

Deployment of federated xrootd infrastructure in ATLAS

Rob Gardner

Storage pre-GDB meeting

Oct 9, 2012

Summary of the deployment thus far

- In ATLAS we have deployed a new xrootd infrastructure suitable for WAN access across a growing fraction of storage resources
 - FAX - Federating ATLAS (storage) using Xrootd
- US: the Tier1, all Tier2 centers (separately, a number of Tier 3 sites)
- UK: three Tier 2 sites
- DE: three Tier 2 sites plus a CZ site
- RU: two Tier 2 sites
- EOS

A different sort of project

- Sites deploy a tandem of xrootd services; in principle it should be as easier than installing apache web server
- But the software, while a proven storage technology, requires development to become a federating technology
- Requires
 - Experiment-site specific file lookup service (i.e. N2N)
 - Customizations for backend storage types
 - Various 3rd party wide-area monitoring services (UCSD collector, ActiveMQ, Dashboards)
 - Security for read-only access: missing initially; still need proxy validation
 - Standardizing monitoring metrics => further development
 - Status monitoring and alert systems for operations
 - Development in production systems: pilot, info sys
 - Site-level caching (file and sub-file) will require more development
- Good news most of these obstacles have been addressed in the R&D phase. In ATLAS, FAX-enabled pilot is undergoing testing at two Tier2 centers; co-located Tier3 users are using FAX endpoints (locally)

Examples: functional status & cost performance between analysis endpoints

Show 200 entries Copy Print Save view: Network

Site Name	Site Info		Network Measurements
	Source	Destination	
ANALY_MWT2_to_ANALY_CERN_XROOTD	MWT2	CERN-PROD	2.074
ANALY_SWT2_CPB_to_ANALY_CERN_XROOTD	SWT2_CPB	CERN-PROD	2.079
ANALY_OU_OCHEP_SWT2_to_ANALY_CERN_XROOTD	OU_OCHEP_SWT2	CERN-PROD	2.083
ANALY_CERN_XROOTD_to_ANALY_AGLT2	CERN-PROD	AGLT2	6.143
ANALY_MWT2_to_ANALY_NET2	MWT2	BU_ATLAS_Tier2	6.6
ANALY_CERN_XROOTD_to_ANALY_NET2	CERN-PROD	BU_ATLAS_Tier2	6.794
ANALY_AGLT2_to_ANALY_CERN_XROOTD	AGLT2	CERN-PROD	7.769
ANALY_NET2_to_ANALY_NET2	BU_ATLAS_Tier2	BU_ATLAS_Tier2	7.982
ANALY_AGLT2_to_ANALY_NET2	AGLT2	BU_ATLAS_Tier2	8
ANALY_SWT2_CPB_to_ANALY_NET2	SWT2_CPB	BU_ATLAS_Tier2	8
ANALY_OU_OCHEP_SWT2_to_ANALY_NET2	OU_OCHEP_SWT2	BU_ATLAS_Tier2	8.25
ANALY_CERN_XROOTD_to_ANALY_CERN_XROOTD	CERN-PROD	CERN-PROD	10.513
ANALY_OU_OCHEP_SWT2_to_ANALY_OU_OCHEP_SWT2	OU_OCHEP_SWT2	OU_OCHEP_SWT2	13.631
ANALY_SWT2_CPB_to_ANALY_SWT2_CPB	SWT2_CPB	SWT2_CPB	14.266

There are many more components as discussed at the Lyon storage federation workshop last month



ATLAS Federated Xrootd Status - 2012-10-08 07:15:51

[Frequently Asked Questions](#)

Host: atl-prod09.slac.stanford.edu (atl-prod09.slac.stanford.edu)

Metric	Last Executed	Enabled?	Next Run Time	Status
org.usatlas.xrootd.grid.xrscp.compare	2012-10-08 07:05:00 CDT	YES	2012-10-08 07:20:00 CDT	OK
org.usatlas.xrootd.grid.xrscp.direct	2012-10-08 07:05:02 CDT	YES	2012-10-08 07:20:00 CDT	OK
org.usatlas.xrootd.grid.xrscp.fax	2012-10-08 07:05:02 CDT	YES	2012-10-08 07:20:00 CDT	OK
org.usatlas.xrootd.ping	2012-10-08 07:05:02 CDT	YES	2012-10-08 07:20:00 CDT	OK

Host: atlas-cm4.bu.edu (atlas-cm4.bu.edu)

Metric	Last Executed	Enabled?	Next Run Time	Status
org.usatlas.xrootd.grid.xrscp.compare	2012-10-08 07:05:01 CDT	YES	2012-10-08 07:20:00 CDT	OK
org.usatlas.xrootd.grid.xrscp.direct	2012-10-08 07:05:01 CDT	YES	2012-10-08 07:20:00 CDT	OK
org.usatlas.xrootd.grid.xrscp.fax	2012-10-08 07:05:01 CDT	YES	2012-10-08 07:20:00 CDT	OK
org.usatlas.xrootd.ping	2012-10-08 07:05:01 CDT	YES	2012-10-08 07:20:00 CDT	OK

Commoditizing

- Unlike other “middleware” services – there is no team of developers, integrators, packagers, pre-production service validations, documenters and the like
- The work has been done with “spare” effort from various groups.
 - With large efforts coming from teams whose first responsibility is to provision resources for the experiments as part of WLCG pledges
 - Mostly in the context of experiments, but with vital and timely development from the software teams (Xrootd, dCache, DPM & EOS)
 - Where possible ATLAS has leveraged developments from USCMS-AAA
- Is there a role for (technical) WLCG group to vigorously pursue **delivery of production federation technologies?**

Advantages

- A point of coordination for the common elements for the basic services that have emerged from the R&D programs
- Help drive requirements and priorities and liaison with software providing groups
- Help with packaging, deployment configurations, documentation, site support
- Coordinate extensions (protocols, caching) and drive consistency in the architecture