CERN-PROD Tier0/1 Site Report

DSS

A common agreed configuration for the DSS Back End software - the user interface part - is being prepared and will shortly be installed in all four LHC experiments.

Networking services

The deployment of AAA (pilot) on the network equipment in b31 and the computer center is now complete. New authentication servers have been installed on the GPN and on the Technical Network to prepare the global deployment. Before this global deployment can occur, we will need to perform a general firmware upgrade of all the network devices on the infrastructure (> 1'000 devices). That will probably occur in December or very beginning of January when the accelerators are not running. When this deployment is complete, the Technical Network will enforce all the MAC addresses to be correctly registered before a system can be connected to this network.

The issue discovered in the way the LCG backbone is handling its load distribution has been understood and a fix has been deployed (no restart of any equipment was required, so no interruption occurred). Since then, the load is very well distributed.

Three LHC experiments are now using intensively the SOAP interfaces to LANDB to document their topology in the Pits, making our operation must more efficient.

We are now participating in a Beta testing of a new HP firmware which integrates a number of new features CERN requested.

Database Services

* 10.2.0.2 bug

A first discussion on a post-mortem analysis was held on Tuesday. A number of points have been identified by CERN and have been transmitted to Oracle. A report will be written.

Fabric Services

CASTOR operations:

* Monday Oct 2 was a busy day:
  - As announced, we decommissioned the SRM endpoint castorgrid.cern.ch, leaving it as a classic SE. We also downsized it to 4 nodes, liberating an IP service.
  - After a one-week test period the Castor client software was upgraded to 2.1.1-1 on the CPU servers (lxplus, lxbatch, etc). The Linux desktops were upgraded during the following night.
  - We upgraded the CASTORLHCB setup to the latest stable Castor server release, 2.1.0-6. The Atlas and CMS setups have been running his version since one week earlier.
  - The configuration of the CASTORCMS setup was updated for the CSA06 exercise. The t0export diskpool now consists of 32 diskservers, providing

150 TB of diskspace.
* We are awaiting a new Castor server release (2.1.1-4), containin support for repack-2 and xrootd. This version will first be installed on the C2TEST validation setup, before being rolled out to CASTORPUBLIC.
  This new version requires us to rebuild the old SRM v1.1 and castor-gridftp software. We are in the process of rebuilding and testing this.
* The setup of an SRM v2.2 pre-production service has started. Together with the Castor development team we are defining a setup, that we plan to deploy by Nov 1.
* We have started to validate the complete software stack needed to run Castor2 diskservers on SLC4 in 64-bit mode. This stack is quite deep; it includes XFS filesystems, the Castor software, gridFTP, VDT, glite, RGMA. Thanks in advance to all those who will be helping us :)
* Both CMS and ATLAS data challenges run OK so far. The atlas challenge has ramped up to relatively high rates and load on the stager, which has revealed some new minor bugs/problems in castor2.

Batch:
* Service re-numbered: IP37(lxb06**) and IP39(lxb12**). IP56 (lxb16** batch) draining. IP52(lxb17xx), IP53(lxb19xx), IP54(lxb20xx), IP57(lxb18xx) are next. Mostly GD machines.
* CMS CSA06 has now 600 CPU boxes. Will be increased to 700 by next week. LSF CMS shares decreased from original 1000 to 200 KSi2K's.

Quattor:
* SLC3/SLC4 software upgrades from Linux team. Glitches regarding ssh configuration caused access problems on lxplus/lxgate for a short time.

Fabric Developments

CASTOR:
* A lot of support time on the CNAF problems (7 days total) and Alice DB
* Repack 2 is now packaged and released. It will be tested next week by the operation team at CERN
* gridFTP v2 as a castor internal protocol is progressing. A prototype is ready. SRM v2 will be modified to test it next week

Quattor:
* the latest CDB is under testing on lxdev23

LEMON:
* RAC hardware has been ordered by DES for Lemon. Should be available early in November
* Preparing for the Lemon tutorial next Monday
* Deployment of the latest version12.0-1 of OraMon on production servers
* optimization of sensors/MSA
* Remedy switch to use lxremedy01 was transparent to Lemon. No issues has been found.
* Issues:
  - with high CPU utilization on the primary Lemon Database server still not understood. Turned off nightly DB statistics gathering job to avoid further NO contact alarms over night. Working with oracle support on finding what is the root cause of it

Remedy:
* Operations contract changes completed and in production
* Currently developing reports (thanks to help from FIO-TSI)
* New Linux-based web server using ARS Mid-Tier v7 now used production

since Friday - No significant problems
* Alleviating load on main Remedy server
* Simplifies upgrade plans for ARS v7
* GGUS Interface
  - A bug, now fixed, prevented updates to GGUS being sent. Unsent notifications were reprocessed
  - sunar01 was revoked access to the technical network, breaking the access control application. It was repaired within hours
  - Begun (re)configuring the failover server prior to testing
  - HMS->SMS interface modified to work with non-quattor managed machines to reduce spurious operator alarms

AFS
---
* afs97 got stuck, no useable core dump, affected home directories. Will get replaced by the newer afs34. Also some similar issue on Solaris (afs31) which should allow to debug this.
* AFS monitoring console now reachable under [12]http://cern.ch/afs-console/

Linux
-----
* sec updates on SLC3/4, see
http://cern.ch/linux/updates/updates-slc3.shtml#02.10.2006
http://cern.ch/linux/updates/updates-slc4.shtml#04.10.2006
http://cern.ch/linux/updates/updates-slc3.shtml#04.10.2006
http://cern.ch/linux/updates/updates-slc4.shtml#04.10.2006

* CASTOR-2.1.1-1 is now the default production version on desktops

* major openssh update on SLC3 to get everybody Kerberos5 credentials caused trouble:
- desktops: some partly-configured machines needed manual intervention
- CDB machines: needed to explicitly configure the new 'default' settings on several clusters (PLUS/GATE/..). We noted that several service managers apparently had copied the LXPLUS template, including special SSH settings that made sense only there... (this points to a more general problem with our CDB template hierarchy, there is no clear 'inherit this if you need that' structure)
- Grid-wide: we killed the Grid (at least temporarily). Apparently >50% of LCG sites use SLC3, and lots of them subscribe the production systems to our desktop updates without further verification on the site. The job submission mechanism used a bug in the old ssh (locked Grid accounts could still login via RSA keys) that no longer worked (as such) in the new version.

This caused some fallout on the LCG-rollout list. We followed due process with the announcements. Following up through IT-GD (thanks to Maite Barroso Lopez / Romain Wartel)

* ATLAS distccd processes are segfaulting, investigation ongoing
* updates to the printer config system to cope with changes in the paper size/duplex capabilities.

Operations and System Administration
----
* A power test in the critical area will take place 05/12 at 07:30. The

physics/normal distribution will be cut for a few minutes in order to check that all equipment in the critical area is properly connected to the critical distribution and check how the critical UPS reacts to the load increase.
* New Operations contract started without any issue.
* 85/85 Transtec Mid-range servers received and being physically installed. Stress tests running on ~50 nodes. Stress tests on other nodes, installed in RAC4, should soon start there after the cabling of RAC4 is finished.
* 24/24 Transtec disk servers received and installed. Stress tests should start before the end of the week.
* 75/115 Elonex mid-range servers received and being physically installed. 37 servers physically installed, 17 severs being stress tested. Nodes arrive in small groups.
* 16/112 Elonex batch servers received, 1 installed and being stress tested. Delays due to disk unavailability.
* 26/26 Elonex disk servers received. All nodes installed, stress tests starting.
* The move of Linux TSM is took place 05/12. AIS TSM move will take place the 17th of October.

Service Challenges

The total outgoing SC4 traffic has ranged from 300 to 700 MB/s as daily average. Alice often doing 100 MB/s or better for a few hours, followed by periods of little activity. CMS have started their CSA'06 challenge on Monday. At the same time Atlas ramped up to a new phase in their SC4 activities and have been producing the majority of the outgoing traffic.

One of the three nodes behind the FTS web service alias managed to fill up the memory allocated to tomcat, causing intermittent misbehavior. We can give tomcat a larger share for now, but there may be a leak that will have to be fixed eventually.

The GridView plots temporarily suffered from negative file sizes being published in the transfer statistics: the familiar 32-bit signed integer overflow, with the usual work-around, viz. adding 2^32. It is clear, however, that the publishers must be fixed to deal with files > 4 GB.

Security report

* All production VOMS services have been upgraded to latest OS and middleware versions.
* Stress tests are being run against test Oracle DBs to determine if latest Oracle patches fix data corruption Oracle bug.

CERN CA

* The CERN Root CA certificate was found on Monday to have a subject different from the one published in the CP/CPS policy document. This error required to re-issue the CERN Root CA certificate and to revoke all previously issued certificates, affecting 50 users and 20 hosts. The service was back to normal at 2pm, however affected users need to request a new certificate manually. This error was corrected in time to finalize Grid approval for CERN CA foreseen this week.

Physics Database Services

The patch for the 10.2.0.2 critical bug has been installed in a rolling way on most of our production RACs (the intervention on the WLCG RAC is scheduled on Monday 09.10, the one on the CMS RAC is postponed as they have just started the CSA06).

We have prepared some draft proposals on database security and backup/recovery policies to review with the experiment and grid deployment teams and be presented at the SCM, 3D and MB meetings. More info on our wiki pages at
https://twiki.cern.ch/twiki/pub/PSSGroup/PhysicsDatabasesSection/ORACLE_backup_policy.doc
https://twiki.cern.ch/twiki/pub/PSSGroup/PhysicsDatabasesSection/ORACLE_database_security_policy_at_Tier_0.doc

LCG 3D project

* 3D Status presented as part of LHCb software week
* The inital streams setup for the LHCb experiment is completed. Streams test from online->offline->T1 now successful after some configuration problems on their (still temporary) online database setup.
* A streams queue corruption in the ATLAS setup after applying a critical patch has been observed and now fixed.
* Incompatibility between restored archive log files with RMAN+ASM and recovery of the capture process has been discovered. Working on a resoultion with Oracle Support.

EIS

ALICE: adding new sites in Mexico, Bratislava and Poland to the production.

ATLAS: validation of new ATLAS software release ongoing on LCG. Production system framework is being modified to ensure more flexibility in job distribution. ATLAS week

CMS: tuning the gLite workload management system configuration for CSA06. Job submission tools adapted to use the bulk submission of jobs.

HARP: the problems with the software installation have been found, now proceeding with it.

LHCb: production running smoothly; reconstruction and stripping phase ready to re-start; CNAF fully recovered now

Savannah

The last security update (on Thu morning) caused problems for users with AFS authentication to login to savannah. These problems were fixed by reverting the /etc/pam.d/php file to the old version (using pam_afs instead of pam_krb5afs from /lib/security/). This was needed as users do _not_ have an entry in /etc/passwd on this machine.

The incidents reported last week did not reappear. In case they do, we do have a 'hack' to minimize downtime of the service.