

Federated ATLAS Xrootd - Overview

Charles G Waldman

Rob Gardner

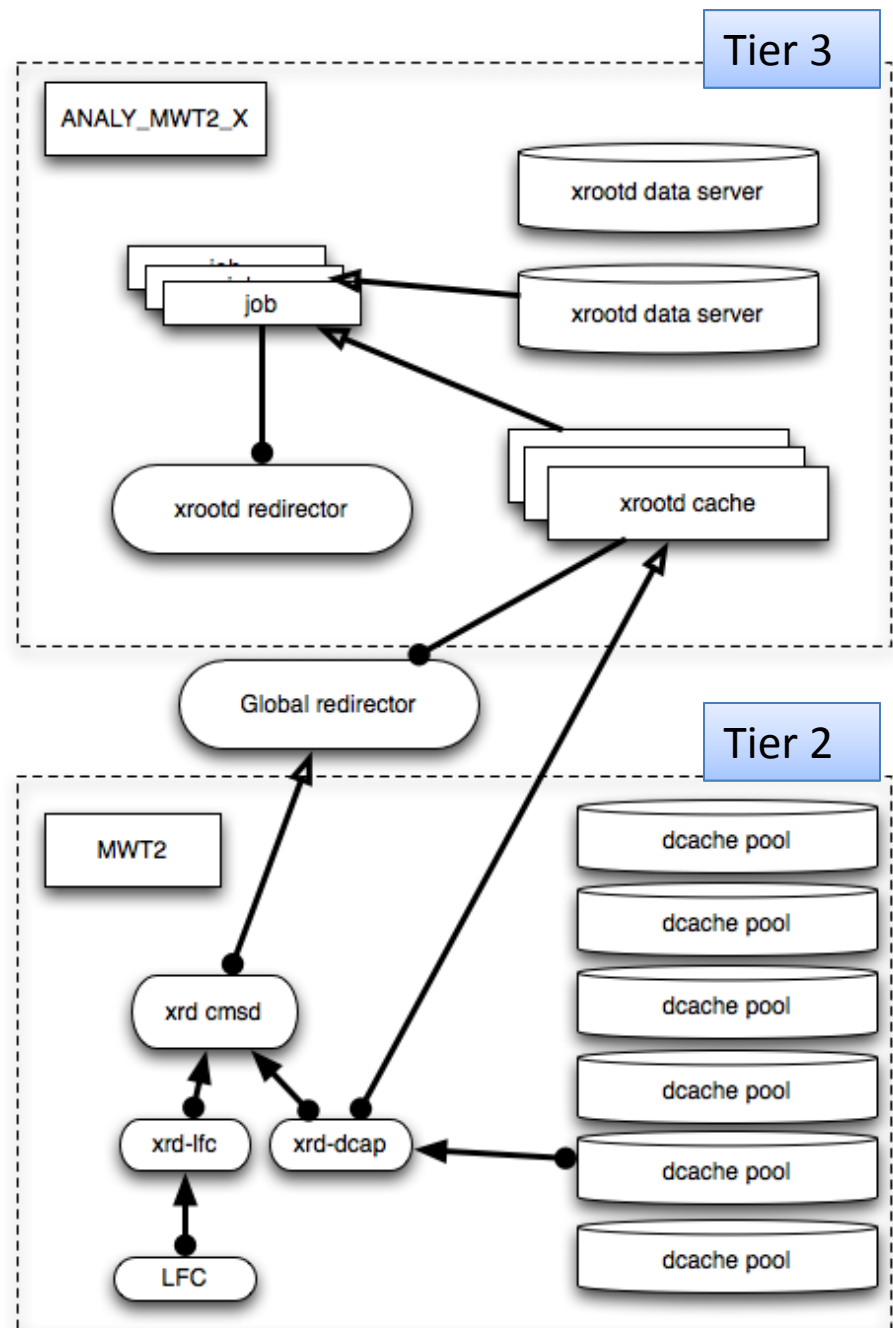
Overview:

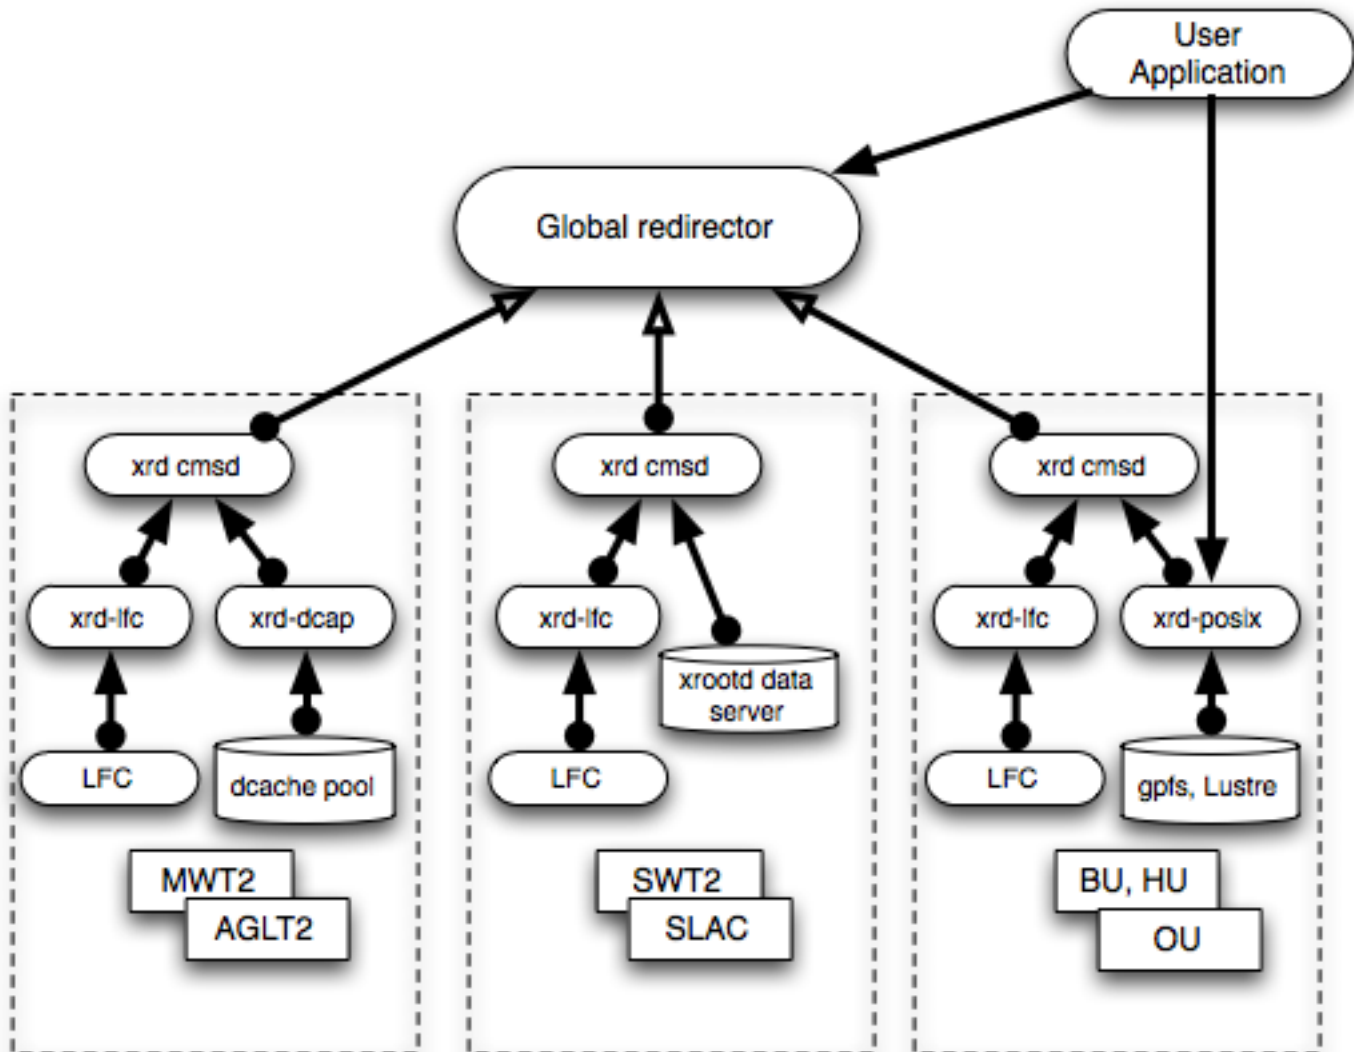
- Idea: Use xroot protocol as common interface to federate storage systems at T2 and T3 sites (and T1)

Motivation

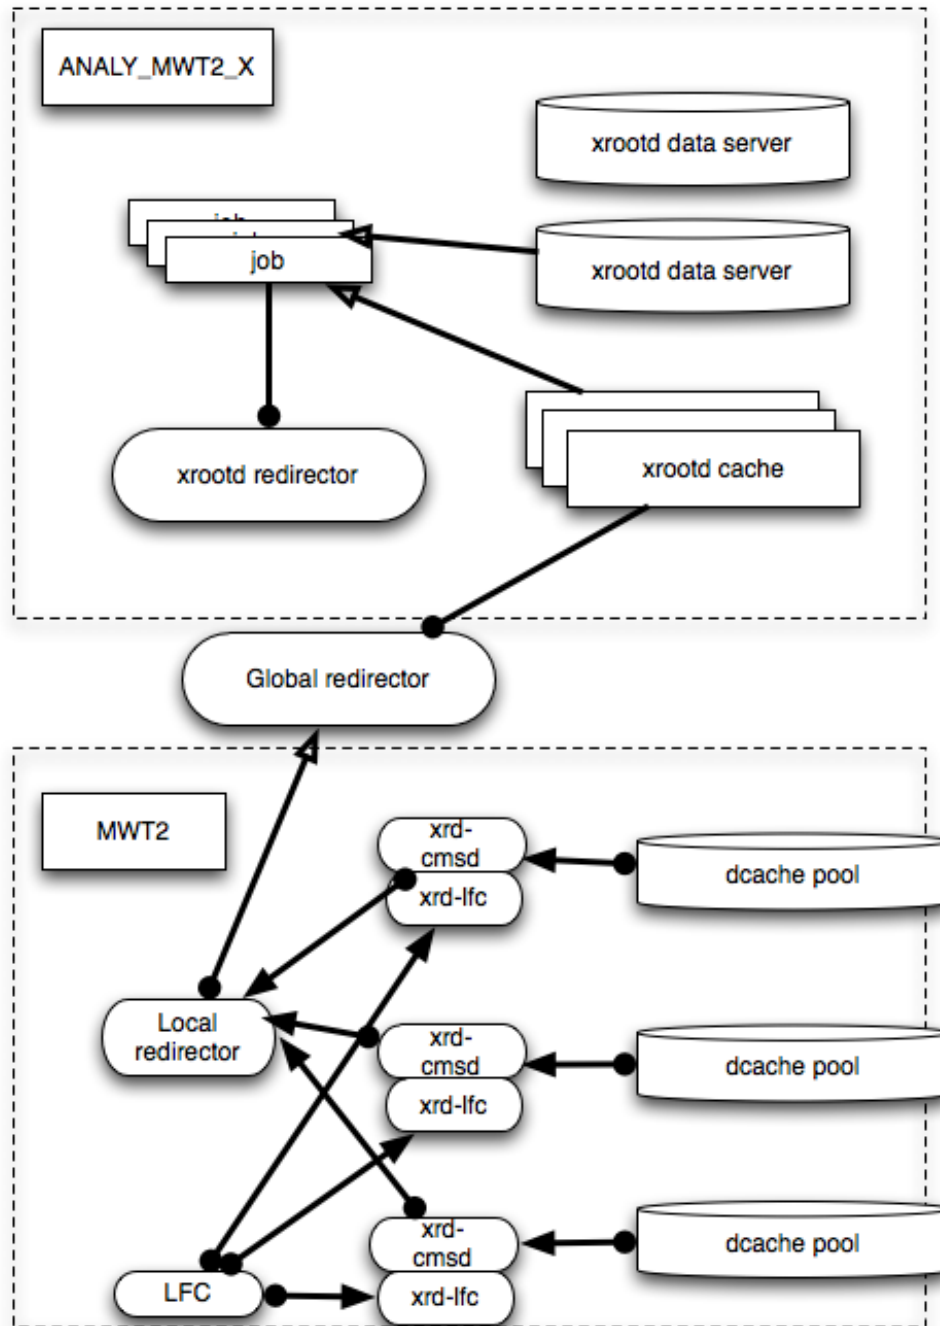
- 1) Simplify data access for T3, reduce data management burden (ToA, FTS, DDM, etc)
- 2) Allow T3 sites to share data
- 3) Provide flexibility for T2 sites, especially complex T2 sites which are “internally federated” (spanning several physical locations)
- 4) Support dCache, xrootd, GPFS, Hadoop, etc as storage backends

Getting data from T2 dCache to a T3 with xrootd





Direct Pool Access



Issues

- Lack of real “global namespace”. We are using modified version of LFC path for this, but
 - Conventions have changed
 - complications with `_dis`, `_sub`, etc
 - Future of LFC is in doubt
 - LFC lookup cannot be avoided due to `_DQ2` timestamps
- DSN/filename identifies file, but LFC does not allow this to be searched directly
- Xrd does not support passing GUID, unless file will be stored under this name (xrd dev required)

Questions

- How does this fit with future plans for LFC?
 - Is the global namespace == LFC?
 - Can we use GUID, DSN, etc to identify files?
 - xrd “opaque data”, in development
 - Will a single central LFC work?
- What is impact on T2 (or T1) site?
 - Minimal additional services required
 - Bandwidth allocation – use dedicated door, or eg. bwctl to control resource use.

Status

- Deployed services at three Tier 2's
 - MWT2(dCache), SLAC, SWT2 (XRootd)
- Functional testing successful:
 - Xrd-native
 - Xrd-dcap
 - Xrd-direct pool
- Performance testing started

Next Steps

- Additional Tier2's and backends: NET2 (GPFS), AGLT2 (dCache)
- Resolution of LFC-Global namespace convention
- Extend client tests to include job-clients: ANALY_MWT2_X queue → HammerCloud
- Testing in the Tier 3 environment – sites, FRM (cf. Wed Tier 3 session)

Thanks

- Doug Benjamin
- Wei Yang
- Andy Hanushevsky
- Hiro Ito
- Patrick McGuigan
- Brian Bockelman