

ATLAS Computing Operations

Ueda I.
WLCG Collaboration Workshop
2013.11.11.

Contents

Site-oriented matters

- WebDAV
- Xrootd
- CernVM-FS
- Multi-Core Resources

ATLAS-oriented matters relevant to the sites

- Rucio commissioning
- FAX
- Cloud Operations

WebDAV

Migration from DQ2 to Rucio requires 'renaming' of SURLS

i.e. either

- WebDAV access to the storage for the Central renaming machinery
- Local renaming in collaboration with DQ2-ops

All the ATLAS sites are expected

- to provide WebDAV access for HEAD, MKCOL, MOVE
- or to discuss with ATLAS a Rucio migration strategy

to complete the **migration** by the end of 2013

To be followed up with

- Savannah tickets, one assigned to each cloud
 - to track each site status/plan
- Reviewing status and plans at Atlas Distributed Computing meetings
- Sites may get GGUS tickets depending on the clouds

ATLAS twiki: <https://twiki.cern.ch/twiki/bin/view/AtlasComputing/StorageSetUp#WebDAV>

WLCG twiki: <https://twiki.cern.ch/twiki/bin/view/LCG/WLCGOpsCoordWebDAV>

Xrootd

According to the studies presented at the ATLAS Software & Computing Workshop

- **DPM Sites are recommended to use**
 - **xrood directI/O for analysis**
 - **xrdcp copy-to-scratch production**
- Further tests with more loads being discussed

CernVM-FS

ATLAS encourages sites to deploy CernVM-FS 2.1

- **WLCG Baseline Version**
 - <https://twiki.cern.ch/twiki/bin/view/LCG/WLCGBaselineVersions>
- To be able to use multiple repositories on WNs
 - ATLAS software and file-base data `/cvmfs/atlas.cern.ch`
 - ATLAS nightly test releases (`/cvmfs/atlas-nightlies.cern.ch/`)
 - SFT repository of lcg software (`/cvmfs/sft.cern.ch/lcg`)
- cvmfs servers at CERN can be upgraded to 2.1 only when all the clients are 2.1
 - Eventually we may set a deadline for sites upgrading to 2.1 in the near future

Checking the validity of CVMFS cache on WNs

- ATLAS has observed some cases with stale CVMFS cache on WNs
 - ATLAS shifters/experts find the issue --> GGUS --> Site admins treat them
- Wouldn't it be useful to have automatic checks on WNs run by the sites?
 - Tool finds the issue --> remove WNs from production --> Site admins treat them
 - CERN has started doing this...
- **WLCG-wide common approaches** to be discussed under WLCG operations coordination

Multi-Core Resources

ATLAS software (AthenaMP) has been validated and ready for production

Still, not going for wide deployment immediately

- Currently fine with all T1s and some more sites providing “static” dedicated resources as multi-core queues
 - NDGF is providing “dynamic” multi-core resources
- “Static” resources would not be optimal
 - unless we are sure about the ratio between multi-core and single-core jobs

Dynamic allocation of multi-core WNs

- **WLCG-wide common approaches** to be discussed
 - how efficient it can be
 - how sites should setup multi-core resources in a dynamic way

Accounting

- What is the situation about WLCG accounting for multi-core job slots?
 - (<https://twiki.cern.ch/twiki/bin/view/LCG/WLCGOpsMinutes130718>)
- APEL entry should be
 - Wall-clock time : (\sim wall-clock time) * (HEP SPEC for N cores)
 - CPU time : (total cputime for the job) * (HEP SPEC for single core)
- Proper accounting needed before wide deployment

Rucio commissioning

The new ATLAS Distributed Data Management system (Rucio)

has been developed

- All the components have been implemented
- Stress tests have been carried out by the development team

and now being validated

- with the volunteer sites
 - those who replied to the call for volunteers at ADC Weekly 24 Sep 2013
- starting with the basic functions
 - file registration, transfers, quota checks, deletion

Commissioning and migration

- Early 2014
- Switching catalog look up from LFC to Rucio
- Full commissioning and migration from DQ2 to Rucio

FAX

Federating ATLAS storage systems using XrootD

Integration is progressing

- **Regular test jobs** to measure remote/local access performances
 - among the volunteer sites
 - building “cost matrix” to be used for job brokerage etc. in future
- **“Input file failover to FAX”** has been activated for some sites and monitored
 - actually rescuing some jobs from failing due to “unavailable” files
 - revealing some issues in some redirectors
- **More tests** (incl. stress tests) being planned
 - eg. to assess the impact of potential failures
 - Should be coupled with failover for StageOut. Implemented, needs testing

Longer term

- Job brokerage can be made beyond site boundaries
 - Sites should foresee accesses to SE from WNs of other sites
- HTTP-based federations to be evaluated

Cloud Operations

ATLAS has successfully integrated cloud resources into the production system

- **“Significant computational resources** were delivered to ATLAS in the form of **IaaS clouds dressed up as PanDA queues”** (ATLAS Software & Computing Workshop, October 2013)
- Started as R&D efforts, by various parties
- Now it is time to unify the methods and proceed towards **“Standardized Operations”**
 - People have started working on it

Additional Note

Quite some of the contents are duplication of what presented at the last WLCG Operations Planning Meeting

▸ <https://indico.cern.ch/conferenceDisplay.py?confId=280057>

- Hoping to draw attention from as many sites as possible
- Apologies for some people listening the same stuff twice...

Looking forward to further coherence between forums (GDB,WLCG Operations..) to minimize the risk of information loss (or duplication)