



CERN Site Report

Helge Meinhard / CERN-IT
HEPiX Fall 2009 LBNL
26 October 2009

- Schedule: beam November 2009, collisions a few weeks later, running until autumn 2010
 - Collision energy: initially 3.5 + 3.5 TeV, later 5 + 5 TeV
 - Short (3 weeks) technical stop at Christmas 2009

- Getting serious about ITIL
 - Several working groups
 - ‘Best practices’ approach
 - Were not too far in many respects, but rigour and coherence missing
 - Dedicated talk by Tony Cass
- Solaris phase-out scheduled for 2010
 - Only Linux and Windows remain

- Current GSM contract extended by 2 years
- Network service desk (NetOps) transferred to IT helpdesk

- TWiki
 - Upgraded to version 4.3.1
 - Remains very popular
- Version control (Subversion)
 - Dedicated talk (Giacomo Tenaglia)
- Oracle 'fun': Big ID problem took months to fix
 - Castor services hit hard (not only at CERN)
- Interested in virtualisation for some services

- Testing 10GE connections for database backup
 - Potential alternative to lanless backup over dedicated fibres
- New Castor releases rolled out (2.1.8, 2.1.9), but deployment slow, in particular outside CERN

- Linux
 - Kernel problems around CPU timers
 - ‘Crazy’ (i.e. unphysical) CPU counters: understood and fixed in latest released RH kernel
 - Wrong user vs system counting: No fix yet. Probably due to register handling (does not happen with newer vanilla kernels), but backport difficult
 - Some minor issues with FC HBA and NIC drivers

- Infrastructure
 - New computer centre
 - Put on hold awaiting complete evaluation of alternative solutions (temporary structures, hosting nearby or far away)
 - Issues: data privacy, licensing, on-site maintenance, ...
 - See Alan Silverman's talk
 - Total of 60 water-cooled racks of 10 kW each basically ready
 - Recent delivery of CPUs using them
 - Uplift of total power: 400 kW (limited by UPS capacity)
- Backup
 - Will probably stay with IBM TSM
 - Following serious negotiations with IBM and careful evaluation of potential alternatives

- Procurements (1): Tenders open or in preparation

Type	Production	Volume
CPU servers	November 2009	50'000 HEP-SPEC06 (about 500 systems)
Small disk servers	November 2009	38 systems
Large disk servers	December 2009	2.7 PB usable (138 systems)
Servers and iSCSI backends	November 2009	1 PB usable
Oracle servers	November 2009	90 systems
Midrange servers	December 2009	116 systems
FC disk arrays	January 2010	About 100 systems
Tape servers	January 2010	60 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 (October 2009)

- 6'300 systems, 39'000 processing cores
 - CPU servers, disk servers, infrastructure servers
- 13'900 TB usable on 42'600 disk drives
- 34'000 TB on 45'000 tape cartridges (56'000 slots), 160 tape drives
- Tenders in progress or planned (estimates)
 - 2'400 systems, 16'000 processing cores
 - 19'000 TB usable on 20'000 disk drives

- Procurements (2) – Fun points
 - Defective memory modules
 - Wrong, inconsistent rails
 - Temporary disk kits
 - Chassis and/or mainboard manufacturer late
 - New PSU hardware revision doesn't work okay with mainboard (same manufacturer!)
 - iSCSI appliances need firmware upgrade before being usable
 - iSCSI appliances don't allow to configure the entire disk space
 - Disk drive problems: thousands of drives upgraded
 - RAID controller / SAS expander incompatibilities
 - Causing data corruptions, required firmware fix for RAID controller
 - RAID controller / disk drive incompatibilities
 - RAID controller firmware problems
 - I/O errors on disk array in JBOD mode
 - Dropped connections (or none at all) via BMC SOL
 - Loss of network connection to BMC

- 10 GigE to the server
 - Getting serious: 48 of 138 disk servers to be delivered this week will be connected via 10 GigE
- R&D
 - iSCSI
 - Lustre pilot project / evaluation
- Console service
 - After years of successful running, decided to replace SLAC console software by conserver

- Identity management
 - Using Microsoft ILM2 to replace CRA
 - Better support for SSO, hence better security
 - Issue with AFS groups, difficult to match to e-groups
 - Workaround implemented for LDAP access from SL5 machines
- Windows clients
 - Vista default for sufficiently powerful machines (e.g. 2 GB or more), but XP remains a choice
 - Windows 7: See Juraj Sucik's talk
 - Migration off Symantec anti-virus to Microsoft Forefront completed
 - IE 8 deployment started

- Mail services
 - Migration of mailing lists and archives to e-groups completed
 - Mailboxes migration to Exchange 2007 stalled
 - Found that pine MUA does not work correctly
 - Will migrate to Exchange 2010 directly
 - Phasing out support for pine etc. considered for end 2010
- Web services: Search
 - Slow progress
 - Carrying on with FAST (purchased by Microsoft)
 - Indico and Twiki being indexed

- Another readiness challenge in May/June: STEP (Scale Testing for Experiment Programme)
 - Data rates well above design values

- Preparing for new desktop / helpdesk contract
 - New contract as of 01-Jan-2010
 - Helpdesk now called service desk (ITIL)
 - Current helpdesk workflow in Remedy re-implemented using BMC components

- Dedicated presentation by Sebastian Lopienski