



CERN Site Report

Arne Wiebalck

HEPiX Spring Meeting
DESY Zeuthen, DE
Apr 18, 2016



CERN and CERN IT

- European Organization for Nuclear Research (Conseil Européen pour la Recherche Nucléaire)
 - Founded in 1954, today 21 member states
 - World's largest particle physics laboratory
 - Located at Franco-Swiss border near Geneva
 - ~2'300 staff members, >12'000 users

- CERN IT to enable the laboratory to fulfill its mission
 - Main data center on Meyrin site
 - Wigner data center, Budapest, 23ms distance
 - Connected via two dedicated 100Gbs links
 - Where possible, services have resources at both sites (plus disaster recovery)



MEYRIN DATA CENTRE

[last_value](#)

● Number of Cores in Meyrin	128,291
● Number of Drives in Meyrin	73,362
● Number of 10G NIC in Meyrin	6,456
● Number of 1G NIC in Meyrin	22,105
● Number of Processors in Meyrin	22,355
● Number of Servers in Meyrin	11,947
● Total Disk Space in Meyrin (TB)	133,730
● Total Memory Capacity in Meyrin (TB)	523

WIGNER DATA CENTRE

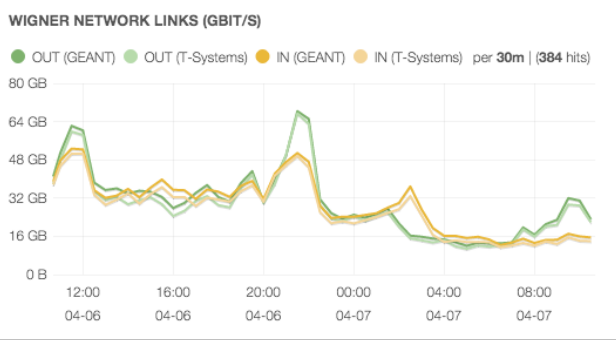
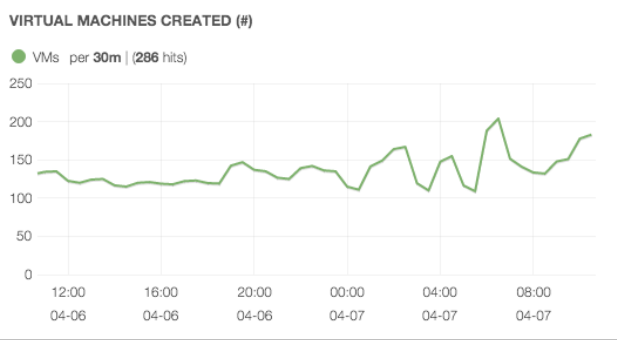
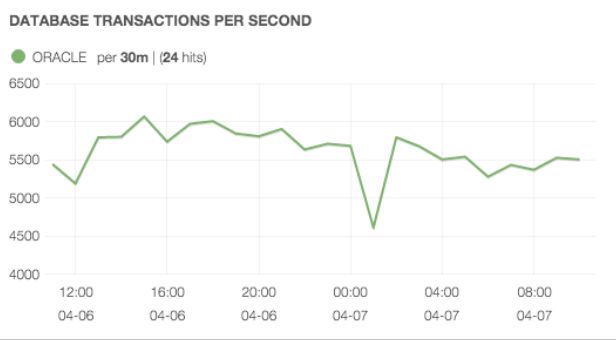
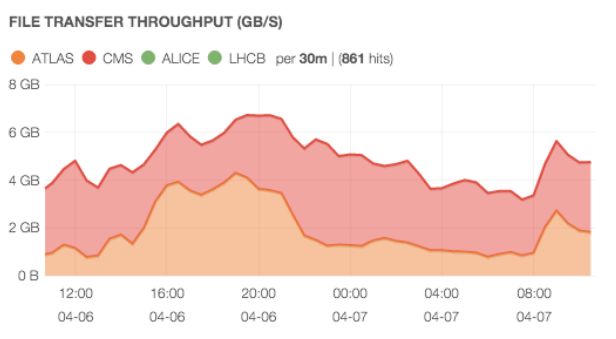
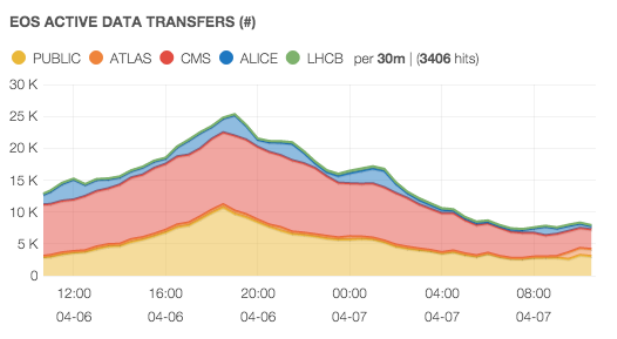
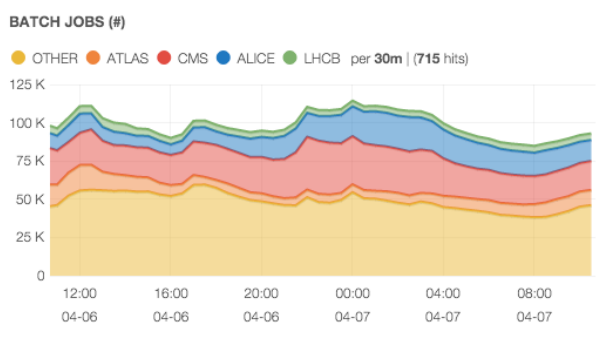
[last_value](#)

● Number of Cores in Wigner	43,328
● Number of Drives in Wigner	23,180
● Number of 10G NIC in Wigner	1,399
● Number of 1G NIC in Wigner	5,067
● Number of Processors in Wigner	5,418
● Number of Servers in Wigner	2,712
● Total Disk Space in Wigner (TB)	71,738
● Total Memory Capacity in Wigner (TB)	172

NETWORK AND STORAGE

[last_value](#)

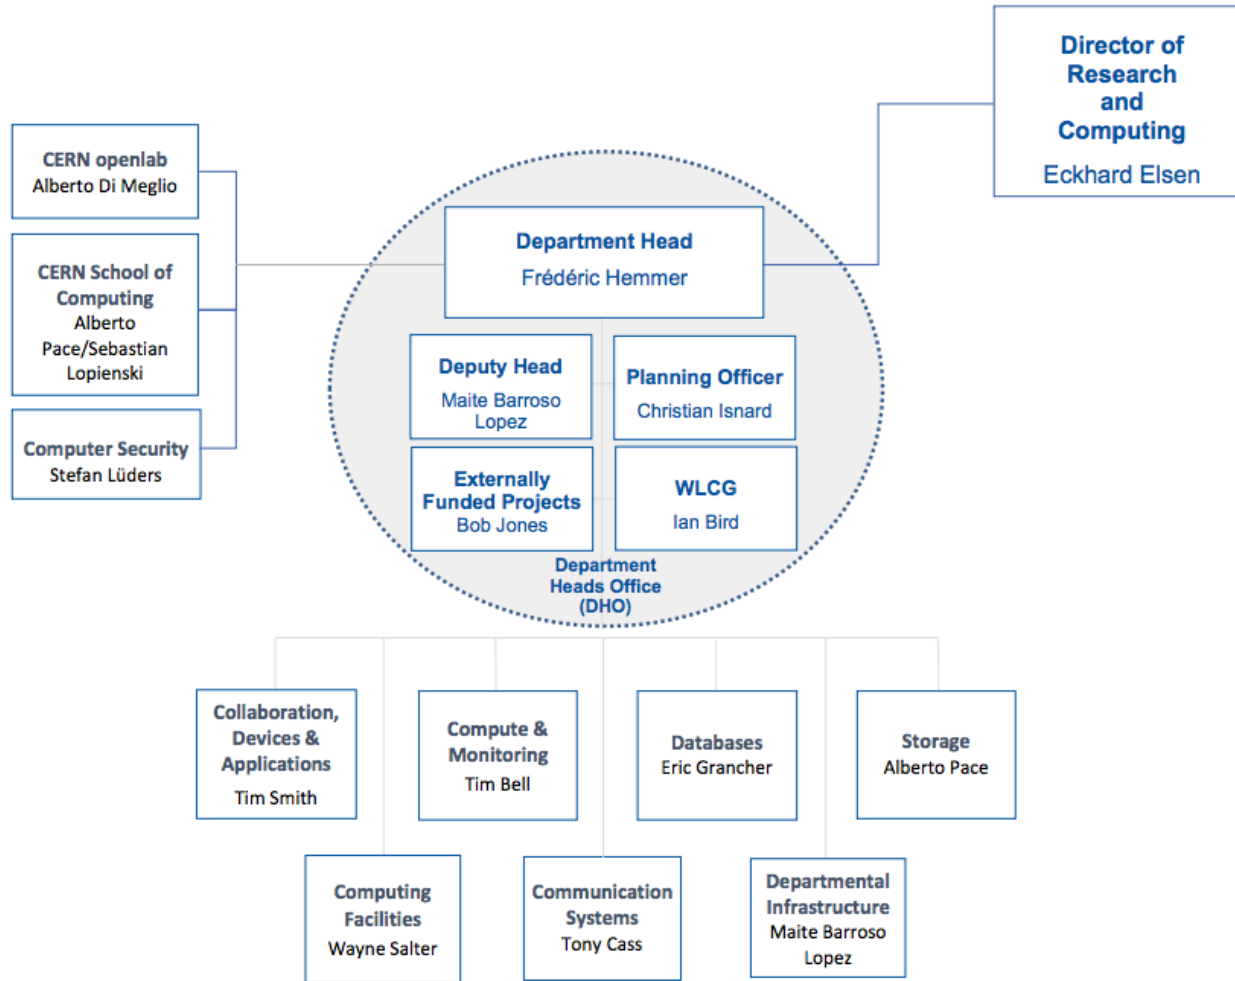
● Tape Drives	104
● Tape Cartridges	18,520
● Data Volume on Tape (TB)	138,415
● Free Space on Tape (TB)	27,773
● Routers (GPN)	138
● Routers (TN)	30
● Routers (Others)	107
● Switches	3,698



meter.cern.ch



2016 CERN IT (Re-)Organization



Storage: AFS Phase-out



- Motivation unchanged, see Alberto Pace's talk at BNL meeting
 - Less of a technical, more a community/support issue
- **AFS phase-out has been decided**
 - No hard deadline yet, expect final clean-up during LS2 (2019)
- No single replacement product, gathering input
 - CVMFS took over software distribution use case
 - EOS(+Fuse) plus CERNBox will cover many use cases
 - Discussions about \$HOME ongoing
- Heads up!
 - Current Windows clients (1.7.31) will break with SHA-1 deprecation end of 2016
 - Reports about potential data corruption in 4.4 kernels (SLC6, CC7 are OK)





Telephony: PABX Replacement

- Private Automatic Branch Exchange
 - Connects cables for office/emergency/lift phones
 - Connections to external providers
 - Network connections for Lync and IP phones
 - Routing, call handling, billing, ...
- Currently: Alcatel OmniPCX Enterprise
 - Four logical nodes in 17 physical locations
 - Installed in 1994
- Move to a software solution (Asterisk)
 - Move to IP or soft phones (Lync)
 - 1H2016: remove unused phones
 - 2H2016: move to s/w PABX
 - 2017: remove active phones
 - Q42018: remove PABX
- Details were given in Francisco Vinagrero's presentation in Oxford



Operating Systems



- Windows 7 remains the mainstream version 
 - MS will provide security updates until Jan 2020
 - Migration out of Win7 will be needed by the end of LS2
- Windows 10 now available as a pilot
 - MS Edge as new default browser
 - Adobe Flash disabled
 - Default PDF app is Edge
- CC7.2 current prod release
 - Since January 11, 2016
 - [CC7.2 details](#)
- Puppet on the desktop
 - Replace Quattor-based setup tools
 - master less configuration
- OpenAFS client issues on OS X, v. 10.11 
 - Now recommending third party packages (provided by SNA)

Hadoop Service



- Technology
 - 3 production + 1 QA cluster: ~100 nodes
 - Current tech stack: HBase, Map/Reduce, Spark, Impala + Parquet, Avro
- Community
 - Notable users: ATLAS Event Index + DDM, CMS WM Archive, WLCG + IT monitoring, Accelerator logs and controls project, CASTOR dashboards, ...
- Hadoop users forum and Analytics Working Group
 - <https://indico.cern.ch/category/5894/>

Version Control Services

- Move from GitoLite to GitLab

- GitLab in production since more than one year
- Hosts ~3'300 repositories (> git.cern.ch + svn)
- Very positive service experience
- git.cern.ch will be shut down by the end of 3Q2016



- GitLab-CI as a useful add-on

- Simpler than Jenkins, fully integrated with GitLab UI and MRs
- Docker infrastructure provided for job execution
 - User provided Docker images supported as well
 - User provided Docker infrastructure supported as well (required for heavy lifting)



- SVN service will be kept alive until LS2

- Feature-frozen



Batch Services

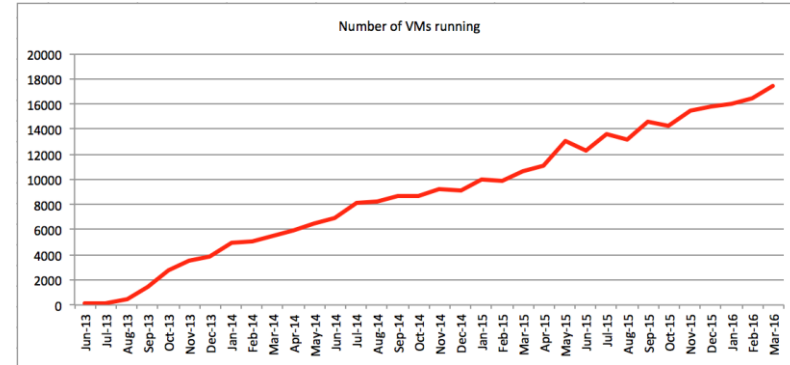


- LSF update to version 9
 - Only minor issues
- HTCondor pool expanded to 15k cores
 - Next increase dependent on migrating local submission (AFS/Kerberos)
 - AFS/Kerberos integration is completed (HTCondor 8.5.4)
- Decommissioned ARC CEs, now 100% HTCondor CEs
- Starting to dedicate multicore resource in HTCondor (initially for CMS)
- Began work adding external cloud resources to HTCondor pool
 - All details in Jérôme's talk on Wednesday

AI: Cloud Service (IaaS)



- CERN private cloud still growing
 - ~5'800 hypervisors, ~36% at Wigner
 - Hosting ~18'000 VMs
 - ~1'800 to deploy in 2Q16, another 400kHS06 4Q16
 - No more spinning rust!
 - Retirement of ~2'000 compute nodes by 3Q16
 - Tricky bit: 800 hosts running services



No of VMs running since July 2013

- Rolling upgrade to Liberty ongoing
 - Glance, Keystone, Ceilometer, Heat, Neutron done
 - Horizon, Cinder and Nova left to do

- NUMA/THP configuration rolled out

- >2'000 hypervisors reconfigured
- ~6'000 batch VMs drained/deleted/recreated!

- Integration of containers moving to production

- Magnum: COE as first class OpenStack resources
- See Bertrand's talk on Thursday

VM sizes (cores)	Before	After
4x 8	7.8%	3.3% (batch WN)
2x 16	16%	4.6% (batch WN)
1x 24	20%	5.0% (batch WN)
1x 32	20.4%	3-6% (bare SLC6 ... batch WN)



AI: Config' Management



- Currently: puppet-server-1.1.3, puppet-3.8.4, foreman-1.6.1, puppetdb-2.3.8
- Catalog compilations dropped from 73 to 25 seconds
 - Ruby2, puppetserver instead of passenger, removal of file stats
 - All changes merged with upstream
- Puppet4 / Foreman 1.11 planned for this year
 - “Future parser” already available within CERN’s infrastructure
- Batch on external clouds (xBatch)
 - Separate Puppet masters and Puppet DB

Miscellaneous Recent Activities

- Set up a centralized Elastic Search service: ESaaS

- There are currently ~20 clusters managed individually
- Conflicting user requirements, Resource planning, Limit user cross-talk, Accounting, ...



- Investigating time series databases

- Several requests (from experiments and IT)
- Proposed implementation: leverage DB on demand infrastructure
- Being prototyped, still collecting requirements from interested users



- Looking at potential cost savings by restructuring the backup service

- ORACLE and DFS main data producer for backup in TSM
- Various options being considered

More CERN talks this week

- Update from **Database Services** (Katarzyna Dzedziniewicz-Wojcik)
- Automating Operational Procedures with **Rundeck** (Daniel Fernandez)
- xBatch: Extending the CERN **Batch** Service into **Public Clouds** (Jérôme Belleman)
- First Experiences with **Container Orchestration** in the CERN Cloud (Bertrand Noel)
- **Storage** at CERN: Towards the Chameleon (Xavier Espinal Curull)
- Evolution of **Printing Services** (Natalie Kane)
- Running Virtualized **Hadoop**: Does it make sense? (Kacper Surdy)
- The **Helix Nebula** Science Cloud Project (Helge Meinhard)
- **Computer Security** Update (Liviu Valsan)
- CERN **Cloud** Update (Arne Wiebalck)

✔ Accelerator Database	✔ ACRON	✔ Administration Database	✔ AFS	✔ AFS Web Hosting	✔ Alerter	✔ Backup and Restore	✔ Batch	✔ BOINC
✔ Bull	Starting at the autumn meeting at LBNL				✔ CERN Search	✔ CERNOBox	✔ Certificate Authority	✔ CIX P
✔ Confe Roc	Jérôme Belleman				✔ Databases Applications	✔ Datacenter Network	✔ DFS	✔ Drupal
✔ E-M	will take over as the CERN representative for HEPiX				✔ E-Fax	✔ File Transfer	✔ FILER	✔ Fixed Line Phone
✔ General Purpose Database	✔ Git	✔ GRID Compute Element	✔ GRID Development	✔ GRID Information	✔ GRID Infrastructure Monitoring	✔ HADOOP	✔ IIS Web Hosting	✔ Indico Event Application Support
✔ Java Web Hosting	✔ JIRA	✔ Linux Operating System	✔ Load Balancing	✔ LXPLUS	✔	✔	✔	✔
✔ Monitoring	✔ MultiMedia	✔ MyProxy	✔ Network Database and Registration	✔ Network for Projects and Experiments	✔	✔	✔	✔
✔ Single Sign On and Account Management	✔ SVN	✔ Technical Network	✔ Twiki	✔ VOMS	✔ WIFI	✔ Windows Desktop	✔ Windows Terminal Servers	✔ WLCG Network

Thanks!

