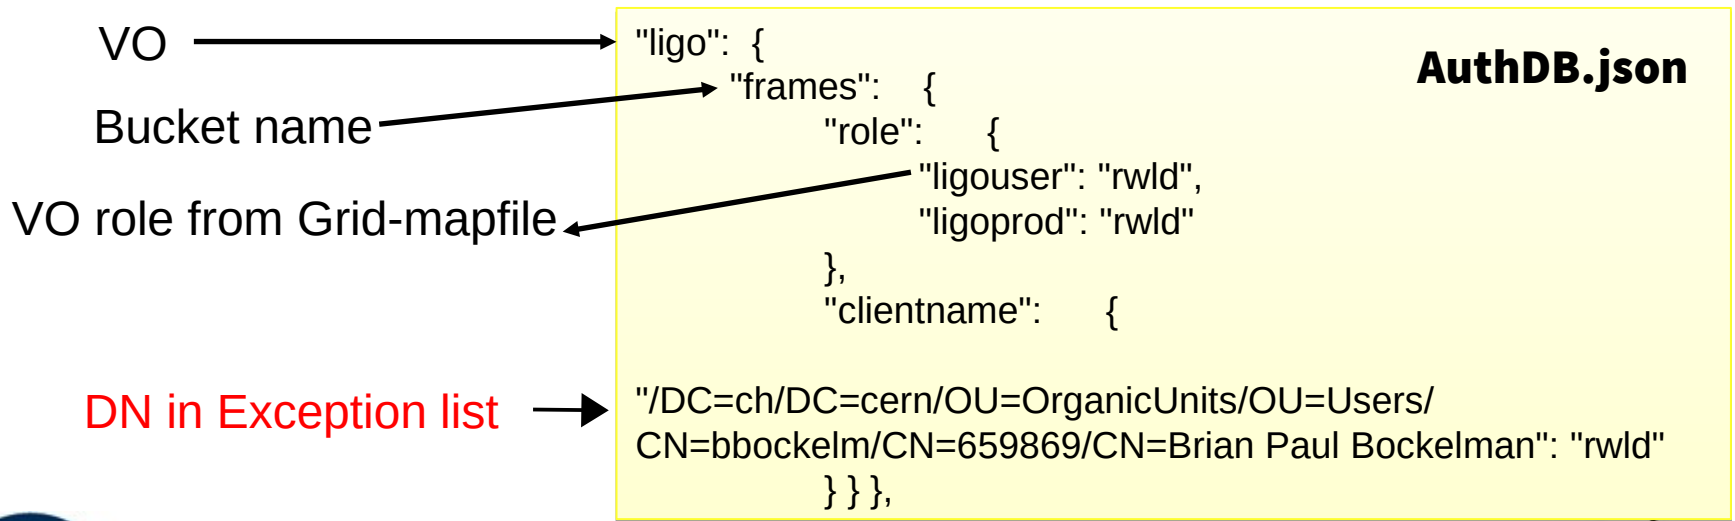
Access Control

- Currently, VO support in the RAL DynaFed instance uses :-
 - SSL client certificate chain, containing identity information to authenticate requests
 - Grid Certificates contain a DN field
 - (Optionally) A VOMS proxy certificate with attributes such as VO role
 - For authorization, we use a grid mapfile to map the VOMS information to a role, and a permissions table to define the operations that role can perform on specific S3 buckets



Grid-mapfile Authorization

- We have written a grid-mapfile authorization plugin.
- github.com/stfc/ral-ceph-tools/tree/master/DynaFed
- Same used method for both CLI access (VOMS attributes) and web browser - **as web browsers don't support VOMS, fall back to lookup identity against an Exception List of DNs.**
- Uses authDB-like format for authorization decision.



Supporting non-HEP users

8

- An important role for the ECHO cluster is to support users from other communities :-
 - LIGO and ENMR (X-Ray crystallography, current)
 - SKA (Radio astronomy, potential)
- Some of these users are happy to use Grid certificates, and may even supply VOMS attributes
 - Others will find certificate enrollment and maintenance a barrier to adoption and usability, and would rather use their institutional login or other methods to authenticate
- Hence, RAL are exploring other authentication schemes, funding from Horizon 2020 project EUDAT



Custom authC for DynaFed

- One use case is supporting requests from S3 clients - this would allow a DynaFed federation to use the same type of authentication as a certain public cloud storage service
- Propose using a separate, trusted authenticator process in front of DynaFed
- S3 client would send requests to the Authenticator, which checks the signature on the HTTP *Authorization* header. If valid, the Authenticator redirects the S3 client to DynaFed with a token, encrypted to prove the identity has been verified by a trusted entity.
- Token can contain a role name, which is consumed by the DynaFed authorization plugin



Other DynaFed Potential

1
0

- Do other institutes wish to support non-HEP communities? If so, what would their preferred authentication scheme be?
- What role can DynaFed play in protocol translation? This has arisen in connection with ATLAS (potential to deploy a whole site in the cloud) and LIGO (interfacing with LIGO Data Replication)

