

Tier-0 Update

Helge Meinhard, CERN-IT
Grid Deployment Board
10-Jun-2015



Outline

- Cloud
- Databases
- Data and storage
- Network
- Platform services
- Infrastructure

Cloud

- OpenStack upgraded to Juno
- Hardware flow
 - Some retirements (first-time from OpenStack)
 - Many new servers – total now exceeding 4'700
- CentOS CERN 7 advantageous for hypervisors – migration campaign ongoing
- Hypervisor tuning for compute use case under way

Databases: Oracle

- Piquet support for experiment database services re-started on 18-May-2015
 - Similar to Run 1
 - Tracking of incidents and interventions in place
 - First two weeks: no critical service incidents
- Replication: Migration to GoldenGate completed
 - Thanks to Tier1s for the good collaboration

DataBase on Demand (MySQL, PostgreSQL, Oracle)

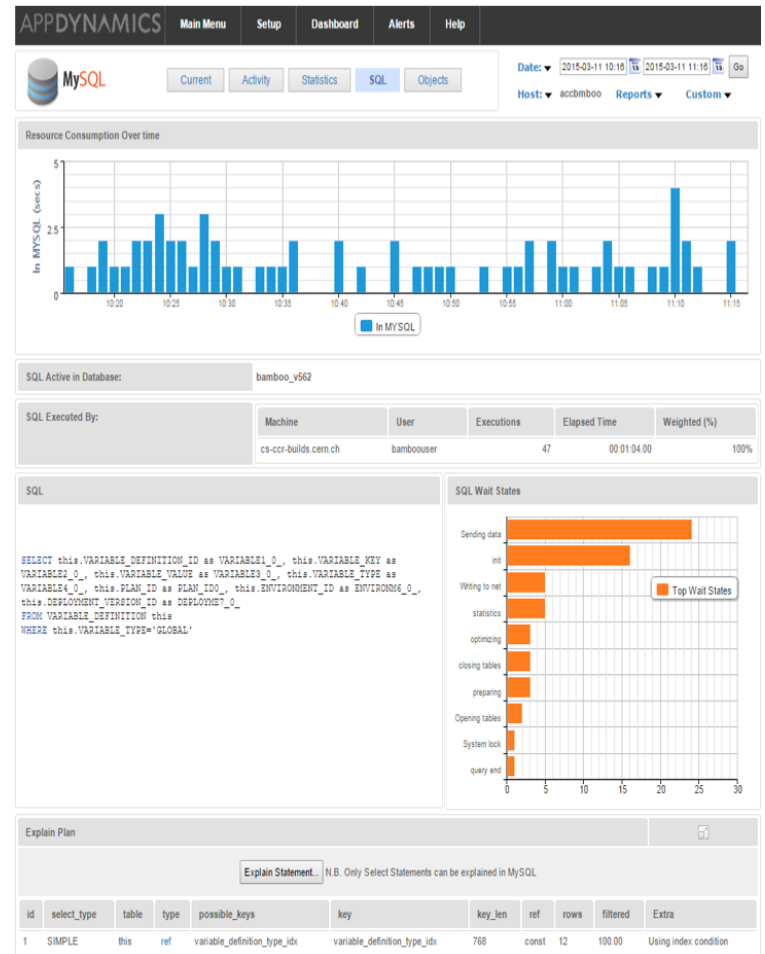
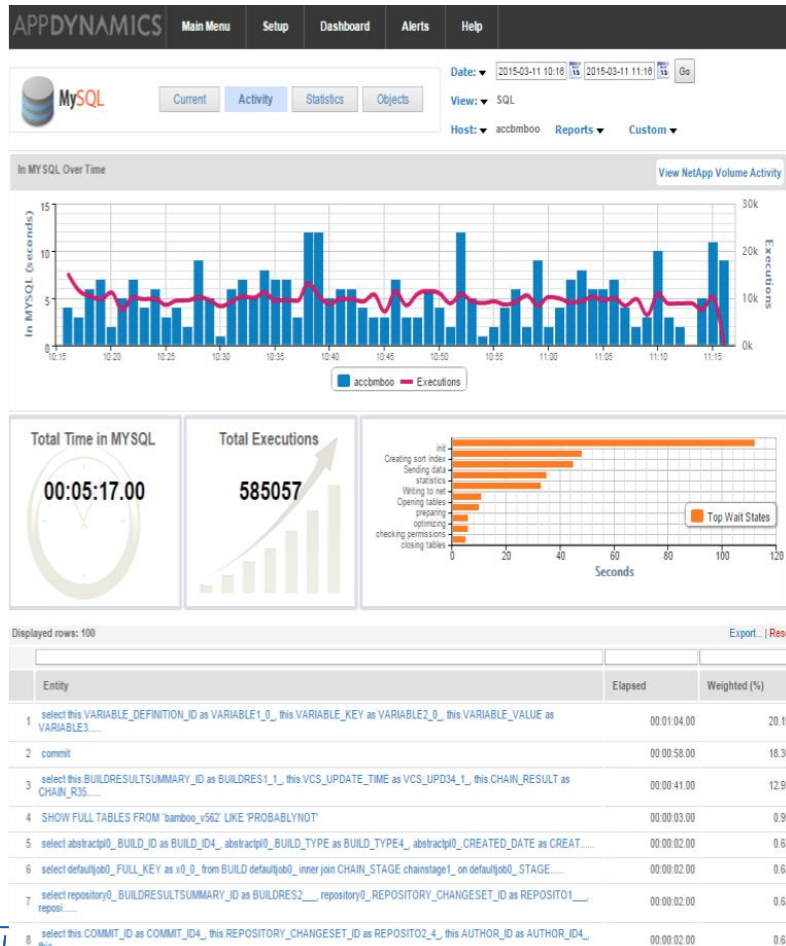
- Offloading many areas: IT, LHC, experiments
 - Used by 38 groups in 10 departments
 - Example: LHCb Dirac, Indico, OpenStack, Drupal, FTS3, ...
- Increasing demand, new requirements from CERN community: HA, replication, data protection
- New monitoring: Appdynamics

(<http://www.appdynamics.com/solutions/database-monitoring/>)

- Allow DBoD users to debug difficult performance problems in a very intuitive way
 - Access requires authentication
 - Working on integration with DBoD web site

DataBase on Demand (DBoD)

- Example AppDynamics display



Data and Storage (1)

- Preparation of data services for Run 2 finished
 - CASTOR v2.1.15 and SRM-2.11-2 for all LHC services
 - Only foreseen change: upgrade of SRM to consolidate the infrastructure (on top of the continuous HW renewal of obsolete machines)
 - Disk configurations drastically simplified (few disk pools, ideally only one per LHC experiment with a typical capacity of 2...3 PB to be used as tape front-end)
 - EOS: "Aquamarine" version (0.3) and BeStMan2-2.3.0-17.osg32 (all using xrootd 3.3.6-4.CERN)
 - Consolidation of headnodes (obsolete hardware) during summer
 - Capacity: about 140 PB (Meyrin:Wigner 50:50) online

Data and Storage (2)

- LHC Run2 data taking
 - ATLAS and CMS workflow now EOS-centric (Exp → CC; CC → T1s; CC → reconstruction)
 - CASTOR will archive data received from EOS (rather than directly from experiment)
 - ALICE and LHCb stay on the previous scheme: CASTOR gets data from experiment and acts as distribution point for further activities
- CERNBOX: Server on version 7, clients upgrading to 1.7.2
 - Service open to all CERN users (default quota: 1TB / user; 1M files / user)
 - Alpha users being enabled to try direct file access (EOSUSER) via dedicated lxplus node
 - Client (and server) use checksums (1.7.2 needed)
 - 2H2015: client going to 1.8.2 (minor fixes, better QA); new applications (notably ROOT.js with PH/SFT to have direct browsing of ROOT files)

Network

- WAN: see Edoardo Martelli's report earlier
- Change of mobile phone operator
 - CERN mobile numbers change from +41 76 487 xxxx to +41 75 411 xxxx on 24 June 2015 at 09:00 h
 - From CERN, calling 16 xxxx will continue to work

Platform and Grid Services (1)

- Significant capacity increases for Run 2: CPU a little rushed before Run 2 collisions (disk, tape in very good time)
- Availability issues showing CERN as low as 90% for ATLAS and CMS understood and contained
- Efficiency studies: working group, analytics project – will be re-enforced by dedicated person-power

Platform and Grid Services (2)

- Service portfolio adapted to needs
 - FTS2, WMS, AFS UI closed
 - LFC to be closed later this month
 - VOMS (SHA-1 only) and VOMRS replaced by SHA-2-aware versions and voms-admin
 - SLC5-based services (plus and batch) stopped, pilot for CC7-based plus established (CC7 supported since quite some time)
- Batch: Condor migration started
 - Small prototype for Grid submission, running jobs successfully (ATLAS, CMS, LHCb)
 - BDII integration done, accounting almost ready to go
 - Capacity increase later this month
 - Ramp-up to full production quality and scale for Grid submission over 2015
 - Second phase: local job submission, requires Kerberos ticket handling
 - Aim: stop LSF services by end of Run 2
- Plus: Performance improvements (hypervisors with SSD cache)

Infrastructure

- Version control: GitLab established
 - Git users encouraged to move; long-term: SVN to be considered for stoppage
- Issue tracking: Savannah → Jira migration completed
- Volunteer computing
 - Very significant uptake by ATLAS, major source of detector simulation
 - Lots of activities in CMS and LHCb, contacts starting with ALICE
 - Data bridge concept addressing issue of volunteer PCs getting access to data without sending credentials to the former
 - Potential BOINC funding issue
- Twiki
 - Running clustered now
 - Considering FOSwiki
- Lots of progress on AI – Quattor closed in time (end 2014)
 - Beware of now unmanaged hosts, in particular in context of service consolidation (CVI)

Acknowledgements

- Tim Bell, Maite Barroso Lopez, Tony Cass, Eric Grancher, Massimo Lamanna, Alberto Pace, Jan van Eldik

Questions?

