


Federating Distributed Storage for Clouds in ATLAS



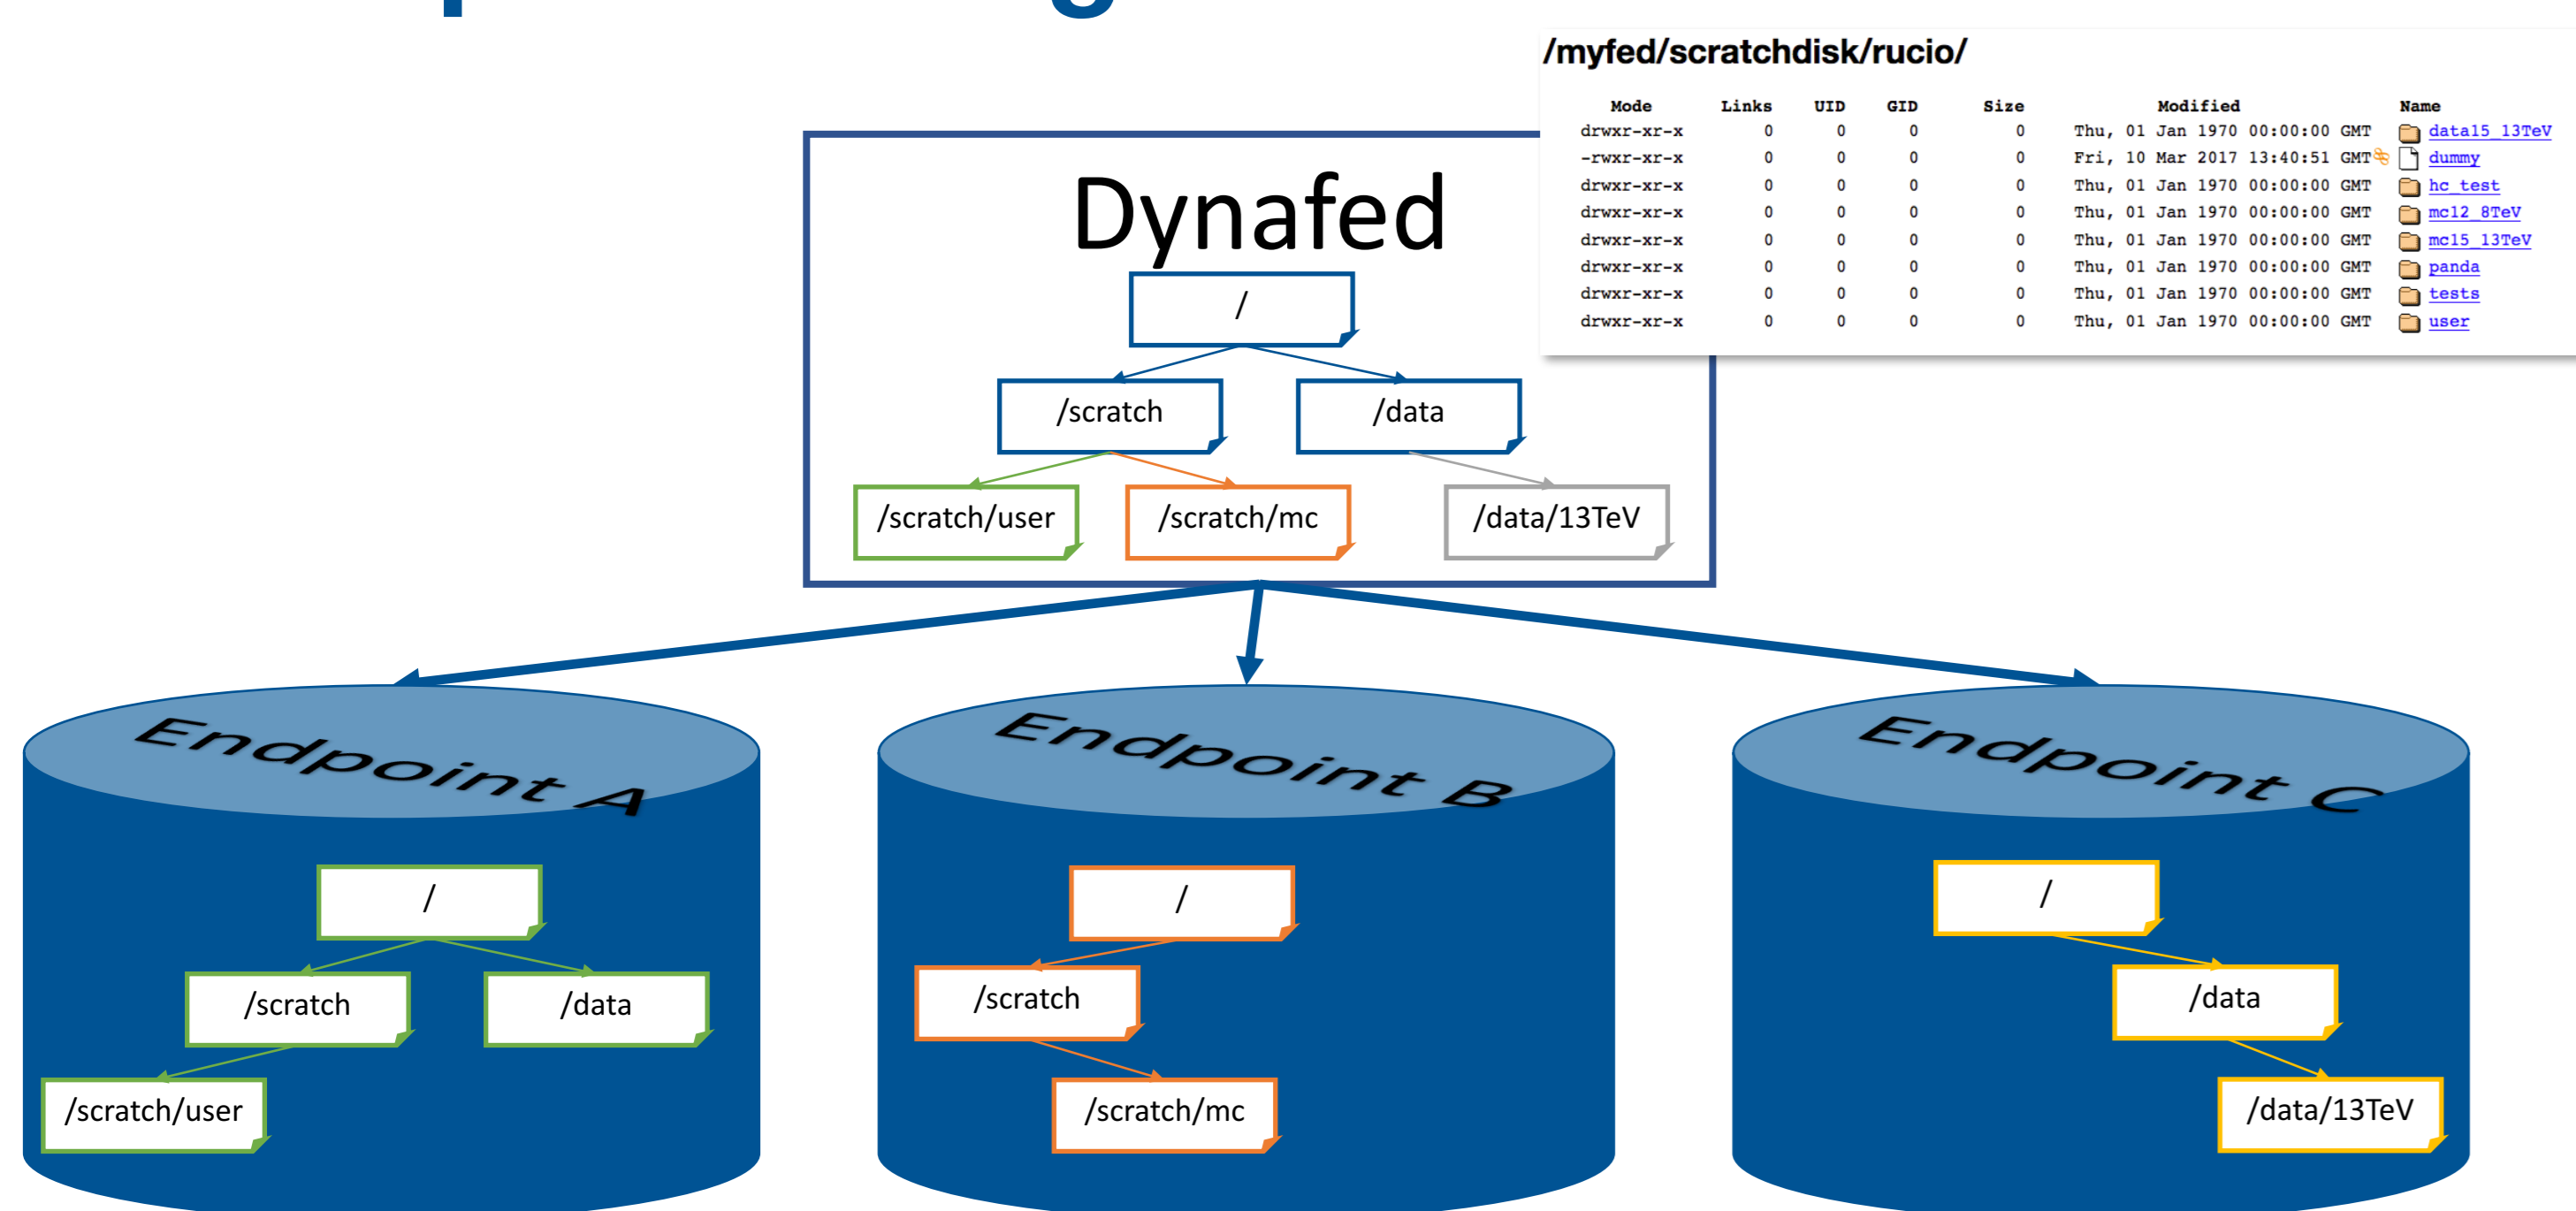
F. Berghaus, K. Casteels, A. Di Girolamo, C. Driemel, M. Ebert, F. Furano, F. Galindo, M. Lassnig, C. Leavett-Brown, M. Paterson, C. Serfon, R. Seuster, R. Sobie, R. Tafirout, R. P. Taylor



Introduction

Our goal is run data-intensive applications on globally distributed opportunistic resources that have no local HEP storage. We want to use a data federation, such as DynaFed [1] to redirect the applications to the optimal storage endpoint to retrieve input or deposit output data.

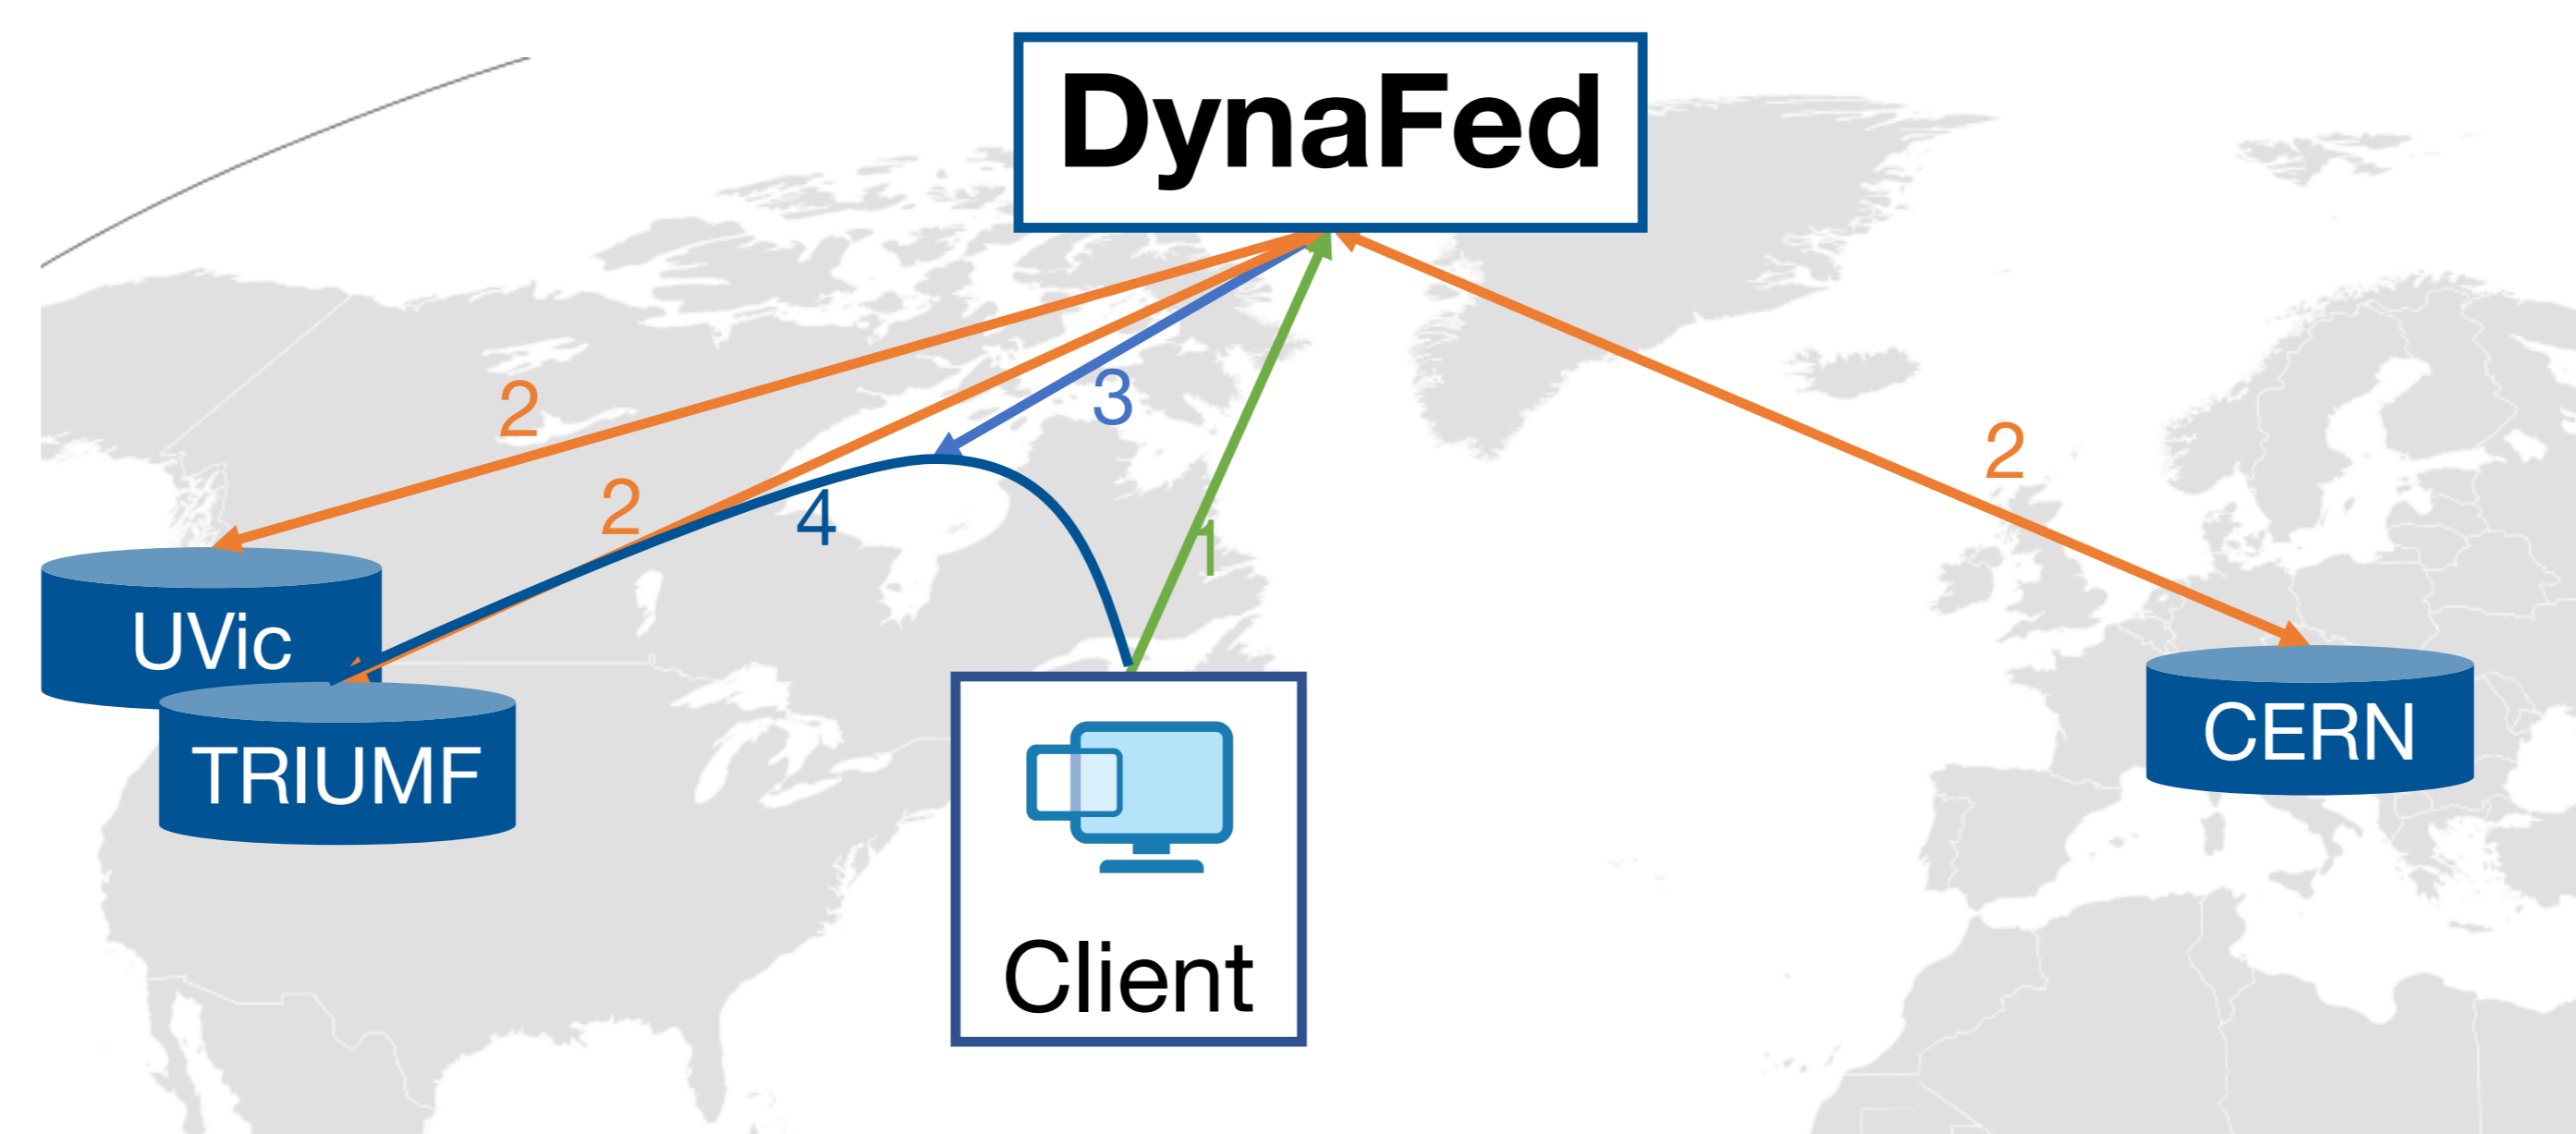
Conceptual Design



- Creates the appearance of a unified namespace from distinct endpoints
- Provides file system catalogue with indexing on standard protocols (HTTP, WebDAV, NFS)

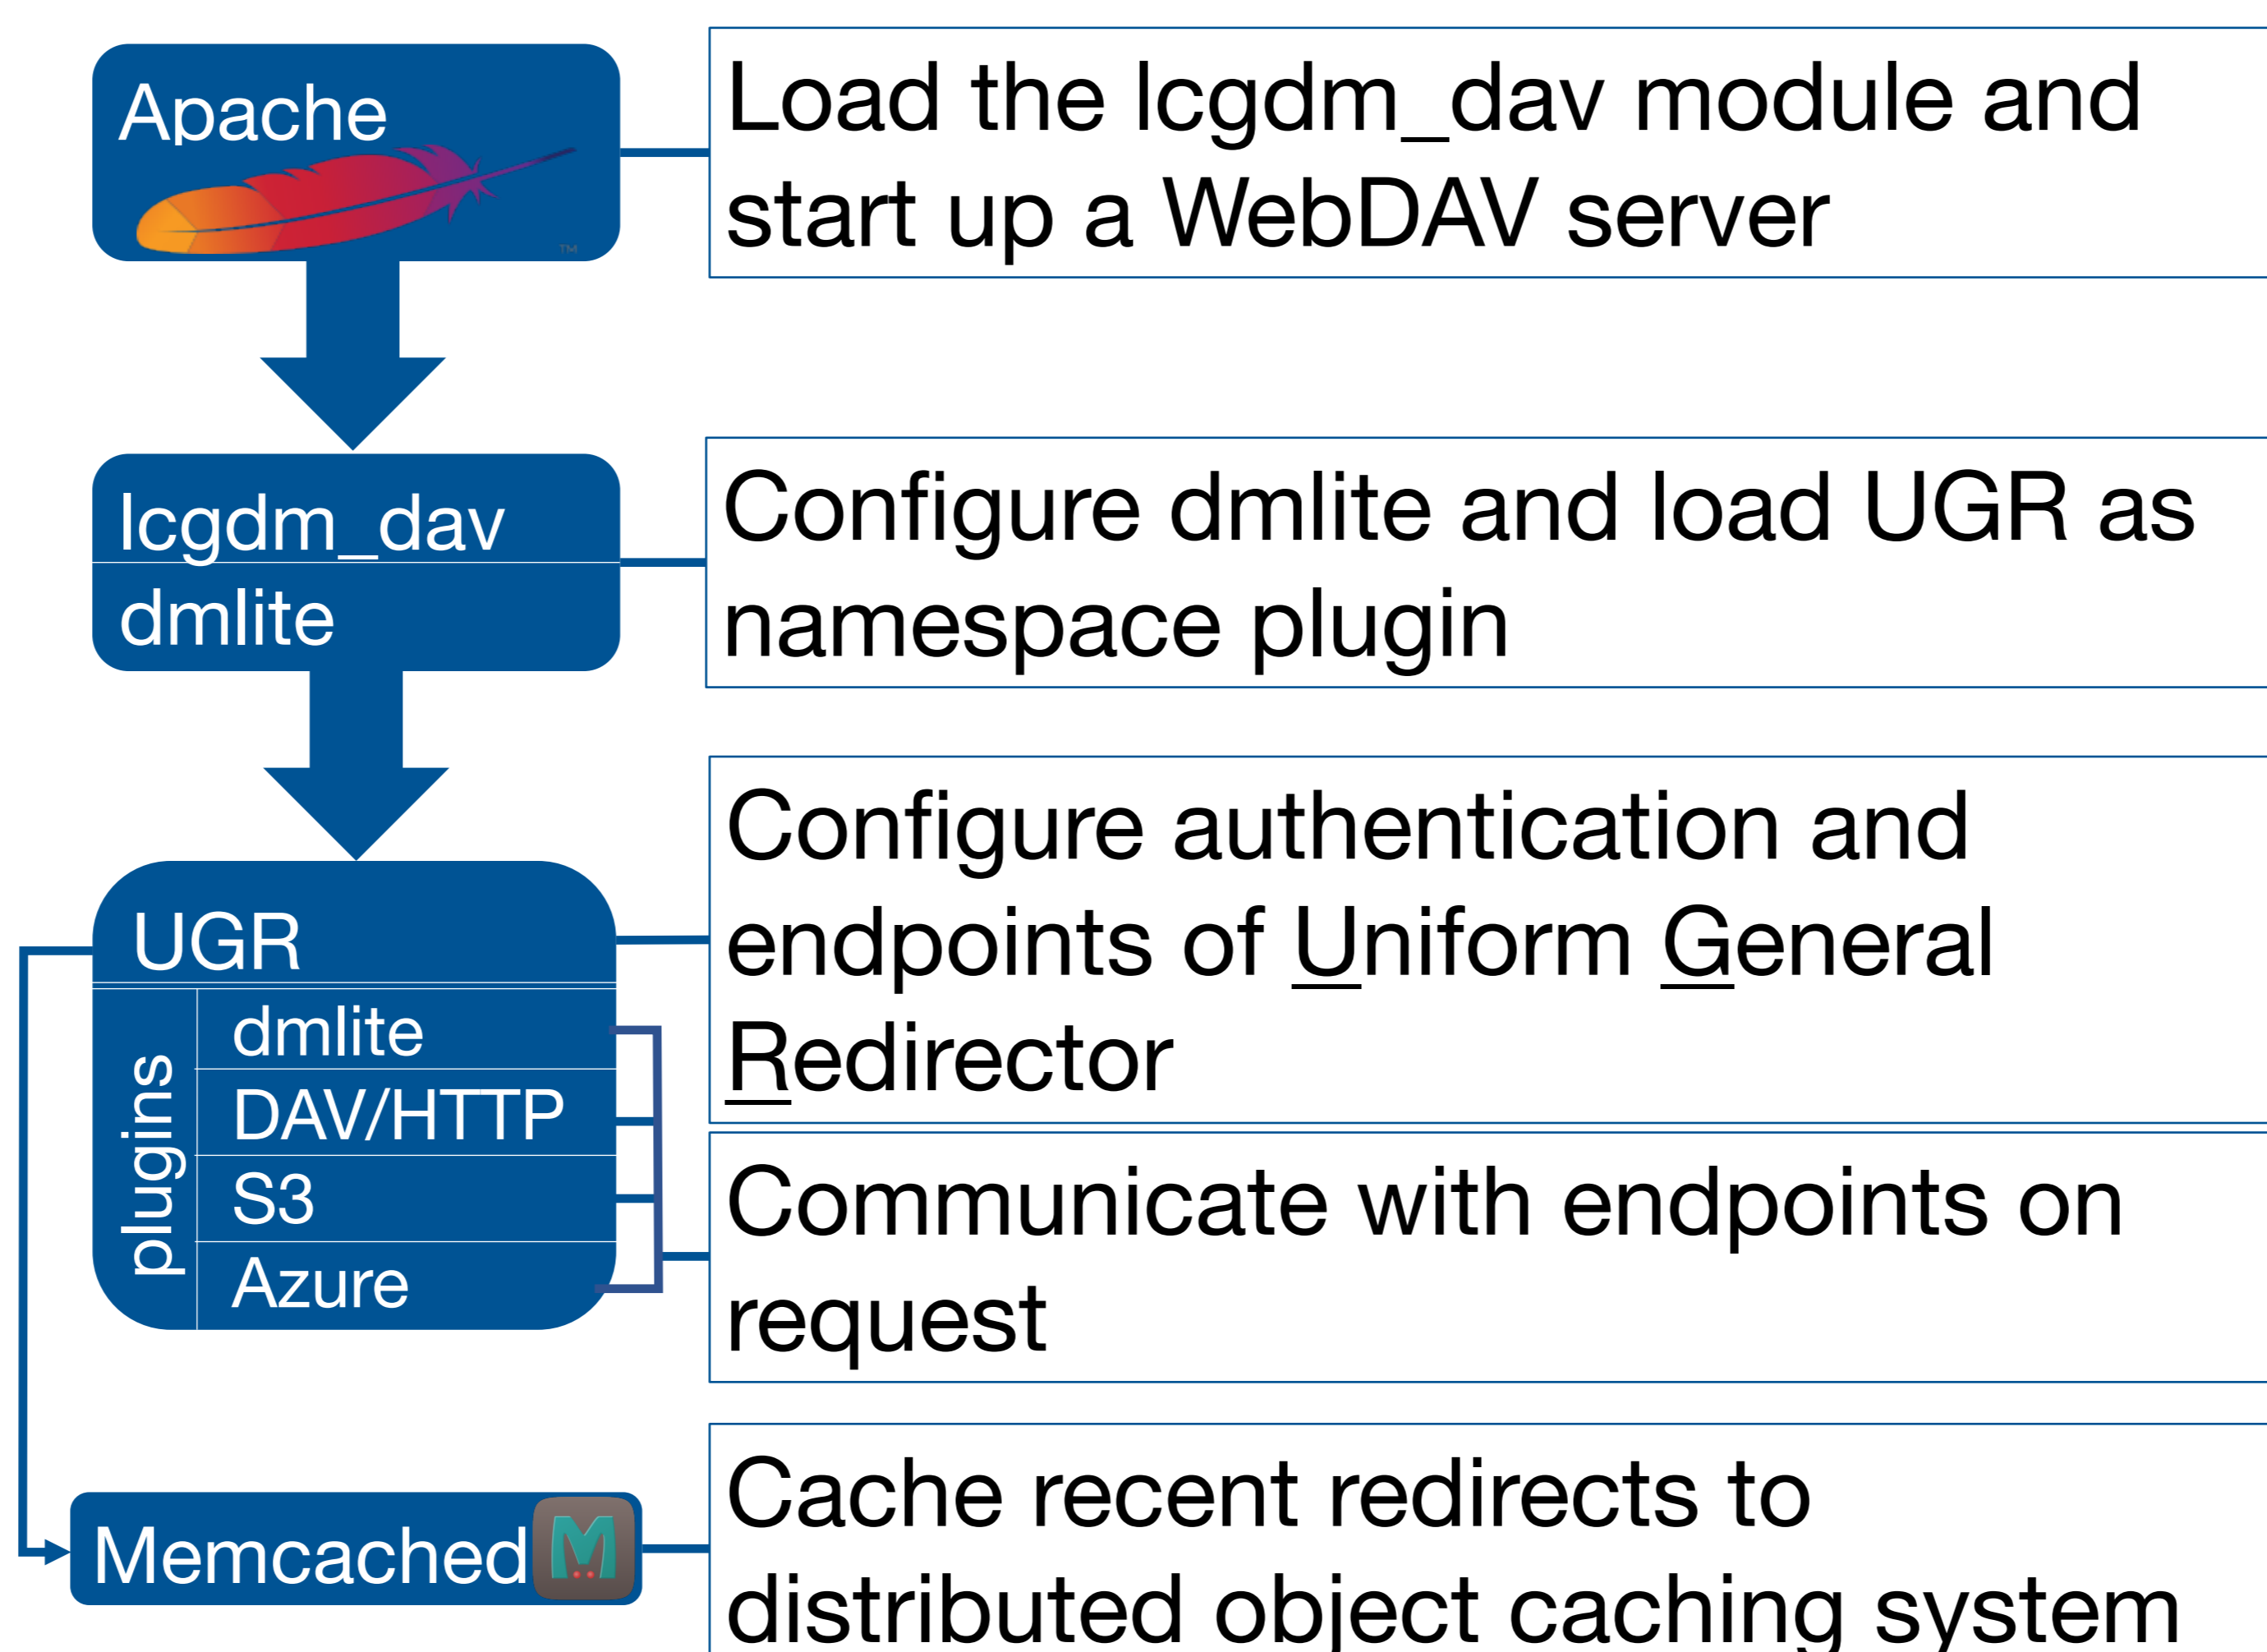
Bibliography

- 1) F. Furano *et al*, **Dynafed - The Dynamic Federation project**, May 24, 2017, cern.ch/lcgdm/dynafed-dynamic-federation-project
- 2) I. Gable *et al*, **CloudScheduler**, May 24, 2017, cloudscheduler.org
- 3) T. Maeno *et al*, **PanDA**, May 24, 2017, pandawms.org
- 4) C. Serfon *et al*, **Rucio**, May 24, 2017, rucio.cern.ch



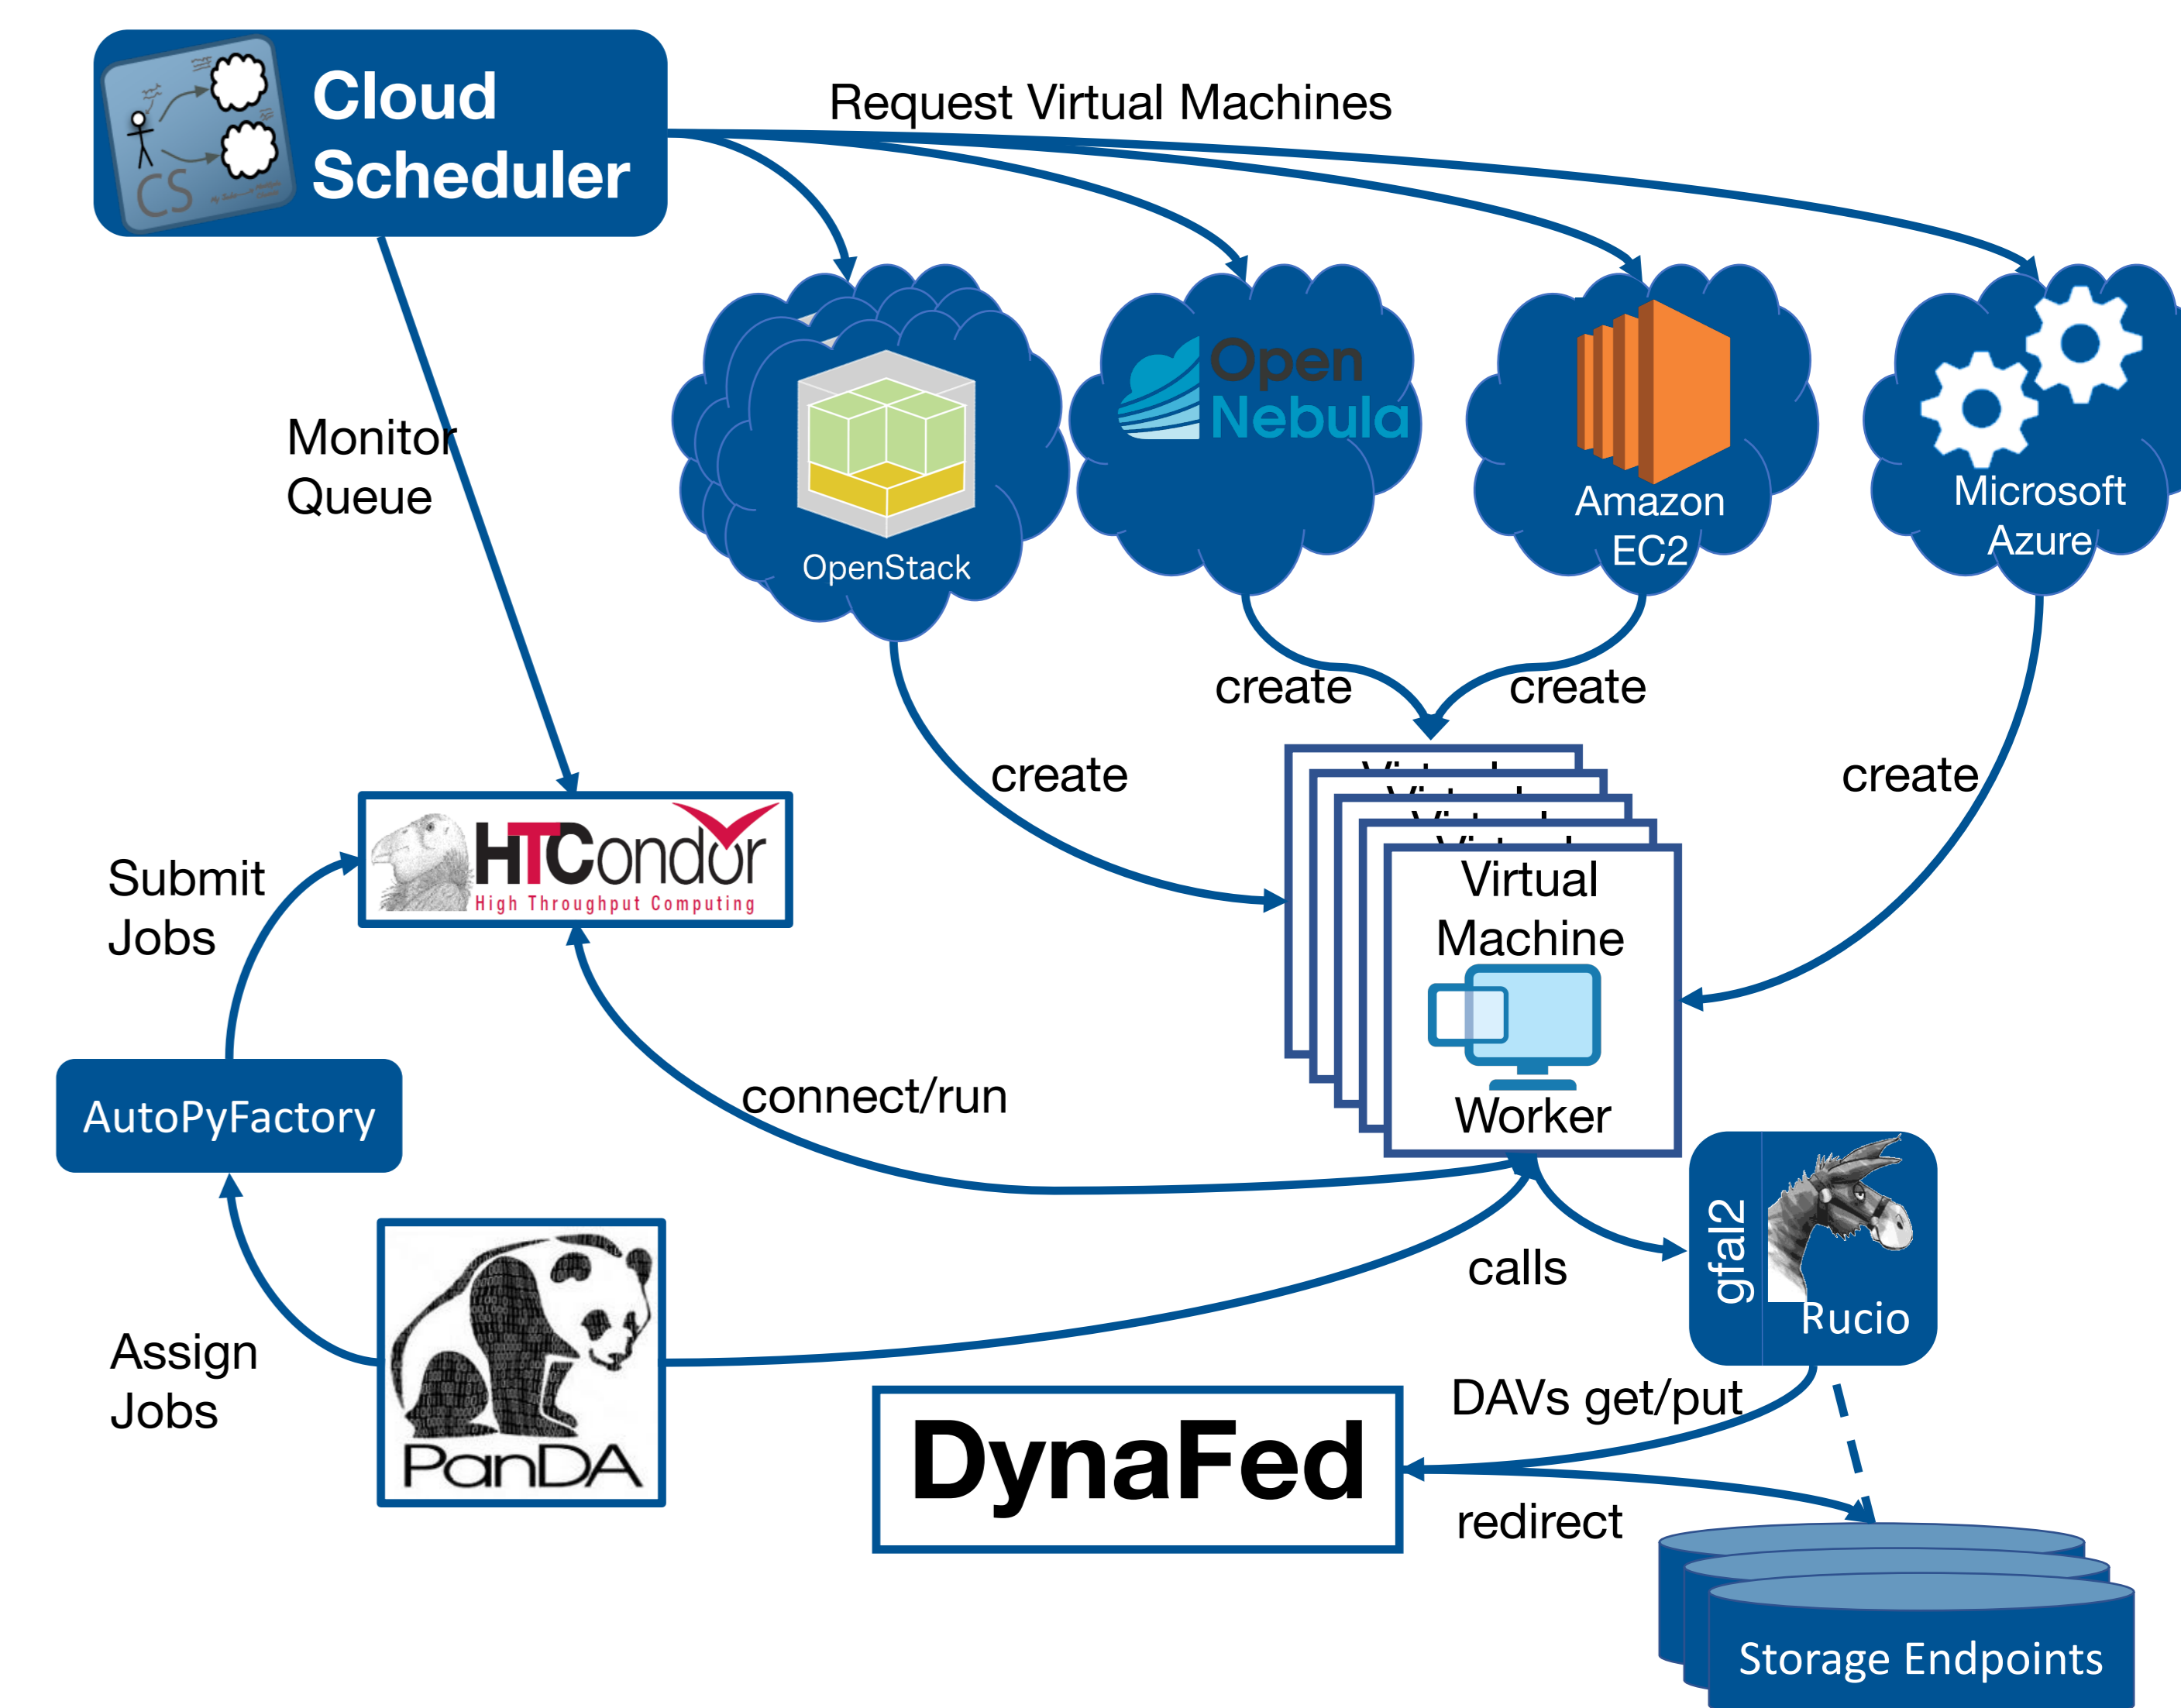
Data Access

1. Regularly poll connected storage endpoints
2. Receive client request for file from federated namespace
 - Authenticate client using X.509/VOMS
3. Poll connected storage endpoints and caches results
4. Redirect (302 Found) client to geographically closest copy
 - Sign redirect URL to endpoints that do not support X.509/VOMS
5. Client connects to endpoint



Application Workflow

1. The AutoPyFactory submits jobs to HTCondor queue
2. Cloud Scheduler [2] requests Virtual Machine (VM) instance on distributed cloud infrastructure
3. VM connects to queue and executes job
4. Pilot job queries Panda [3] for workload
5. Rucio [4] fetches input data from DynaFed
6. Rucio uploads job output to DynaFed



Summary

A dynamic storage federation supporting HTTP and WebDAV protocols has been integrated into ATLAS production environment. Full integration into the distributed data management system, Rucio, is in progress. The next steps are performance tests and integration into DIRAC and Belle II computing.