

CERN Site Report

Andrei Dumitru

CERN IT Department

14th of October 2019

CERN

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

Founded in 1954, 23 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

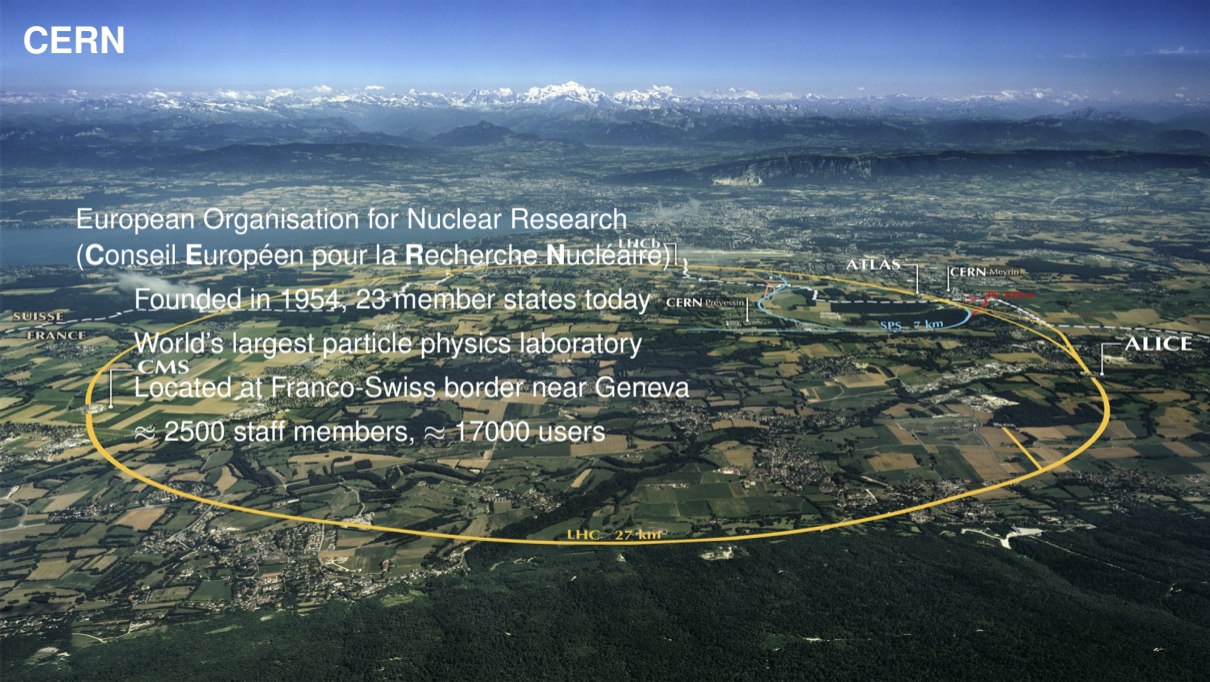
CERN-Meyrin

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)

10.6 K

Cores (Meyrin)

174.4 K

Servers (Wigner)

3.9 K

Cores (Wigner)

63.1 K

STORAGE

Disks (Meyrin)

60.5 K

Tape Drives

77

Disks (Wigner)

30.5 K

Tape Cartridges

29.5 K

NETWORK

Routers

288

Star Points

722


Switches

4.3 K


Wifi Points

4897


Batch Jobs




EOS Active Data Transfers



File Transfer Throughput



Cloud Virtual Machines



Wigner

Metric	Avg
Servers	4 K
Processors	7 K
Cores	63 K
Disks	30 K
Memory Modules	32 K
1GB NICs	7 K

Wigner Network Links



BATCH, HPC and CERNMegabus

Batch migration to CERN CentOS 7 (99%)

LSF retired now

- ▶ HTCondor in production for all batch

Backfill work ongoing to improve resource efficiency

Wigner batch resources

- ▶ will be drained by end of year

Majority of HPC work now migrated to Linux from Windows

- ▶ SLURM for MPI applications
- ▶ HTCondor for single-node, large memory applications

Looking at Batch on Kubernetes to improve resource utilisation and automation

CERNMegabus

- ▶ Upgraded CERNMegabus python libraries to stomppy-4.1.21.
- ▶ Redesigned the decision-taking process for the Power Cut Event, in order to adapt it to the different generations of UPS.

Monitoring

Lemon phase-out

- ▶ Migration to Collectd was completed. All Collectd data was integrated in the MONIT infrastructure.
- ▶ Lemon decommissioning by end of September.

Measuring IT Services

- ▶ A new MONIT workflow to gather service KPIs and SLIs was put in place.

Migration of ATLAS and CMS monitoring

- ▶ The migration of the old ATLAS Data Management, ATLAS Job Accounting, and CMS Job/Task Monitoring dashboards to the MONIT infrastructure was completed and the old infrastructure retired.

Messaging moving to open source

- ▶ Ongoing migration from Oracle Java + Red Hat A-MQ to OpenJDK + open source ActiveMQ
- ▶ expected to be completed by the end of the year

This Tuesday

*CERN Linux services status update
by Daniel Abad*

This Friday

*CERN Cloud Infrastructure update
by Daniel Abad*

Streaming and Analytics Services

Hadoop

- ▶ Current setup: 120 physical machines, 30TB RAM, 4k cores, 26PB of storage
- ▶ Evaluated Hadoop3 features, including Erasure Coding - [details here](#)
- ▶ Migrated all monitoring to central service based on Collectd
- ▶ Data migration of biggest Oracle database - accelerator logging, 1PB (full production next year)

SWAN Integration (Service for Web based ANalysis)

- ▶ Possibility to offload Spark computation to Kubernetes cluster
- ▶ On-going work to port SWAN service to run on Kubernetes.

Machine Learning

- ▶ Implemented a full data pipeline with distributed deep learning, for a physics use-case, based on Spark ([details here](#))

Streaming service (kafka) multitenancy

- ▶ self-service portal for topic management + ACLs with e-group support now in prod

Database Services

Oracle databases

- ▶ Upgrading databases from 11g/12c to 18c/19c

Weblogic

- ▶ Launched migration from VMs to Kubernetes
- ▶ Preparing for ver. 12.2 upgrade

Oracle Cloud evaluation

- ▶ Developed open days ticketing application entirely hosted in Oracle cloud
- ▶ Angular application deployed to Kubernetes
- ▶ Oracle database deployed as Autonomous Transaction Processing
- ▶ Generally good experience

Oracle signed agreement with Geant. Soon all Geant users will be able to get "fast track" connectivity to Oracle Cloud.

This Thursday

*CERN Database on Demand Update
by Ignacio Coterillo Coz*

Storage

Wigner decommissioning status

- ▶ 90PB of EOS capacity → transfer/replication of 68PB of data so far (12PB to go)
- ▶ entire Ceph capacity (2PB)

FTS (V3.9)

- ▶ Scheduler performance improvements: up to 10-15% reduced scheduling latency under load
- ▶ Several enhancements related to CTA integration

XrootD (V5.x)

- ▶ Significant performance improvements for multi-stream transfers
- ▶ Implementation of root/xroot protocol over TLS finalised

Disk Storage

EOS

- ▶ Deployed new namespace technology for scaling + new, robust FUSEX clients

AFS

- ▶ Less pressure from upstream to decommission the service, shrinking continues

CEPH

- ▶ Deployment of new Nautilus release ongoing, service optimizations for Block storage and CephFS

S3

- ▶ New service with object store backend

CERNBox

- ▶ Integration of several collaborative applications ("application hub")
- ▶ Service optimizations and new functionalities rolled-out

CVMFS

- ▶ Shortened publication time for LHC experiments with the new S3 backend

Tape

LTO in full deployment

- ▶ >55PB data on tape, ongoing library tender

Oracle (almost) decommissioned

Preparing ramp-up for Run-3

CTA validated, preparing switchover from CASTOR for LHC instances

This Thursday

*Current status of tape storage at CERN
by Dr Steven Murray*

*CERN Storage Evolution
by Giuseppe Lo Presti*

Computing Facilities

Procurement

- ▶ June adjudication, storage tender, delivery during fall: 192 JBODs each with 24 disk drives of 12TB providing a total of 55PB+
- ▶ June adjudication, compute tender, delivery during fall: 708 servers
- ▶ September adjudication, storage and compute: 6x blanket contracts valid for 18 months are being put in place

Repatriation of hardware from Wigner is progressing well

- ▶ 1772 CPU servers + 320 disk arrays and 160 front-end servers
- ▶ All CPU servers have been installed in one LHCb computer module, which is now full
- ▶ The final batch has been drained and is being prepared for repatriation in October
 - ▶ 636 CPU servers
 - ▶ 120 disk arrays and 60 front-end servers
- ▶ For the remaining capacity: the batch queues will be stopped on 1st of November

This Wednesday

*How to provide required resources for Run3 and Run4
by Wayne Salter*

*Open Compute Project 2019 Global Summit Report
by Luca Atzori*

Networking

Feasibility of expanding the LHCOne community

- ▶ CERN and Fermilab will establish a "DuneOne" network
- ▶ teams at CERN will study options to dynamically allocate shared resources to the appropriate VRF

Datacentre network

- ▶ updated, to provide more capacity, flexibility and redundancy
- ▶ extended, including a test WDM setup - to support the IT services in LHCb containers at point 8

Transfer broker developed (NOTED project)

- ▶ collect information about upcoming and ongoing data transfers
- ▶ used as basis for SDN optimization of network connections

Wi-Fi pilot

A "Wi-Fi mostly" pilot is underway for IT staff

- ▶ Mobile devices and NUCs are not allowed to connect to structured cabling
- ▶ Wi-Fi capable devices (including some Canon printers) moved as well
- ▶ More than 150 Ethernet outlets are no longer being used on a regular basis
- ▶ Issue identified so far: inability to install Linux systems over Wi-Fi

This Wednesday

*NOTED: Network-Optimized Transfer of Experimental Data
by Coralie Busse-Grawitz*

*CERN Computer Center Network evolution
by Vincent Ducret*

*LHCb containers - Network overview
by Daniele Pomponi*

Collaboration, Devices and Applications

Indico

- ▶ Indico 2.2 released (improved room booking module) in June
- ▶ Licence plates handling for visitors' cars available in Indico
- ▶ Coming next:
 - ▶ improved paper editing/review workflow
 - ▶ Doodle-like scheduling tool (Newdle for MALT)

Telephony: ongoing service development (MALT), starting pilot

Mail: preparing migration to Kopano (MALT), starting pilot

E-Fax

- ▶ Service shutdown: end Nov. 2019
- ▶ E-fax server is not compatible anymore with the TONE infrastructure

Collaboration, Devices and Applications

Identity and account management

- ▶ New SSO service based on Keycloak - first use cases being migrated
- ▶ New Authorization service - tests ongoing

Mattermost

- ▶ Integrated with CERN e-groups

Documentation available on

- ▶ Creating Markdown-based documentation web sites (details [here](#))
- ▶ Making short videos explaining your services (details [here](#))

This Monday

*CERN Fixed Telephony Service Development
by Thomas Baron*

*Unified Home folders migration or
chronicles of a complex data migration
by Vincent Nicolas Bippus*

This Monday

*Challenges and opportunities when migrating
CERN e-mail system to open source
by Thomas Baron*

*Self-service web hosting made easy
with containers and Kubernetes Operators
by Alexandre Lossent*

This Tuesday

*Zero-touch Windows 10 migration: dream or reality?
by Vincent Nicolas Bippus*

CERN Open Data, REANA and CDS

CERN Open Data

- ▶ CMS released, on 18 Jul 2019, open data for Machine Learning, details [here](#)
- ▶ total volume of CMS open data over 2PB now
- ▶ release includes datasets prepared specifically for use in Machine Learning or in Data Science

REANA - [v0.5.0 released](#)

- ▶ Allows to explore workflow results by running interactive Jupyter notebook sessions on the workspace files
- ▶ Supports now: HTCondor compute backend / Hybrid HTCondor-Kubernetes workflows / User Secrets such as Kerberos to access private data
- ▶ integrated with GitLab in order to run REANA Workflows from GitLab's CD/CI

CERN Document Server (CDS)

- ▶ Last phase of the on-going development work for a new version of the Open Source Library Loan and Book Search software.
- ▶ Working ongoing on UI and Migration.

Digital Memory Project, Invenio and Zenodo

Digital Memory Project

- ▶ Digitization: 565 new videotapes, 580 audiotapes and 8'000 negatives shipped to the different scanning companies
- ▶ 2442 videotapes already treated; 1860 files checked, 1000+ people identified and 1175 public and 185 restricted videos loaded to CDS
- ▶ Collaboration(WIPO & CERN): run GPU server for speech to text transcription AI tool applied to HEP conferences in English: ~400 video subtitles ready

Invenio

- ▶ InvenioRDM kicked off with over 11 partners
(number of partners is increasing due to high interest)
- ▶ Invenio User Group Workshop at Open Repositories 2019 conference

Zenodo - Global roadmap published [here](#)

- ▶ DRYAD and ZENODO signed a Collaboration Agreement to provide a solid and unified solution for publishing of data, software and additional artifacts.

This Wednesday

*Computer Security Update
by Liviu Vâlsan*

Talks from CERN this week

Talks from CERN this week

Computer Security Update

- ▶ by Liviu Vâlsan

CERN Fixed Telephony Service Development

- ▶ by Thomas Baron

Challenges and opportunities when migrating CERN e-mail system to open source

- ▶ by Thomas Baron

CERN Linux services status update

- ▶ by Daniel Abad

CERN Cloud Infrastructure update

- ▶ by Daniel Abad

Self-service web hosting made easy with containers and Kubernetes Operators

- ▶ by Alexandre Lossent

Talks from CERN this week

CERN Computer Center Network evolution

- ▶ by Vincent Ducret

CERN Database on Demand Update

- ▶ by Ignacio Coterillo Coz

Current status of tape storage at CERN

- ▶ by Dr. Steven Murray

CERN Storage Evolution

- ▶ by Giuseppe Lo Presti

LHCb containers - Network overview

- ▶ by Daniele Pomponi

Open Compute Project 2019 Global Summit Report

- ▶ by Luca Atzori

Talks from CERN this week

WLCG/OSG Network Activities, Status and Plans

- ▶ by Marian Babik

Network Functions Virtualisation Working Group Update

- ▶ by Marian Babik

Unified Home folders migration or chronicles of a complex data migration

- ▶ by Vincent Nicolas Bippus

Zero-touch Windows 10 migration: dream or reality?

- ▶ by Vincent Nicolas Bippus

NOTED: Network-Optimized Transfer of Experimental Data

- ▶ by Coralie Busse-Grawitz

HEPiX CPU Benchmarking WG: status update

- ▶ by Domenico Giordano et al.

How to provide required resources for Run3 and Run4

- ▶ by Wayne Salter



home.cern