


Middleware Deployment

Simone Campana

ATLAS Jamboree, December 2012



CVMFS

- ATLAS requests CVMFS to be deployed at ***ALL*** sites
 - Being followed by the WLCG-OPS task force on CVMFS deployment
- Sites with small/no local disk in the WN can use CVMFS over NFS
 - This is ***NOT*** the recommended way for production clusters
 - CVMFS 2.1 will be production ready by the end of January
 - The CVMFS 2.1 NFS exporter is stable enough if one is prepared to have a single point of failure
 - cvmfs/nfs server must be SL6 or any other recent OS
 - Issues observed in case of Squid failure. Better deploy a reliable squid setup.
- ATLAS deadlines (out of production/analysis and no auto SW install)
 - The deadline for ATLAS sites to migrate to CVMFS is April 30 2013
 - For sites on SL6, only SW installation via CVMFS will be supported

Squids and Frontier

- We have no plan to change the Squid/Frontier deployment
 - O(5) Frontier servers, tentatively one squid service per site
 - Could be multiple servers
 - Preserve the current fallback scenarios for both squid and frontier
- Number of Frontier servers could slightly decrease (see tomorrow' session)
- Squid Monitoring under discussion (WLCG-OPS task force) but preliminary:
 - We ask all sites to publish their Squid service in GOCDDB/OIM
 - We plan to keep the current monitoring (SAM/MRTG) and extend SAM test
 - We ask for the squid *monitoring* port (default 3401) to be open for UDP from CERN IP addresses
 - 128.142.0.0/16 (CERN)
 - 188.185.0.0/17 (Hungary)

More on Squids

- Squids become an important service at sites for ATLAS
 - Especially for CVMFS where access to a remote Squid could be problematic (files could be big)
- We suggest a reliable setup of Squids at sites with 2 servers (or more) in load balance
 - resiliency and possibility to upgrade transparently
- As far as we know and are concerned, squids for Frontier and CVMFS can be shared

gLexec

- ATLAS restarted the test of utilization of gLexec from the ATLAS pilot
 - There is a non critical SAM test in place, will be refined following the ATLAS utilization scenario
- ATLAS is not ready to use gLexec at production scale
- While deployment of gLexec is encouraged (at least to facilitate testing at scale), ATLAS does not perceive this as priority
- No deadline/target date can be established today

SL6 worker nodes

- We discuss running on SL6 in SL5 compatibility mode
 - Need extra RPMs, see <https://twiki.cern.ch/twiki/bin/view/Atlas/RPMCompatSLC6>
- Production can run on SL6: passed physics validation
- Analysis has an issue with compilation
 - Except for 17.2 which has been validated
- Issues should be sorted out before Christmas
 - Testing the fix on 15.2 (sorting out last issues with KitValidation)
 - Most likely new caches and no new full releases
- Timescale for a green light is Jan 2013 TBC
- For “desperate” site something more aggressive can be negotiated with ATLAS
 - Run only production and analysis on 17.2. Remove all other SW tags

/etc/machinefeatures

- ATLAS supports this initiative
 - <https://twiki.cern.ch/twiki/bin/view/LCG/WMTEGEnvironmentVariables>
 - Basically a way to publish important per-node information
- Panda pilot is instrumented to collect the information
 - Does not expose them to the server yet, but will soon
- After validation, ATLAS plans to use them for:
 - Monitoring (e.g. HS06 of the job slot)
 - Improve the knowledge of the pilot runtime environment (e.g. CPU and memory limits)
- ATLAS would like to ask sites to follow the Hepix WG recommendations
 - Followed up at the GDB

perfSONAR

- ATLAS requests perfSONAR to be deployed at all sites
 - Central coordination from WLCG-OPS Task Force
- More details on Wed at the GDB but in short
 - We suggest the deployment of 2 perfSONAR machines (latency and throughput) per site
 - We ask sites to publish them in GOCDB/OIM
 - We ask to start configuring mesh tests in the cloud
- We ask one ATLAS contact person per cloud
 - Offering some help in the central configuration and following up regional deployment
- We plan to use perfSONAR measurements for
 - Monitoring of network links
 - Complement the SONAR tests possibly even in decision making in DDM

Contacts

- perfSONAR deployment
 - wlcg-ops-coord-tf-perfsonar@cern.ch
- Squid Monitoring
 - wlcg-ops-coord-tf-squidmon@cern.ch
- CVMFS deployment
 - wlcg-ops-coord-tf-cvmfs@cern.ch
- General ATLAS questions on Operations and Development
 - atlas-project-adc-operations@cern.ch
 - atlas-project-adc-development@cern.ch