

T1 reports – 05/03/2007

ASGC

CPU

- * requested job priority configuration is

Storage

* Castor2 preproduction service nodes have been deployed and couple disk servers will be deployed this week.

- * This will be provided for CMS for testing.
- * remaining disk servers will be deployed after additional LSF licenses are procured

Atlas

- * DDM AOD replication error fixed on Feb 26th. This required the following action to recover:
- * LFC acl were updated (/grid/atlas have wrong acl configured)
- * DQ2 agents restarted (help resolving the cert error observed in subscription logfile).

Event Log

Mar 02, 2007

- * castorsc disk server, castor001.grid.sinica.edu.tw with high load
- * restart forcely at Fri, 2 Mar 2007 03:13:36 GMT
- * simple validation pass at Mar 02 12:13:58 CST 2007
- * [lcg00122] ~> srm-get-request-status httpg://sc004.grid.sinica.edu.tw:8443/srm/managerv1 978126474
- * submitTime :Fri Mar 02 12:13:58 CST 2007
- * startTime :Fri Mar 02 13:13:58 CST 2007
- * finishTime :Fri Mar 02 12:14:06 CST 2007

FNAL

This Saturday, March 3, bandwidth for Fermilab's offsite network access will be significantly reduced during the day, due to hardware upgrades for the optical network equipment that supports most of the Laboratory's wide area network connectivity. The upgrades are an initial step in enhancing the Laboratory's offsite network access infrastructure. The maintenance period will start at 8:00am, CST, and is estimated to last until 5:00pm. During the maintenance period, offsite network access will be carried on a 622Mb/s link, instead of the 10Gb/s links normally utilized. Once the upgrades are completed, normal offsite access capacity will be restored.

CERN

Networking services

On Tuesday the 13th of March, a new firewall system will go into production. At 7h30 the old firewall will be put off-line and the traffic will move to the new system. This will cause a short instability for the external network connections. Until the 19th of March, the new system will be carefully monitored in order to quickly detect any anomaly. Users are invited to contact extip@cern.ch in case they notice any connectivity problem with external destinations. In order to reduce the number of variables during this monitoring phase, from 12h00 of Thursday 8/3/07 to 9h30 of Monday 19/3/07 only critical firewall changes will be made.

=====
Fabric Services
=====

- An issue with rootd requiring home directory to be available on the disk servers for grid pool accounts was resolved by the ROOT team who provided a new RPM. LHCb are testing the new version before on their castor instance, we hope to deploy it more widely next week

- We have liberated ~30 disk servers to allow Western Digital and Transtec engineers to upgrade disk firmware. Upgrades for another

~90 servers will need to be planned for the coming weeks.

- A hardware problem on the diskserver used by NA48 as their CDR/reprocessing server is currently under active investigation. The server will be taken by Elonex for refurbishment, and we hope to avoid data loss.

- A configuration update on the information provider for the SRM endpoints (where we added an optional field to satisfy a failing SAM test) unfortunately introduced a trailing white-space where there should have been none... Therefore, no SRM information was published for 2 hours on Thursday afternoon.

- Atlas tier-0 challenge started.

- the configuration and infrastructure for reporting back VOMS group/roles back to the grid using the information providers is ready for deployment. We currently support the software, lcgadmin and none roles. These roles are mapped to special grid pool accounts which are part of LSF groups with dedicated shares. Awaiting green light...

- Since last week end we have seen occasional high load alarms on lxmaster01. These were affecting all batch related activities, including ATLAS T0 activities and job submissions from the CEs. The origin of the high load could be traced back to two non-Quattor managed machines which are used for Grid software evaluation. The information system on these nodes went berserk and hammered the batch system with useless requests. With the help of GD the services on these nodes were stopped which fixed the problem.

- Batch Capacity as of today, in state "production":

1) ATLAS T0 dedicated resources:

HW model	slc3	slc4	Sum
nodes (Ksi2k)	nodes (Ksi2k)	nodes (Ksi2k)	nodes (Ksi2k)
Elonex_2.8GHz	10 (20.74)	0 (0.00)	10 (20.74)
E4_NOC_2800	8 (16.59)	10 (20.74)	18 (37.33)
e4_07_1	0 (0.00)	170 (2048.16)	170 (2048.16)
nc_07_1	0 (0.00)	35 (421.68)	35 (421.68)
ex_05_1	5 (10.37)	0 (0.00)	5 (10.37)
Sum	23 (47.70)	215 (2490.58)	238 (2538.28)

2) Public Batch resources available on the Grid:

HW model	slc3
nodes (Ksi2k)	nodes (Ksi2k)
Transtec_2.8GHz_64	13 (26.96)
Elonex_2.8GHz_64	3 (6.22)
Elonex_2.8GHz	447 (927.08)
SEIL_2.4GHz	589 (1003.66)
SEIL_1GHz	45 (41.49)
E4_NOC_2800	334 (692.72)
tt_06_3	9 (104.80)
mI_06_1	2 (23.29)
cv_06_1	359 (746.72)
ex_05_1	135 (279.99)
Sum	1936 (3852.92)

3) Public Batch resources on SLC4 (all 64bits architecture except a few boxes):

```

HW model ||      slc4 |
|| nodes (Ksi2k) |
-----+-----+
Transtec_2.8GHz_64 || 6 ( 12.44) |
Elonex_2.8GHz_64 || 74 ( 153.48) |
Elonex_2.8GHz || 3 ( 6.22) |
E4_NOC_2800 || 2 ( 4.15) |
e4_07_1 || 30 ( 361.44) |
se_07_1 || 162 (1951.78) |
nc_07_1 || 163 (1963.82) |
ex_06_9 || 45 ( 486.00) |
ml_06_1 || 5 ( 58.22) |
ex_05_1 || 51 ( 105.77) |
-----+-----+
Sum || 541 (5103.32) |
-----+-----+

```

Grid deployment

- Total number of Sites (*): 229

- Software -> Num. Sites (*):

```

gLite-3_0_2 -> 194
gLite-3_0_1 -> 2
gLite-3_0_0 -> 6
LCG-2_7_0 -> 14
unknown/down -> 13

```

- Status -> Num. Sites (*):

```

ok -> 161
degraded -> 9
down -> 59

```

- Average of concurrently running jobs during this week (+): ~30k

(*) Sites that are Certified _and_ Production _and_ Monitored by SAM:

[12]<https://lcg-sam.cern.ch:8443/sam/sam.py>

To see this page one needs a grid certificate loaded in the browser.

The calculation of the Site availability (Status) is described at:

http://goc.grid.sinica.edu.tw/gocwiki/SAM_Metrics_calculation

(+) Job statistics taken from GStat:

[13]<http://goc.grid.sinica.edu.tw/gstat/>

[14]http://goc.grid.sinica.edu.tw/gstat/total/GIISQuery_Usage_cpu_.html

For the time being we do not report CPU numbers:

1. Not all the reported CPUs are actually available for grid jobs.
2. Sites with multiple CEs may have their CPUs double-counted.
3. GStat includes sites that are not considered by the SFTs.

CERN Grid Operations managed by GD:

* Middleware upgrade (update 15 for gLite 3.0) on all the glite WMS 3.0.

* One new LCG RB put in production for SAM.

* 3 UIs still managed by GD-OPS have been removed from production definitively.

* Jobs submission blocked on all the gdrb nodes (LCG RBs with old hardware).

WLCG Services for Dress Rehearsal:

* The list of services required for the experiments dress rehearsals

starting in summer have been agreed upon. These need to be in place by Apr/1 to allow proper preparation.

WLCG Transfer Service:

- * Transfer ranging from 60 to 200 Mb/s, averaging around 100MB/s per day.
- * Involving all major T1 sites except ASGC (Castor upgrade).
- * Mostly traffic from CMS, with Atlas switching on in the middle of the week.
- * Little activity from Alice and LHCb.
- * dTeam tests to stress test FTS2.0 continue.
- * The Tier-2 FTS service has now been split away from the Tier-0 export service: expecting experiments to migrate by 1 April.
- * Around 15 tickets have been submitted this week for a variety of problems.
- * All problems detected by the transfer service on tier-1 SRMs are now being actively followed-up using the EGEE operations infrastructure.
- * Throughput plots:
[15]<http://gridview.cern.ch/GRIDVIEW/>

GridView :

- * Completed the validation of data for the newly deployed Web Service Transport for Gridftp data at CERN. Fixed a few bugs in the server module. Now it is working perfectly ok, without any data loss.
- * Deployed the Server for Web Services based Transport in production environment, now it is archiving CERN Gridftp data in production Oracle Database.
- * We had discussions with Users and internal discussions in team and formulated a way in which VO specific test results could be displayed in Gridview.

gLite 3.x Integration and Build:

- * Certification repository
 - .. Details can be found at:
[16][https://savannah.cern.ch/patch/?group=jra1mdw&func=browse&status_id\[\]=1](https://savannah.cern.ch/patch/?group=jra1mdw&func=browse&status_id[]=1)
[17]&resolution_id[]=109&resolution_id[]=111&resolution_id[]=102
[18]&resolution_id[]=103&resolution_id[]=108&report_id=183
- * PPS repository
 - .. Two new patches released to PPS on Monday as 3.0.2 PPS Update 21
 - .. Six new patches scheduled for next update on Monday for the moment
 - .. Details can be found at
[19]https://savannah.cern.ch/patch/index.php?go_report=Apply&group=jra1mdw
[20]&report_id=186&summary=&custom_sb1=104
- * Production repository
 - .. Set of two new patches have been released to production as 3.0.2 Update 15 on Monday:
 - .. Next set of patches to be released next Monday:
 - .. Details can be found at
[21]https://savannah.cern.ch/patch/index.php?go_report=Apply&group=jra1mdw
[22]&func=browse&report_id=188&release=3.0.2+Update+09

gLite 3.x Certification and Testing:

- * Patch overview

.. Certified
#996 (edg-mkgridmap 2.8.1 more robust)
#1003 (Dcache 1.7 upgrade)
#1013 (WMPProxy does not create all the needed job directories)
#1014 (UI Python commands)
#1036 (gLite WMS respects \$EDG_WL_SCRATCH)
#1045 (DPM/LFC 1.6.3)
#1045 (Matchmaking fix for a field containing a ".")
#1051 (yaim 3.0.0-38)
.. Rejected
#969 (glite-yaim-3.0.1-7 update)
#1034 (UI configuration bugfix)
#1027 (aggregate minor bug fixes)

* SL4

.. Installation and configuration testing of the natively built SL4 gLite 3.1 WN. Still chasing problems with packaging and dependencies.
.. Update to PPS to fix the problem with glite-WN_sl4compat not being maintainable after installation.

* Tests

.. Further UI tests written, including for lcg_utils and tcsh support
.. glite-CE tests revealed a problem with job cancellation ;
[23]https://savannah.cern.ch/bugs/index.php?func=detailitem&item_id=24209

* Configuration

.. Configuration for dynamic-info-scheduler plugins for lsf and condor being worked on.
.. Work has started on bash configuration for the glite-CE and WMS, intended for gLite 3.1.

* Other Work

.. Maintenance of experimental 3.1 WMS.
.. Production LB troubleshooting

Data Management:

* LFC/DPM 1.6.3 is now certified (supports bulk methods in client for LFC). LFC 1.6.3 has been installed on the CERN LFC nodes.
* FTS service split completed. We now have a separate service for T0/T1 and T0/T2 transfers. Creation of FTS T2 service (tiertwo-fts-ws.cern.ch) to be announced shortly for public consumption.
* Tests of the FTS using SRM 2.2 end points are continuing.
* Beta sites for dCache 1.8 (supports SRM 2.2) agreed: DESY, FNAL, IN2P3. Others to be added.
* Creation of FTS T2 service (tiertwo-fts-ws.cern.ch) complete to be announced shortly for public consumption.

ETICS:

* Most of this week has been dedicated to helping the gLite developers and integrators in producing WN and UI distributions natively compiled on slc3 and slc4/32 machines.
* Our focus will now shift for slc4/x86_64.
* Several new external packages have been registered in the Externals project for use by gLite components (db4, hibernate, struts, dCache-client, jaxb, derby, jxta, exist).
* A new project has been registered in the ETICS db: GridICE
* Support tickets received: 8, details can be found in <https://savannah.cern.ch/support/?group=etics>

* A new server is being installed at CNAF to deploy temporarily the ETICS 1.0 RC, until suitable hardware can be found at CERN. Several of our nodes are reaching or have reached their end-of-operation deadline and we are experiencing more and more failures (mainly hard-disks and power supplies).

Physics Database Services

The migration to new hardware and extension of the production RAC for the LHCb experiment from 4 to 6-nodes was carried on last Monday morning from 8.00 to 11.30.

The upgrade to the ORACLE server version 10.2.0.3 was also performed at the same time.

With this intervention, we have completed our hardware extension plan and doubled the capacity of all the experiments/grid RACs. We have also uniformly upgraded our hardware to RHEL 4 and to the latest ORACLE release 10.2.0.3.

We observed more occurrences of the ITRAC servers hangs for the CMS and WLCG production RACs (all our RACs have been upgraded to RHEL 4). The workaround proposed by ORACLE is to change the default value of the `_high_priority_processes` parameter, modifying the priority of the cache fusion processes (LMS0,1).

We already applied this configuration change to the ATLAS production RAC two weeks ago. We will now apply this change to all our production RACs next week in a rolling (transparent to the user) way.

The 3D OEM repository has been migrated to the new hardware.

New sites have been added to the LHCb COOL streams setup: CNAF and SARA (the configuration using now 5 destination sites). IN2P3 has installed new databases for the ATLAS and LHCb experiments. The streams configurations have been modified to include the new databases.

Streams propagation reporting ORA-600 KWQPCBK179 (propagation jobs get stuck - Bug5529797) have been found on the ATLAS Streams setup using downstream. Oracle will release a patch in about two weeks.

ARDA/EIS

ALICE:

Production ongoing. Issues with the gLite-CE found and reported. Finished to implement the SAM tests for the VOBOX.

ATLAS:

Installation and configuration of ATLAS agent nodes at CERN for the T0 throughput exercise (started Monday at low rate). Foreseen to move the ATLAS central catalogues service under IT services (load balanced). Major changes in the Production System framework carried on. Some changes in the ATLAS monitor (users feedback and fix of minor issues)

LHCb:

~10,000 concurrent Grid jobs this week, 30% at CERN. Reconstruction running OK at all Tier-1 and CERN. VOBOX transfers working at full regime.

CMS:

still testing gLite 3.1 WMS, problems due to frequent locks of the daemon which logs events to the Logging&Bookkeeping.
Survey of CE accessibility, time limits and consistency with published information for CMS tiers. Testing job submission to OSG via LCG RB, also via SAM.

Dashboard:

More views for Grid reliability being added (after discussions with some Tier1). Monthly report format being discussed.

Other development:

VOMS proxy renewal client is being integrated into the VOBOX middleware. Currently only renewal on non-VOMS proxies is supported by the VOBOX.

Tier1 GridKa (FZK)

Main issues this week were:

on 24/2/07 the SRM part of dCache became unresponsive for many hours.
glite to update 14 lead to some CE info system instabilities on 1/3/07
again SRM instabilities on 1/3 which were mitigated with a dcache fix. Next week will tell if this improves the situation

SARA

Around 10:00 misconfigured ACLs resulted in replication failures.

RAL

dCache SEs has been upgraded to 1.7

PIC

-lcg-CE service:

+A bad configured acted as a black hole. Some of the jobs that ran in this particular node didn't succeed to finish properly. The problem itself was due to the ssh keys that were incorrect so the job output was not properly transferred to the Computing Element. The problem is now solved after having reconfigured the node.

+We still have a quite serious problem with pbs "status=Waiting" jobs at pic. We have asked for help all around but it seems it's a known bug which we will hope it will be solved in of torque's next release.

+We have closed one of our 3 CE head nodes (ce04.pic.es) to all the VO's except for the monitoring ones but we have continued to receive jobs. Some vo's like lhcb have white lists that ignore the information we publish and they submit jobs to unpublished resources. We had an open ggus ticket for this. This ticket was immediately closed (ggus 19135)

-LFC service: There was a problem with PIC's pre-production central catalog that was published under CERN's production top BDII. It seems the problem was not ours. All the PPS sites were published under Cern's top-bdii, including PIC. Since we have the central catalog at PIC's preproduction site there two central catalogs published at the same time. Judit Novak at cern is solving the problem.

-Top-Bdii service: We still get some errors in our top Bdii. It seems applications like lcg-cr commands fail while pointing to our top bdii. This causes GSTAT tests to fail from time to time (finding zero entries). We are investigating it. The problem is serious because a lot of sites in the region use our top-bdii. We must say instead that we have almost 100% of availability on the top-bdii sam tests, so we do not understand wether some of the failures are false negatives. Problem under investigation.

-SRM-tape service: no issues.

-SRM-disk service: more stable this week, did not have to restart gridftp doors that often. But the usage was not very high neither.

