Table of Contents

WLCG Tier1 Service Coordination Minutes - 8th April 2010

Attendance

Site	Name(s)
CERN	Julia, Nicolo, Miguel, Dirk, Patricia, Zbyszek, TIm, Maria, Jamie, Flavia, Maarten, Roberto, Alex K, Andrea V
ASGC	
BNL	Carlos
CNAF	Luca, Barbara
FNAL	Jon
KIT	Angela
IN2P3	Osman
NDGF	Vera
NL-T1	Ron
PIC	Gonzalo
RAL	Carmine, Andrew
TRIUMF	Andrew

Experiment	Name(s)
ALICE	
ATLAS	Dario
CMS	
LHCb	

Interventions foreseen during LHC stop (26 - 28 April)

Site	Intervention(s)
CERN	
ASGC	no interventions planned
BNL	no interventions planned
CNAF	Tape library intervention < 4 hours; migration of DB to new hardware
FNAL	no interventions planned
KIT	
IN2P3	
NDGF	no interventions planned
NL-T1	no interventions planned
PIC	no interventions planned
RAL	no interventions planned - may do a small network intervention (part of UPS room network)
TRIUMF	no interventions planned

glexec deployment status

Nagios glexec test results for "ops" on EGEE/EGI: here

Site	Status
CERN	
ASGC	OK for end May

BNL	
CNAF	
FNAL	Fully deployed, published monitored and used by CMS
KIT	OK for end May - have deployed, ready and working but didn't see any user of this service yet.
IN2P3	
NDGF	gLite related? NDGF have issues with pilot job concept (as stated at MB).
NL-T1	
PIC	
RAL	
TRIUMF	

Tentatively ok for all except BNL and NDGF where we are expecting more news.

Maarten - milestones on Tier1 sites first to make available and pass OPS tests. Should also configure for VOs supported. Other VOs will have to ensure by running same Nagios test that it also works for them. Discuss again towards end May when most sites have it working for OPS to see where we are with tests for experiments.

Other sites may also join at this stage but current focus is on Tier1s. In US-CMS glexec has been in use for a much longer time - in Europe this is new!

Data Management & Other Tier1 Service Issues

Storage systems: status, recent and planned changes

Site	Status	Recent changes	Planned changes
CERN	CASTOR 2.1.9-4 (all) SRM 2.8-6 (ALICE, CMS, LHCb) SRM 2.9-2 (ATLAS)	None	None
ASGC	CASTOR 2.1.7-19 (stager, nameserver) CASTOR 2.1.8-14 (tapeserver) SRM 2.8-2		
BNL	dCache 1.9.4-3		
CNAF	CASTOR 2.1.7-27 (ALICE) SRM 2.8-5 (ALICE) StoRM 1.5.1-2 (ATLAS, CMS, LHCb)		
FNAL	dCache 1.9.5-10 (admin nodes) dCache 1.9.5-12 (pool nodes)	none	none
IN2P3	dCache 1.9.5-11 with Chimera		
KIT	dCache 1.9.5-15 (admin nodes) dCache 1.9.5-5 - 1.9.5-15 (pool nodes)		
NDGF	dCache 1.9.7		
NL-T1	dCache 1.9.5-16 (SARA), DPM 1.7.3 (NIKHEF)		
PIC	dCache 1.9.5-15	xrootd doors enabled and published (request from LHCb)	none
RAL	CASTOR 2.1.7-27 (stagers) CASTOR 2.1.8-3 (nameserver central node) CASTOR 2.1.8-17 (nameserver local node on SRM machines) CASTOR 2.1.8-8, 2.1.8-14 and 2.1.9-1		

	(tape servers) SRM 2.8-2		
TRIUMF	dCache 1.9.5-11 with Chimera namespace		

Other Tier-0/1 issues

CASTOR news

Nothing to report.

dCache news

Nothing to report.

StoRM news

LFC news

The production version of LFC is now 1.7.3.

FTS

Experiment issues

WLCG Baseline Versions

- Release report: deployment status wiki page
- WLCG Baseline versions: table

Conditions data access and related services

Frontier/Squid

- The minutes of the last meeting can be found at the usual URL:ATLAS weekly FroNTier meetings
- Release 2.7.STABLE9-3 of frontier-squid has been announced. The release notes can be found here. The relative rpm has been made available for tests on Tuesday this week. Feedback received from BNL and CMS and integrated. A new rpm release will be announced soon.
- Squid caches are needed at CERN to alleviate stress on launchpads at other sites (namely Lyon). Information requested about the number of batch slots allocated to ATLAS and CMS analysis jobs since the number of needed squid caches depends on the number of slots. Squid caches at CERN will be installed for ATLAS by the VOC as soon as this information and the new rpm will be available.
- Squid caches can be installed on VMs provided that the physical machine hosting the VMs comes with multi-Gigabit network connectivity (1Gb/sec-link per Squid).
- Dave Dykstra requested more resources to monitor Squid and Frontier launchpad in ATLAS. The request is being put forward by the ATLAS VOC.
- Squid caches information will be stored in the ATLAS AGIS. Details on how to extract information from AGIS will be made public by the AGIS developers.
- CNAF have asked if they should install a frontier server for ATLAS or just squid caches. The recommendation is to install squid caches. CNAF has already 2 squid caches for CMS installed.

COOL and CORAL

- The LFC read-only instance at CERN for LHCb was unreachable on Tuesday timing out all requests and causing many jobs to fail. This is again due to the sub-optimal use of LFC in the Persistency replica service component. The problem is known since a long time and had been avoided with a workaround for production jobs, but it reappeared this week in the analysis jobs submitted by individual users. Various actions have been taken in parallel to mitigate and eventually fix the problem:
  - A workaround has been deployed by LHCb on Wednesday to avoid LFC access from user analysis jobs submitted through the DIRAC backend of Ganga. If necessary, this might be extended next week to the whole LHCb software environment (including interactive jobs).
  - An SQLite snapshot produced on Thursday with all conditions taken so far will allow users to analyse the LHCb data collected before the LHC stop, bypassing the access to Oracle and hence to the LFC replica service.
  - A Persistency patch prepared last week has passed preliminary tests on Wednesday and will be tested more thoroughly next week by LHCb when the relevant experts are back, in view of its release and deployment.
- A new release of COOL, Persistency and POOL (LCGCMT_56f) was prepared for ATLAS last week. The main motivation for this new release was to pick up some bug fixes and enhancements in the POOL collections package. Several bug fixes and improvements in Persistency and COOL were also included. The release notes are available on https://sftweb.cern.ch/persistency/releases.
  - Some problems with hanging connections in Persistency have been reported by ATLAS on Wednesday during the validation of the LCGCMT_56f release prepared last week and are currently being investigated.
- Two patches have been received from Oracle Support to fix issues reported in the 11.2.0.1.0 client software. The patch for the first issue ('cannot restore segment prot after reloc' when loading the 64bit OCI library with SELinux enabled) has been fully validated. The patch for the second issue (crashes in ATLAS production jobs on AMD Opteron quadcore nodes), which had triggered a downgrade to the 10g client for ATLAS a few weeks ago, has passed tests by the Persistency team on an ATLAS node in Ljubljana, but is still pending a more complete validation by ATLAS. A new client software installation '11.2.0.1.0p1', including these two patches and a third one previously received for the 32bit OCCI library on SELinux, has been prepared in the LCG AA software installation area in AFS.

Database services

- Experiments reports:
  - ALICE: ntr
  - ATLAS: A new version of the job responsible for cleaning up the DB audit table (usermon) has been developed, tested and deployed into production. Previous version of this job combined with high activity of atlas_t0 service caused transient performance problems on atlas offline cluster
  - CMS: ntr
  - LHCB: Intervention on the main controls router

- Tier0 Streams: On 7th of April ATLAS replication to CNAF suffered from apply process failover bug as there was rolling intervention without stopping of the apply process. The stream was split from the main replication in order to resynchronize missing gap and will be merged in the nearest time.

- Sites status:
  - RAL: Upgrade of OS kernels is in progress
    - No news about required licenses.
  - Gridka: ntr
  - SARA: network intervention scheduled on 20th of April. Whole cluster will be stopped.
  - CNAF : Migration of ATLAS database has been postponed until end of April.

    - ◊ ATLAS conditions replication will be merged back with main one
- TRIUMF: ntr
- ASGC: Problem with archived logs will be solved next week
- NDGF: ntr
- IN2P3: crash of one node due to memory problems. Second instance of DBAMI and DBATL were affected. Node is up again but the root of the problem is unknown.
- PIC: ntr
- BNL (Carlos): BNL agents has been patched with latest PSU

AOB

-- JamieShiers - 30-Mar-2010

This topic: LCG > WLCGTier1ServiceCoordinationMinutes100408
Topic revision: r24 - 2010-06-11 - PeterJones