

CERN Site Report

Andrei Dumitru

CERN IT Department

8th of October 2018

CERN

European Organisation for Nuclear Research
(Conseil Européen pour la Recherche Nucléaire)

Founded in 1954, 22 member states today

World's largest particle physics laboratory

Located at Franco-Swiss border near Geneva

≈ 2500 staff members, ≈ 17000 users

SUISSE
FRANCE

LHCb

ATLAS

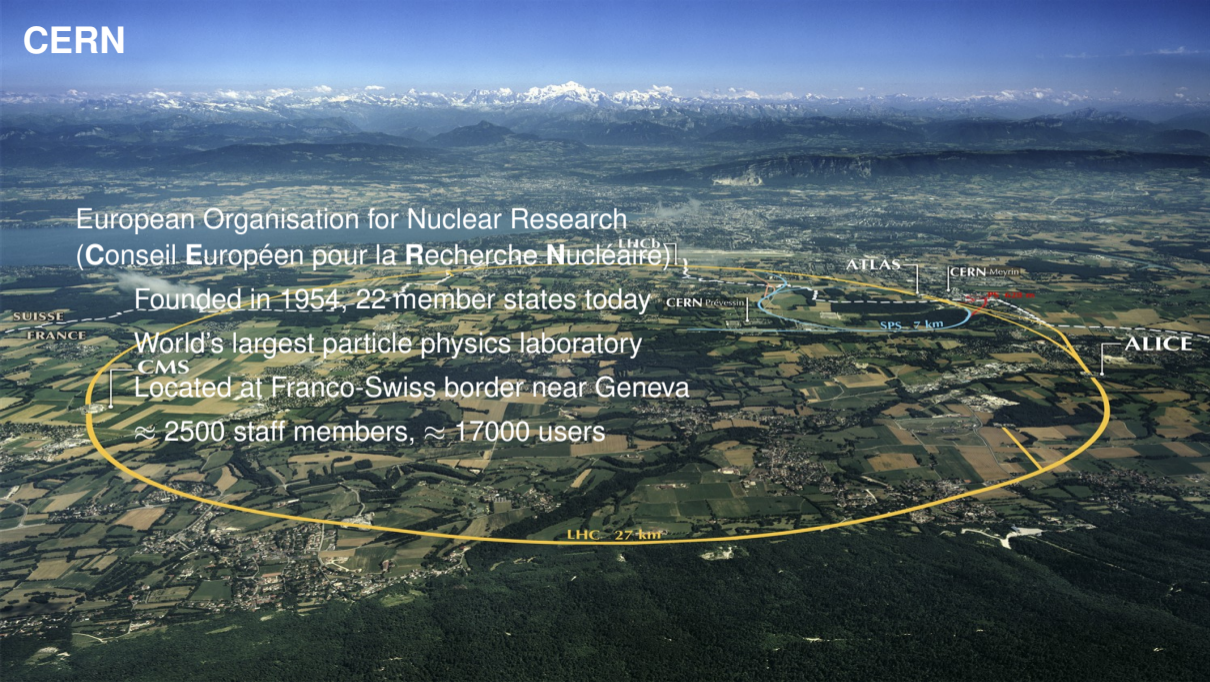
CERN-Meyrin

CERN Prévessin

SPS 7 km

ALICE

LHC 27 km



CERN IT Department



Enabling the laboratory to fulfill its mission

Main data centre in Meyrin site

Wigner data centre in Budapest, 23 ms away

Connected via 3 dedicated 100-Gb/s links

Service resources in both sites where possible
(+ disaster recovery)

COMPUTING

Servers (Meyrin)

11.5 K

Cores (Meyrin)

174.3 K

Servers (Wigner)

3.5 K

Cores (Wigner)

56.0 K

STORAGE

Disks (Meyrin)

61.9 K

Tape Drives

104

Disks (Wigner)

29.7 K

Tape Cartridges

32.2 K

NETWORK

Routers

250

Star Points

692


Switches

4.1 K


Wifi Points

1.2 K


Batch Jobs




EOS Active Data Transfers




File Transfer Throughput




Cloud Virtual Machines



Databases Transactions



Wigner Network Links



Cloud Service Update

Expiration of personal VMs

- ▶ Lease of 180 days (renewable)
- ▶ Expiration warnings via mail
- ▶ Upon expiration: stop & locked, then deleted
- ▶ Workflow based on OpenStack Mistral

File Share service in production

- ▶ CephFS as the backend
- ▶ Complements block storage (volumes) offering
- ▶ Management done via OpenStack Manila



MISTRAL
an OpenStack Community Project



MANILA
an OpenStack Community Project

Cloud Service Update

Prototype: Pre-emptible VMs

- ▶ Increase resource utilisation with opportunistic instances
- ▶ Integration with OpenStack Nova ongoing



Bare Metal Provisioning

- ▶ In production, preparing enrollment of 10k servers
- ▶ Based on OpenStack Ironic



L1TF: Cloud reboot campaign

- ▶ Impacted: 1100 hosts, 11.5k VMs
- ▶ Unconditional flushing + SMT off
- ▶ Performance impact on high overcommit hosts under investigation



Centralised Elasticsearch Service

Smooth production running since last HEPiX
Completed migration to SSD based hypervisors
Current ES versions

- ▶ 5.5.2, upgrading to 5.6.9
- ▶ Added support for 6.2.4 and started migration
- ▶ Added additional clusters for 6.x to ease up the migration path

Reduced intervention time by basing ES replica placing policy on tenants rather than hypervisors

New Kibana plugins for 6.x:

- ▶ Enhanced table
- ▶ SSO logout

More information: <https://cern.ch/esdocs>



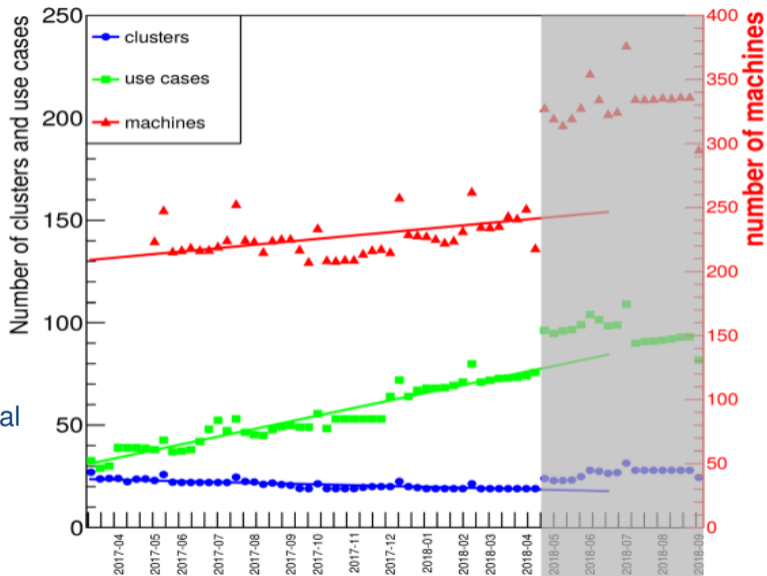
Centralised Elasticsearch Service

Blue: decreasing number of clusters

Red: Slow increase of machines

Green: 3 new use cases / month

Grey: Deployment of additional resources and new test clusters for ES6



Compute and Monitoring

This Tuesday

CERN Linux services status update

by Thomas Oulevey

Workflows automation @CERN IT with CERNMegabus

by Zhechka Toteva

Compute and Monitoring

This Wednesday

Too much BEER and the Kubernetes dance
by Ben Jones

Data integration / reporting service

Challenging count and complexity of data sources

Call for centrally hosted data-integration solution

- ▶ Optimizing organization resources
- ▶ Data extraction, transformation, loading
- ▶ Create, store, share and publish reports
- ▶ One solution for many use-cases
- ▶ Self service design

Scale-up and migration of existing non-IT deployment



Oracle WebLogic on Kubernetes

Improve the provisioning making it:

- ▶ Immutable
- ▶ Repeatable

Demonstrate that almost all CERN Java critical applications are working on WebLogic running on Kubernetes clusters

- ▶ OpenStack Magnum to provision Kubernetes Cluster

What's next : On Premise & Public Cloud

- ▶ Minimise effects of disruptive events
- ▶ Taking advantage of Oracle Cloud Credits

Streaming and Analytics Services

Apache Spark on Kubernetes with physics data test (scaling 60 to 180 VMs)

- ▶ Investigations with RDataFrame for tight integration with ROOT framework
- ▶ Improvements on xrootd connector/spark-root (allows to read EOS directly from Spark and integrate with other "Big Data" tools) - alternative to HDFS
- ▶ Load test 500 VMs with over-tuned read-ahead parameter. Network saturated!

Hadoop

- ▶ 74 physical machines, planning to add 40 additional
- ▶ Moving out from Cloudera distribution to Apache

SWAN Integration (Service for Web based ANalysis)

- ▶ Automatic generation of hadoop credentials (delegated tokens support in hadoop).

Streaming service (Kafka) multitenancy

- ▶ Pilot self-service portal for topic management and ACLs with e-groups support

Machine Learning: planned testing on Oracle cloud and GPUs

This Thursday

*LHC Long Shutdown 2 and database changes
by Ludovico Caldara*

Storage

2018 data taking so far: +66PB to tape (295PB total), record peak (August): 13.8PB

- ▶ Ready for Heavy Ions (successful Data Challenge)

Repacked 100PB in ~6 months

- ▶ Rewriting IBM enterprise from 10TB to 15TB; 2 lost files (out of 135M)
- ▶ Using a small, local SSD disk cache layer (80x3.5TB SSD)
- ▶ More repacking (80PB out of T10K media) during LS2

New LTO library with 10K slots and 20 LTO-8 drives in production

- ▶ Preparing tender for another library ~2019
- ▶ Sparse read performance

CTA (CERN Tape Archive) field test deployments

Storage

EOS service

- ▶ Instabilities observed in Spring under control
 - ▶ New catalogue needed on the long run (in prod after Heavy-Ion run)
- ▶ BEER in prod for ATLAS and LHCb (tested on CMS)
 - ▶ "containerised" Condor jobs running on EOS (eligible) disk serves

Progress with the new FUSE and new EOS catalogues

- ▶ New deployment model (scale-out instances)
- ▶ EOSBACKUP close to 1B entries; reboot time from $O(1000)$ s to less than 10 s

EOSUSER \implies EOSHOME migration in full swing (CERNBox backend)

Preparation of Heavy-Ion migration

- ▶ Demonstrated at scale (ALICE 8GB/s + CMS 6GB on top of "regular" activities)

Update of the FILER migration

- ▶ Move from the VM+CephRADOS Block Device \longrightarrow Manila shares (CephFS)
- ▶ Remove single-point of failure. Almost completed

Storage

This Thursday

Storage at CERN

by Cristian Contescu

LTO experiences at CERN

by German Cancio Melia

Latest developments of the CERN Data Management Tools

by Edward Karavakis

Storage

This Friday

Backup Infrastructure at CERN

by Aurelien Gounon

CERN Tape Archive initial deployments

by Julien Leduc

IT Facilities

Summary of 2018 new hardware

- ▶ 377 servers
- ▶ 167 JBODs (each with 24 disk drives, disk capacity 6TB and 12TB (majority))
- ▶ 1 server with nVidia V100 GPU for Techlab
- ▶ 4 servers, each with 4x nvidia V100 GPUs for Openstack (order in preparation)
- ▶ 85 "intelligent" PDUs
- ▶ 96 servers and 17 JBODs for non-IT (experiments and other departments).

The System Administration service will be stopped by the end of 2018.

IT Facilities

This Tuesday

Challenges, solutions and lessons learnt in 7 years of Service Management at CERN


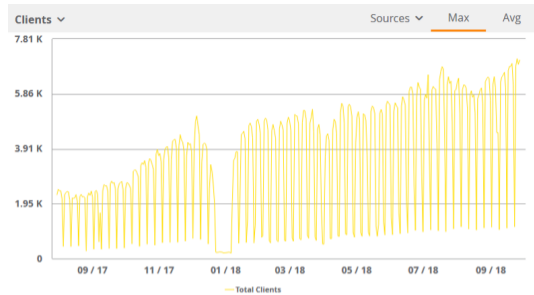
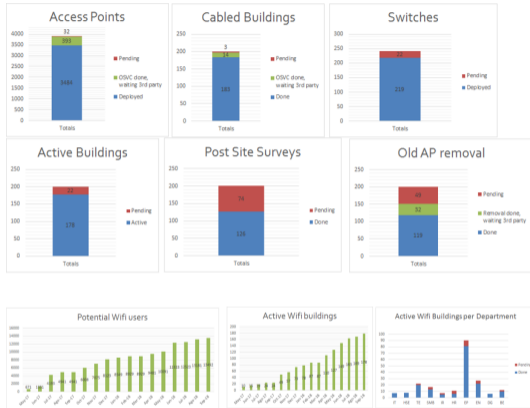
by David Martin Clavo

This Thursday

CERN procurement update

by Olof Barring

Wi-Fi progress



Outdoor coverage now included in project:

- ▶ most pedestrian routes on the Meyrin site
- ▶ CCC & Restaurant 3 areas on the Preveessin site

Networking

Second network hub now provides full redundancy for internal operations

- ▶ Many external fibre connections to the second hub are still not redundant
- ▶ Will be addressed gradually over the next year

Successfully demonstrated for LHCb operation of a remote HLT subfarm in the Computer Center

- ▶ Operationally and performance-wise: same as running a subfarm at LHCb point 8

Tested a low cost DWDM solution from SmartOptics which can multiplex up to 40 100Gbps channels into a single fibre pair

- ▶ 200Gbps and 400Gbps options are expected by 2021

Networking

FTTO technology could be an interesting alternative to structured cabling in experimental areas

- ▶ Will purchase equipment for a test in 2019
- ▶ Not seen as an alternative to the existing structured cabling in office buildings

Faxes no longer part of the universal service obligation for telephony providers in Switzerland

- ▶ Correct transmission of faxes is no longer guaranteed

Networking

This Monday

Network configuration management at CERN: status and outlook
by Arkadiy Shevrikuko

Update about Wi-Fi service enhancement at CERN
by Adam Sosnowski

WLCG/OSG Networking Update
by Marian Babik

Networking

This Monday

Network Functions Virtualisation Working Group Update

by Marian Babik

CERN campus network upgrade

by Carles Kishimoto Bisbe

Challenges for connecting Internet of Thing devices at CERN

by Adam Sosnowski

This Tuesday

Network configuration management tools evolution
by Nikolaos Papoutsis

Digital Repositories

CDS

- ▶ Development on-going for a new version of the Open Source Library Loan and Book Search software (targeting the release at the end of the year)

CERN Open Data

- ▶ First non-LHC experiment using CERN Open Data portal
- ▶ OPERA collaboration released tau neutrino and multiplicity studies electronic detector and emulsion data

REANA

- ▶ REANA v0.2.0 released
- ▶ Adds support for running Common Workflow Language (CWL) computational workflows on Kubernetes clouds
- ▶ Simplifies REANA command line client usage scenarios

Zenodo

- ▶ Usage statistics launched: find the number of views and downloads on record pages and sort search results by most viewed

Collaboration tools

Vidyo Service

- ▶ New VidyoConnect client (with whiteboard functionality) deployed in August

Indico

- ▶ Visitors CERN site access formalities can now be done through Indico (through a dedicated plugin)
- ▶ Release 2.1 includes importation of registrants and contributions, and creation of customized roles
- ▶ Next release (2.2) will be a redesign of the room booking features

Applications and Devices

Ongoing work on a unified home directory using CERNBox

- ▶ For Windows, macOS, Linux personal devices (using sync client)
- ▶ With access from iOS and Android (mobile apps)
- ▶ In addition: CERNBox web access and SMB/FUSE mounts to EOS

Started migration of user home directories from DFS to CERNBox

Mobile Application Management (MAM) solution for iOS, Android, macOS and Windows

- ▶ Investigating open source and commercial solutions available on the market
- ▶ PoC with MobileIron

Any experiences with MAM/MDM/EMM suites ? (please get in touch)

Collaboration, Devices and Applications

This Tuesday

Container orchestration as key to the evolution of CERN Web Services

by Alexandre Lossent

Indico - Event management in HEP

by Pedro Ferreira

MAlt Project

by Emmanuel Ormancey

The new CERN Authentication and Authorization

by Paolo Tedesco

This Wednesday

A Data Lake Prototype for HL-LHC

by Xavier Espinal

This Thursday

Cost and system performance modelling in WLCG and HSF: an update

by Andrea Sciabà

Computer Security

This Wednesday

Computer Security Update by Liviu Valsan

Data Privacy at CERN IT

This Wednesday

Data Protection at CERN IT by Andrei Dumitru

Talks from CERN this week

Talks from CERN this week

Network configuration management at CERN: status and outlook

- ▶ by Arkadiy Shevrikuko

Update about Wi-Fi service enhancement at CERN

- ▶ by Adam Sosnowski

WLCG/OSG Networking Update

- ▶ by Marian Babik

Network Functions Virtualisation Working Group Update

- ▶ by Marian Babik

CERN campus network upgrade

- ▶ by Carles Kishimoto Bisbe

Challenges for connecting Internet of Thing devices at CERN

- ▶ by Adam Sosnowski

Talks from CERN this week

Workflows automation @CERN IT with CERNMegabus

- ▶ by Zhechka Toteva

The new CERN Authentication and Authorization

- ▶ by Paolo Tedesco

Network configuration management tools evolution

- ▶ by Nikolaos Papoutsis

Challenges, solutions and lessons learnt in 7 years of Service Management at CERN

- ▶ by David Martin Clavo

CERN Linux services status update

- ▶ by Thomas Oulevey

Indico - Event management in HEP

- ▶ by Pedro Ferreira

Talks from CERN this week

MAlt Project

- ▶ by Emmanuel Ormancey

Container orchestration as key to the evolution of CERN Web Services

- ▶ by Alexandre Lossent

A Data Lake Prototype for HL-LHC

- ▶ by Xavier Espinal

Computer Security Update

- ▶ by Liviu Valsan

Data Protection at CERN IT

- ▶ by Andrei Dumitru

Technology watch working group

- ▶ by Helge Meinhard

Talks from CERN this week

Cost and system performance modelling in WLCG and HSF: an update

- ▶ by Andrea Sciabà

CERN procurement update

- ▶ by Olof Barring

Storage at CERN

- ▶ by Cristian Contescu

LHC Long Shutdown 2 and database changes

- ▶ by Ludovico Caldara

Latest developments of the CERN Data Management Tools

- ▶ by Edward Karavakis

Talks from CERN this week

LTO experiences at CERN

- ▶ by German Cancio Melia

Backup Infrastructure at CERN

- ▶ by Aurelien Gounon

CERN Tape Archive initial deployments

- ▶ by Julien Leduc

Too much BEER and the Kubernetes dance

- ▶ Ben Jones



home.cern