

Site Report: CERN

Helge Meinhard / CERN-IT
HEPiX, GSC St Louis MO USA
06 November 2007

LHC: The Inner Triplet story

- March 27: Quadrupole broke during pressure test at CERN
 - Support structure holding cold mass in cryostat did not stand the longitudinal forces
 - Cryogenic and power distribution box damaged as well
- Review of repair actions on 24/25 April
 - Proposed repair actions can be performed underground, except broken inner triplet
- No statement yet concerning potential impacts on LHC startup schedule

From my talk in Hamburg

LHC schedule

- Following inner triplet problems, engineering run in autumn 2007 at injection energy cancelled
- First circulating beams in May 2008, ramping up to 7 TeV, physics as soon as beam is clean
 - No delay of LHC physics
- Open days in 2008 for public, VIPs, press, CERN families
- No impact on IT procurement planning

IT structure changes 01-Jan-2008

- Ian Bird succeeding Les Robertson as LCG project leader
- New Data Management group (Alberto Pace)
 - Physics databases and developments
 - Castor developments
 - Grid data management tools
- New Experiment Support and Services group (Jamie Shiers)
 - LCG service coordination
 - Experiment integration, distributed analysis
- Grid services operation moves to FIO

Database and Engineering Support

- CVS services
 - Looking at streamlining general services (AFS-based) and LCG services (local disks)
 - AFS too slow for LCG projects
 - Considering to offer a service based on subversion, and to stop cvs services
- RAID 5 is no guarantee
 - EDH blackout for three days – four disks died in a RAID array!

Fabric Infrastructure and Operations (1)

■ Infrastructure

- Power and cooling... again and again...
 - Insufficient envelope in CC by 2010/2011
 - Need a solution by then, even if ad interim
 - Various activities ongoing, no concrete plans yet
- Even more pressing problem is “critical power”
 - Backed by dual UPS, Diesel generators, total of 300 kVA
 - Current (well reasoned) requests exceed this threshold
- Incidents with cooling
 - Cooling units did not restart automatically after power glitch
- Simulations
 - Team wants CFM values for our servers...
- Cold aisle confinement started in main room

Fabric Infrastructure and Operations (2)

- Castor performance problems addressed
 - Task force terminated
 - Very significant improvement applauded by CMS and Atlas
- In LCG statistics of September, CERN had 100% availability
- TSM: 1 billion files
- AFS: More than 2 billion hits on a server per day
 - Flash disks would not have survived for longer than a day...
 - Client hit rate 2 orders of magnitude higher

Fabric Infrastructure and Operations (3)

- All machines (re-)installed under SLC4 that support it are installed in 64-bit mode
 - Support for 4+ GB of memory
 - Applications mostly, if not all, 32 bit
- Oracle applications run on “disk servers” previously moved to Oracle RAC with NAS boxes
- TSM backup service expanded
 - According to needs which by far exceeded projected requirements
 - IBM certified hardware
- Tape services expanded
 - Generally new robotics and drives go well, but quite some disruptive interventions due to hardware and/or software problems

Fabric Infrastructure and Operations (4)

- Procurements (1)
 - Big orders for CPUs and disk servers out, delivery in November 2007
 - 620 dual Clovertown 2.33 GHz (E5345), 16 GB of RAM, all as Supermicro twin boxes (2 systems in 1U chassis sharing a power supply)
 - Very power-efficient solution
 - Total about 6 Million SPECint2000 (LCG conservative measurement)
 - 280 disk servers, 5U, 22 500GB disks, two Clovertown 2.00 GHz, 8 GB
 - Total about 1.6 PB usable
 - Next round of orders being prepared, to be approved by Finance Committee in December
 - 6 Million SPECint2000, 3 PB usable

Fabric Infrastructure and Operations (5)

- Procurements (2)
 - Sizeable orders for dedicated servers (Oracle RAC, midrange, small disk servers, ...)
 - First big order out for >100 servers from a Tier1 supplier, more likely to come
 - Considering blades as a realistic option
 - Round of evaluation of boxes on loan
 - Procured a precise 8-channel power analyser
 - Technical student working on a procurement workflow system, about to be put into production

Fabric Infrastructure and Operations (6)

- Identified hardware problems (1)
 - Buggy firmware on disks causing disk dropouts and (sometimes) data corruption
 - One drive manufacturer, different models
 - Drive manufacturer flashed 5'500 drives, swapped 100 out
 - 600 1U and 2U machines suffering from high vibration level in chassis
 - Higher than normal disk failure rate
 - Chassis supplier coming on site to swap out the fan assemblies

Fabric Infrastructure and Operation (7)

- Identified hardware problems (2)
 - Bad batch of Nocona mainboards
 - 25 of 45 mainboards dead after a year
 - Supplier swapped out entire batch
 - RAID controller problems
 - NCQ commands on the bus even though NCQ switched off
 - Disks invisible after unclean shutdown
 - BBUs reportedly failing massively
 - Power supply problems
 - 10% power supply failures on a batch of servers from a Tier1 supplier within a few weeks
 - Supplier swapped the PSUs out

Fabric Infrastructure and Operation (8)

- Identified hardware problems (3)
 - Number of batches with bad BMC firmware, mainly wrong thresholds
 - In all these cases, collaboration of suppliers and manufacturers was between rather good and excellent
 - What if vendors react less well?
 - Seems that problems with batches are on the rise
 - Not convinced that Tier1 material much better than white boxes

Fabric Infrastructure and Operations (9)

- Linux support
 - Stopped support of SLC3 i386 on 31-Oct-2007
 - Community expressed no significant demand for SLC5
 - SLC4 is a big problem on laptops, rather impossible with new machines
 - Do other sites support Linux on laptops? If so, which distro?

Fabric Infrastructure and Operations (10)

- Miscellaneous

- Migrated to LSF 7 and at the same time to much more performant hardware for LSF master
- Now deploying IPMI fully
 - We require RMCP+ in tenders
 - No physical serial connections any more
 - Seamless integration with console service

Internet Services

- Office 2007 now default for new installations; Vista on offer
 - Only if HW supports it
- SharePoint deployed and widely used
- NNTP stopped
- NICE authentication now used for AIS, Twiki, ... as well
 - “CERN password”
- New alerter based on RSS feeds
 - Replaces zephyr
- Acrobat Reader saga
 - Version 8 not compatible with some PS drivers
 - Version 7 suffering from security vulnerabilities
 - Replacing PS drivers by PCL ones on affected printers

Computer Security

- Firewall: Additional TCP ports closed on 29-May
- Kerberos 4 phased out, moved to ssh v2 as default
- Following popular demand, evaluation of Skype Enterprise ongoing
- Botnets moving to p2p technology
- Security training for developers
- VPN to be closed in January 2008
- A number of compromised machines, most of which locally “managed”

Miscellaneous

- Acceptance of CERN Computing Rules can now be confirmed electronically
- Single sign-on making progress
- Single entry point for account-related interactions: <http://cern.ch/cernaccount>
 - Change account-related info
 - Change/reset password
 - Request new account
 - Manage external accounts
 - ...

Next HEPiX

- May 05 to May 09, 2008
- At CERN
- Accomodation and conference rooms reserved
- Starting to work on other organisational and programme items
- Hope to see you there... and to whitness the first beams in LHC together!