



CERN Site Report

Helge Meinhard / CERN-IT
HEPiX Spring 2009 Umea
25 May 2009

- CERN Directorate
 - Rolf Heuer (Director General)
 - Sergio Bertolucci (Research and Computing)
 - Stephen Myers (Accelerators and Technology)
 - Sigurd Lettow (Administration and General Infrastructure)
- IT Department
 - Department head: Frédéric Hemmer
 - Deputy: David Foster
- Two former IT groups changed department
 - CO (now in Engineering Department as Industrial Controls and Electronics)
 - AIS (now in General Infrastructure Services Department)

- First circulating beams in LHC: 10 September
 - Very large press coverage
 - Web servers hit 20 times more often than usually
- Incident on 19 September 2008
 - During attempt to ramp beam energy up to 7 TeV, a leak occurred in the cold mass causing significant loss of helium
- Repair work is ongoing
 - Instrumentation for detecting this kind of problem being added
- Schedule: beam end September 2009, collisions end October 2009, running until autumn 2010
 - Collision energy: 5 + 5 TeV
 - Short technical stop at Christmas 2009
- Work on schedule

- TWiki
 - Started at CERN in 2003
 - Now 5'000 active users, 44'000 topics, 30'000 monthly updates
 - Atlas and CMS are heavy users
 - Security enhancements: write access to CERN accounts only, campaign to disable anonymous read
- Subversion
 - “A better version control system”
 - Following successful beta test, went into production in January 2009
 - Coexistence with cvs for quite some time

- Linux
 - Certification etc.
 - SLC4 is current production version, mostly x86_64 on servers
 - Laptops problematic since 2006, desktops since a little later
 - SLC5 ongoing, formal certification delayed because of CERNLIB
 - Because of demand for SLC5, lxplus and (very sizeable) lxbatch facilities established, publicly open now
 - Recommendation for experienced users without any known dependency on SLC4 is SLC5
 - DNS incidents: Found that with RH5, the DNS lookup is conforming more closely with RFC 3484
 - No longer the first IP, or a random choice, but the “geographically next” IP is returned
 - Change of responsible: Jan van Eldik takes over from Jan Iven

- Infrastructure
 - New computer centre
 - New directorate very favourable towards CC project
 - Green solution: link new CC to heating of respective site (Prévessin)
 - 4 detailed design studies delivered in 2008
 - Tender for “client advisor”
 - Subsequently (unique) tender for construction
 - Total of 60 water-cooled racks of 10 kW each: Infrastructure ready
 - Delivery of CPUs and disks in September 2008 will use them
- Backup
 - Growth remains higher than expectations (~ 35%)
 - Users don't know what they want...
 - Alex Iribarren took over from David Asbury as service manager

- Procurements (1)

- Change of LHC running schedule required procurement planning to be adapted
 - Initially: 2 big rounds for 2009, 2 big rounds for 2010
 - After September 19th: Only one round for 2009, one for 2010
 - February: “Emergency” procedure added to have additional equipment in October 2009
- Tenders open or in preparation

Type	Production	Volume
CPU servers	October 2009	50'000 HEP-SPEC06 (about 500 systems)
Small disk servers	October 2009	50 systems
Large disk servers	October 2009	4 PB usable (about 200 systems)
Servers and iSCSI backends	October 2009	1 PB usable
Oracle servers	October 2009	90 systems
Midrange servers	> October 2009	50 systems
FC disk arrays	> October 2009	50 systems
CPU servers	March 2010	150'000 HEP-SPEC06 (about 1'500 systems)
Large disk servers	March 2010	10 PB usable (about 500 systems)

CERN CC currently (March 2009)

- 5'700 systems, 34'600 processing cores
 - CPU servers, disk servers, infrastructure servers
- 13'900 TB usable on 41'500 disk drives
- 34'000 TB on 45'000 tape cartridges (56'000 slots), 160 tape drives

- Procurements (2) – Fun points
 - Defective memory modules
 - Yet another supplier went out of business in January
 - > 100 systems in CC under warranty no longer maintained
 - Memory manufacturer insolvent
 - Late deliveries
 - Spurious error messages due to defective firmware
 - Almost 1'000 systems upgraded: BIOS and BMC
 - Disk drive problems: thousands of drives upgraded
 - drives dropped out of RAID under heavy load
 - I/O errors on disk drives
 - Corruptions on disk drives
 - Risk to make drives inoperable
 - Default expiration date
 - ...
 - RAID controller / SAS expander incompatibilities
 - Causing data corruptions, required firmware fix for RAID controller
 - I/O errors on disk array in JBOD mode
 - Dropped connections (or none at all) via BMC SOL
 - Unphysical (aka crazy) CPU accounting
 - Loss of network connection to BMC

- Tapes

- Following successful tests, drives upgraded to newest generation supporting ~1 TB, media being repacked
- After many years of dealing with tapes, Charles Curran retired
 - Vlado Bahyl has agreed to take over

- R&D

- Lustre pilot project (see Arne Wiebalck's talk)
- iSCSI tests (see Andras Horvath's talk)

- Windows services
 - Terminal services remain very popular, number steeply increasing (now ~ 60 in production, i.e. 40% growth over 2008)
 - Ivan Deloose taking management over from Michael Kwiatek
- Windows clients
 - Vista default for sufficiently powerful machines (e.g. 2 GB or more), but XP remains a choice. Issue on Technical Network because of activation
 - Michael Kwiatek taking over from Rafal Otto

- Mail services
 - Migration of mailing lists and archives to e-groups completed
 - Mailboxes being moved to Exchange 2007
 - Plagued by Thunderbird “bug” exhausting IMAP connections

- Preparing for re-tendering the desktop / helpdesk contract
 - New contract as of 01-Jan-2010
 - Current helpdesk workflow in Remedy being re-implemented until September 2009
- Printing under consideration
 - (Far) Too many printers, too many different printer models, high (but mostly hidden) cost
 - Proposal for small number of multi-function devices, pay-by-sheet

- Incidents include
 - Well-known attack on network protocol on Linux machines back
 - Increased phishing attacks – some users really give their passwords away...
 - Users posting sensitive information on open channels (mailing lists, Web including TWiki etc.)
- Measures taken include
 - Tor added to list of blocked applications
 - Access to remote DNS servers blocked
- Skype: Trial ongoing under certain conditions
 - “Tolerated as long as configured correctly and not more problems than expected occur”
- Training courses for users and developers well received

- Information Technology Infrastructure Library
- CERN-IT getting serious about it...
- In some places not far off, but rigour sometimes missing
 - Meetings are working well, but invested effort can't be compared with measurable improvements of services
 - No clear separation yet of change management, incident management, service level management
- Hope is to reduce stress and 'heroism' to run services
- KPI difficult in our environment
- Aim is to demonstrate that best practices are applied
- IT Service Review meeting being instantiated (role of DTF changes)
 - Need to unify approach of service review, avoid multiple meetings

- CERN Openlab
 - Phase III started
 - Siemens joined as a partner for automation and controls
 - Oracle, Intel and HP renewed membership
- SURE monitoring retired after 15 years