




www.cern.ch



Status of WLCG Tier-0

Maite Barroso, CERN-IT
With input from T0 service managers
Grid Deployment Board
9 April 2014



Outline

- Agile Infrastructure (AI)
- Facilities
- Services

Agile Infrastructure

- In production since July 2013 in Meyrin and Wigner
- Weekly meetings to provide a discussion forum for providers and users of the AI-related services
- VM provisioning (Openstack)
 - Offer a block volume service, with a CEPH backend
 - Growing at a high rate from the currently 60,000 cores
 - Service is used by 600+ people, 100+ projects
 - 2600+ hypervisors, 5500+ VMs today.
- Configuration management (Puppet etc.)
 - Ongoing migration from Quattor to Puppet, deadline this autumn (31 October 2014)
 - 8216 host in Puppet, 80-150 commits/week
- Monitoring infrastructure
 - LAS alarms gradually moving to GNI

Facilities (1)

- Wigner (Budapest)
 - 720 servers: 560 for openstack, 160 for EOS (as front-end servers)
 - about 5% of the EOS disk capacity deployed at Wigner
 - 160 disk arrays
 - Plus additional servers for Windows and AFS

Facilities (2)

In June we will install the following in addition:

- 80 CPU servers
- 376 front-end servers
- 376 disk arrays
- 36 DB servers

Generally good experience with operations to-date at Wigner and also with the reactivity of the Wigner team

Services (1)

- All grid services at WLCG baseline
- WMS decommission: draining of experiment instances ongoing
 - Expect SAM instances to be shut down by the end of June
- FTS: agreement to shut down FTS2 in summer, 1st of August
- VOMS:
 - Ongoing VOMS-admin to VOMRS transition, active testing with VO admins
 - Discrepancies in mail address field btw VOMRS-HR DB understood, fix/workaround being discussed
 - New version of VOMRS being prepared
 - VOMRS service can deal with SHA-2 certificates

Services (2)

- Batch services
 - Looking at replacing LSF over time
 - HTCondor looks promising
 - Large scale testing ongoing
 - Resources made available to ATLAS to run dedicated tests in view of their requirements for Tier-0 processing during Run 2.
 - Checking CPU accounting
- Job efficiency/throughput:
 - Meetings with IT experts and the 4 experiments to measure job efficiency
 - Conclusions will be presented next week:
 - No indication of a genuine Wigner problem
 - Multidimensional issue: CPU architecture, data access, VM/physical hw, SLC5/SLC6
 - Some derived actions already in place: monitor systematically job efficiencies on lxbatch

Services (3)

- SLC6 migration
 - Capacity at 65%, 5% more expected this week
 - Combined with migration to AI
 - From physical HW to VMs
 - Go through all resource provisioning chain: re-import the machines into various databases, configure them as hypervisors
 - Check point once we reach 80% to decide to completely shut down SLC5 capacity

Services: Network

- Experiments DAQ connections to the computer centre ready for Run2:
 - Alice: 40G
 - Atlas: 80G
 - CMS: 60G
 - LHCb: 20G
- Network to Wigner
 - Long standing issue on one Geneva-Budapest 100G link identified and fixed in the ISP network.
- IPv6 deployment
 - DHCPv6 roll out: leases given by default to all devices in Geneva and Wigner Datacentres

Services: Databases

- **Databases: Oracle version migration**
 - Database upgrades to 11.2.0.4 for Tier 1s are in progress. TRIUMF has already upgraded. IN2P3 is expected to upgrade in April. Other sites are expected to migrate by the end of summer.
 - Upgrades to 11.2.0.4 (majority) and 12c (LHCb offline, ATLARC) together with hardware upgrade well advanced at CERN.
 - End of 11.2 support in January 2018; looking at moving to 12c at one of the "end of year closures" (to be further discussed).
- **Databases: replication (Tier1 sites: BNL, IN2P3, TRIUMF and RAL)**
 - The workshop together with Tier1 DBAs in preparation for the migration from Streams to Golden Gate will take place the first week on June (during the ATLAS software week).
 - Migration scheduled to happen in June-July for online-to-offline replication. Migration is expected to be transparent to the users and will affect ATLAS and LHCb replication.
 - Migration CERN->Tier 1s is scheduled after that (September-October).
 - The details of the scheduling will be further discussed during the next ATLAS software week and Golden Gate workshop with Tier1s.

Questions?



