




Accélérateur de science



**CERN Storage and Data
Management Group (“IT-SD”)**

Overview

**Oliver Keeble
On behalf of the group
Storage and Data Management**

Including large chunks of a slide deck from Luca Mascetti

The three elements of the LHC program

- Accelerators, detectors and computing
- The collisions produce particles
- The particles disappear
- The data remains

ACCELERATORS

DETECTORS

COMPUTING

The LHC detectors are analogous to 3D cameras



The detectors measure the energy, direction and charge of new particles formed.



They take 40 million pictures a second. Only 1000/s are recorded and stored



The LHC detectors have been built by international collaborations covering all regions of the Globe.

A Laboratory for the world

A LABORATORY FOR THE WORLD

The Laboratory has emerged invigorated from 2022 – a year that was marked by an upsurge in the number of scientists coming to CERN and in ideas and innovation, fostered by a favourable health situation and the launch of the third run of the LHC. The efforts of the 17 000-member-strong CERN community were once again focused on the present, as many major projects were implemented during the year, as well as looking to the future, as the Laboratory continued to develop its plans for the coming decades. On closer inspection, this vast community is made up of 2700 staff members carrying out CERN's mission and 12 000 users affiliated with institutes in 82 countries. Many of these users have been able to take advantage of the LHC's restart to pursue the research conducted by the scientific collaborations based at CERN.



(CERN-PHOTO-202304-106-55)

BREAKDOWN OF CERN USERS ACCORDING TO THE COUNTRY OF THEIR HOME INSTITUTE, AS OF 31 DECEMBER 2022

NUMBER OF USERS: 11 860

MEMBER STATES (7147)

Austria 85 – Belgium 129 – Bulgaria 43 – Czech Republic 244 – Denmark 49 – Finland 90 – France 844 – Germany 1225 – Greece 119
Hungary 73 – Israel 64 – Italy 1527 – Netherlands 169 – Norway 79 – Poland 305 – Portugal 100 – Romania 109 – Serbia 33 – Slovakia 70
Spain 383 – Sweden 103 – Switzerland 406 – United Kingdom 898

ASSOCIATE MEMBER STATES IN THE PRE-STAGE TO MEMBERSHIP (69)

Cyprus 15 – Estonia 30 – Slovenia 24

ASSOCIATE MEMBER STATES (382)

Croatia 38 – India 132 – Latvia 16 – Lithuania 14 – Pakistan 35 – Türkiye 122 – Ukraine 25

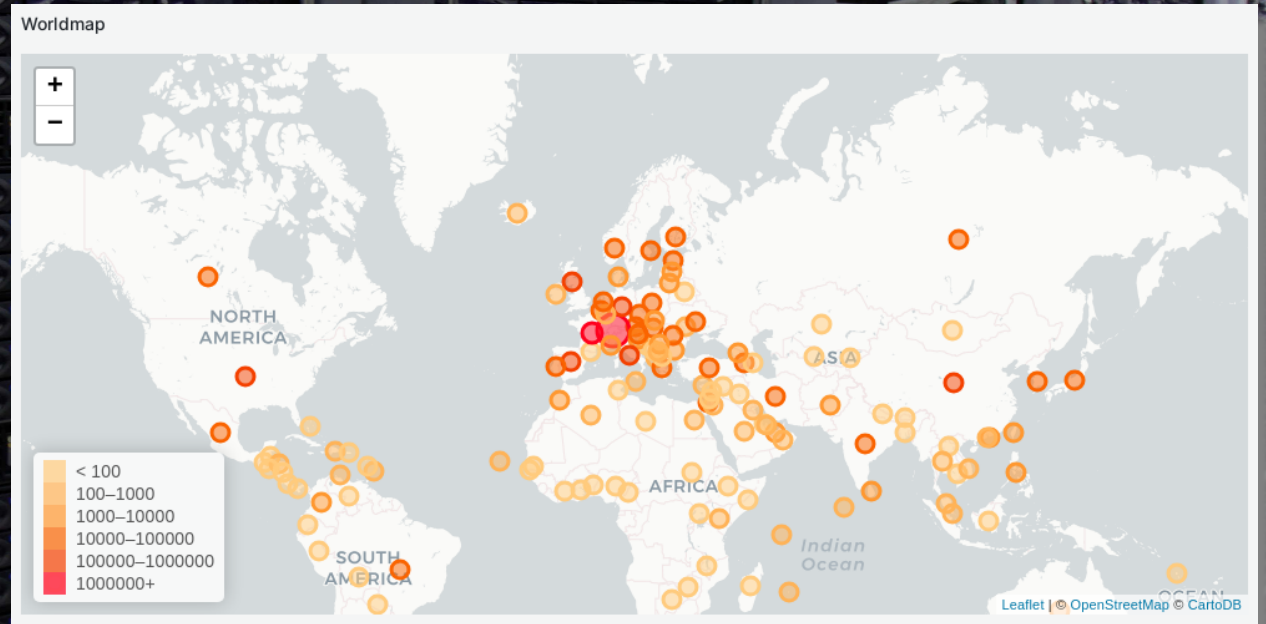
OBSERVERS (2991)

Japan 216 – Russian Federation 873 (the Observer status of the Russian Federation has been suspended in accordance with the Resolution adopted by the CERN Council on 8 March 2022) – United States of America 1902

OTHER COUNTRIES (1271)

Algeria 2 – Argentina 13 – Armenia 8 – Australia 21 – Azerbaijan 2 – Bahrain 4 – Belarus 18 – Brazil 122 – Canada 199 – Chile 34
Colombia 21 – Costa Rica 2 – Cuba 3 – Ecuador 4 – Egypt 20 – Georgia 32 – Hong Kong 15 – Iceland 3 – Indonesia 5 – Iran 11 – Ireland 5
Jordan 5 – Kuwait 4 – Lebanon 13 – Madagascar 1 – Malaysia 4 – Malta 1 – Mexico 49 – Montenegro 4 – Morocco 19 – New Zealand 5
Nigeria 1 – Oman 1 – Palestine 1 – People's Republic of China 333 – Peru 2 – Philippines 1 – Republic of Korea 147 – Singapore 2
South Africa 52 – Sri Lanka 10 – Taiwan 45 – Thailand 17 – Tunisia 2 – United Arab Emirates 7 – Vietnam 3

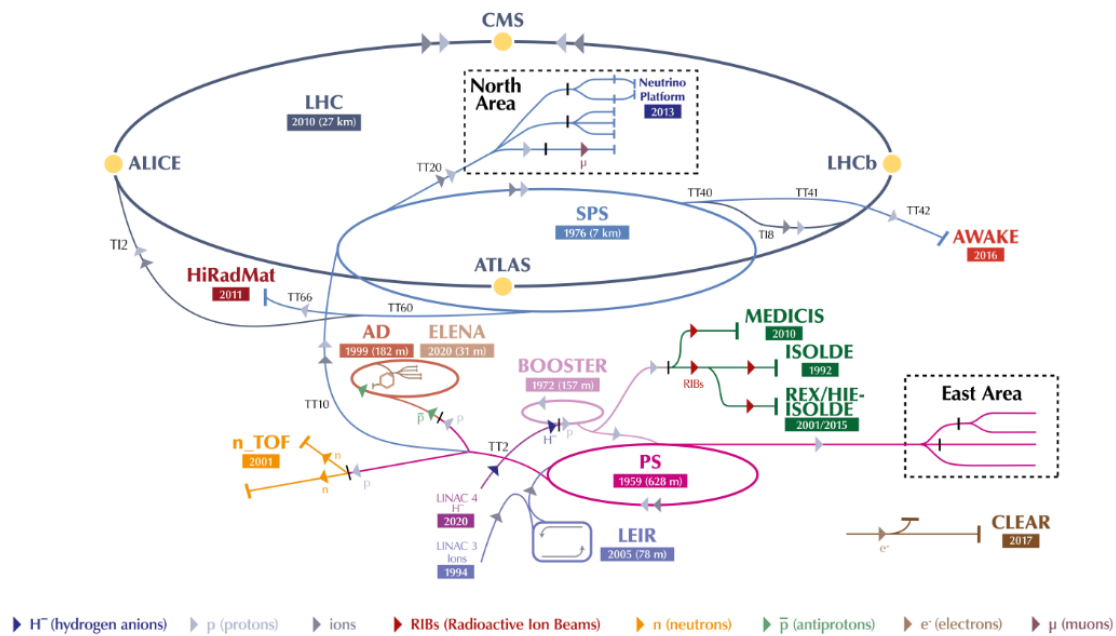
CERNBox client locations



What brings us here?

CERN's Scientific Data Management stack in context

The CERN accelerator complex Complexe des accélérateurs du CERN



LHC - Large Hadron Collider // SPS - Super Proton Synchrotron // PS - Proton Synchrotron // AD - Antiproton Decelerator // CLEAR - CERN Linear Electron Accelerator for Research // AWAKE - Advanced WAKEfield Experiment // ISOLDE - Isotope Separator OnLine // REX/HIE-ISOLDE - Radioactive Experiment/High Intensity and Energy ISOLDE // MEDICIS // LEIR - Low Energy Ion Ring // LINAC - LINear ACcelerator // n_TOF - Neutrons Time Of Flight // HiRadMat - High-Radiation to Materials // Neutrino Platform

Open Science Policies

CERN Open Science Policy

The CERN Open Science Policy covers all elements of the Open Science relevant to CERN. This includes, in particular open access to research publications, data, software and hardware, as well as research integrity, infrastructure, education and outreach activities supporting or enabling open science practices.

“ Supported by long term financial investments from its Member and Associate Member States, with significant contributions also from non-Member States, CERN is committed to the advancement of science and the wide dissemination of knowledge by embracing and promoting practices making scientific research more open, collaborative, and responsive to societal changes. ... CERN accordingly recognizes the holistic practice of open science as one of its guiding principles. ”

Last revision: Oct 2022

[Access to the complete policy text](#)

[Access to the OS Policy: Implementation Plan](#)

CERN Open Access Policy

The CERN Open Access Policy defines the principles and processes through which CERN authors can publish their peer-reviewed articles Open Access. A dedicated [website](#) also provides authors with additional resources to find the easiest route to comply with the policy.

“ CERN authors are required to publish all of their peer-reviewed primary research articles open access (by default under a Creative Commons attribution license, i.e. CC-BY-4.0) ”

Last revision: May 2021

[Access to the complete policy text](#)

CERN LHC Open Data Policy

This policy relates to the data collected by the LHC experiments for the main physics programme of the LHC — high-energy proton-proton and heavy-ion collision data. The foreseen use cases of the Open Data include reinterpretation and reanalysis of physics results, education and outreach, data analysis for technical and algorithmic developments, and new physics research.

“ Making data available responsibly, at different levels of abstraction and at different points in time, allows the maximum realisation of their scientific potential and the fulfillment of the collective moral and fiduciary responsibility to member states and the broader global scientific community. ”

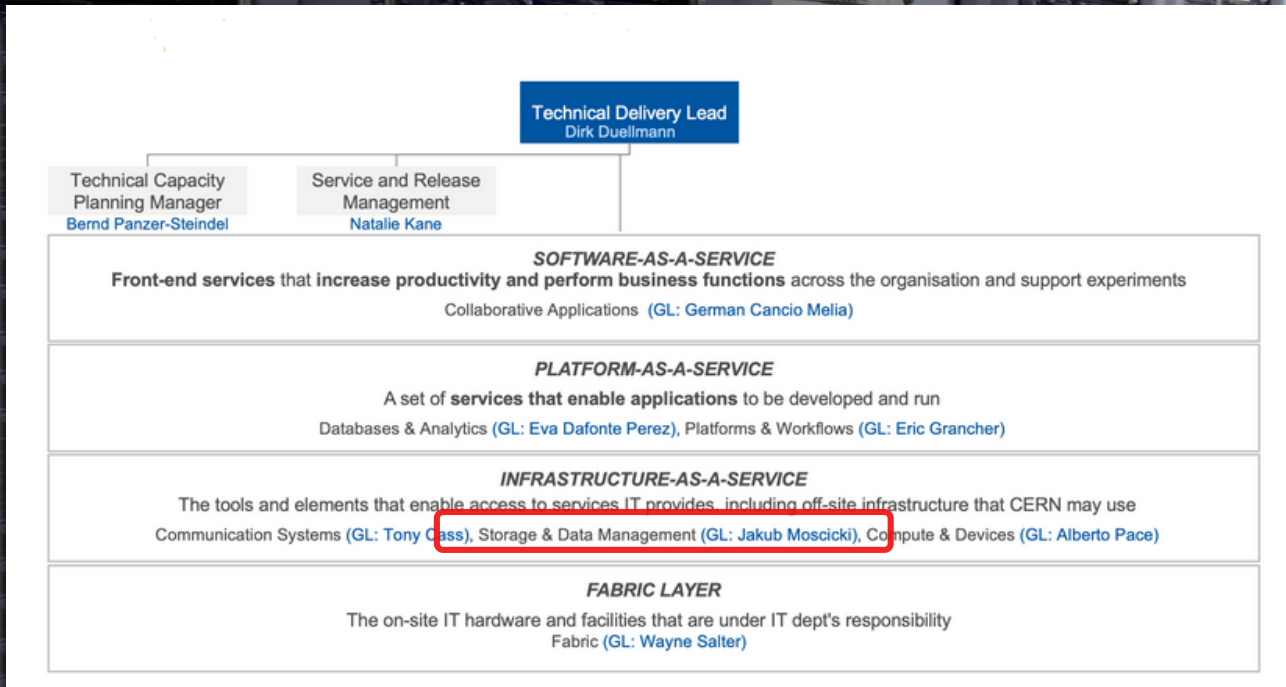
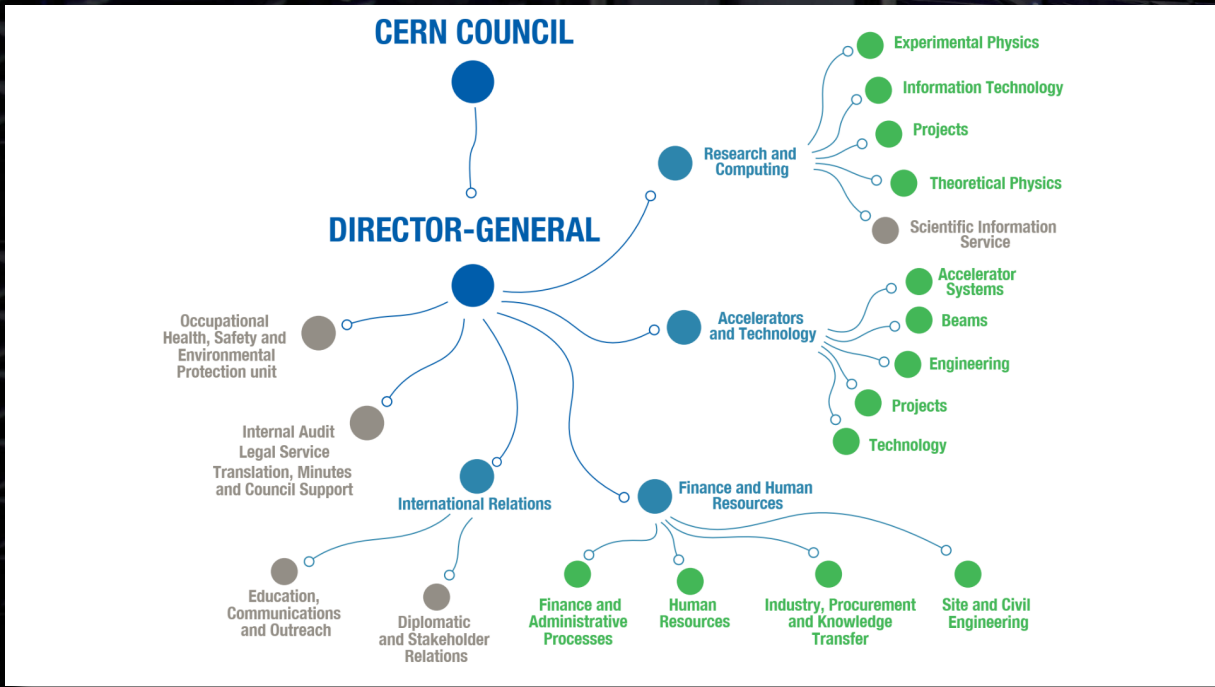
<https://opensource.web.cern.ch/>



CERN's structure

The Organisation


The IT dept ("technical delivery")



Storage in High Energy Physics




Archival & Backup Storage



Storage for Data Acquisition



Storage for Home Directories




Storage for HPC



Storage for Applications



Private Cloud Storage



Storage for GRID Computing




Public Cloud Storage



Storage for Software Distribution



Storage for Data Analytics



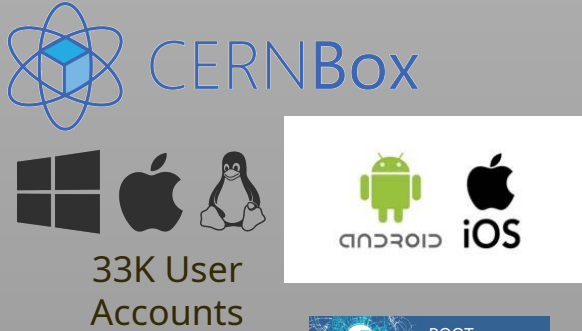
Storage for Physics Analysis



Storage for Sync&Share

General Home, Media & Project spaces

CERNBox



33K User Accounts

ROOT Data Analysis Framework

Microsoft 365

diagrams.net

CodiMD

indico

zenodo

restic

CDS

AV



physics analysis
xrd

Windows

Samba DFS

fusex

EOS File Store

Data Analysis & Physics Infra



WLCG

~26PB Physics Data worldwide

150K+ clients worldwide

Backups Preveessin Zone

restic

indico

DPM

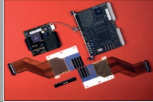
CVM FS

S3 RBD Ceph FS

CEPH Object Store

General Infrastructure & Services

EP-ESE



OPENSIFT

puppet

openstack.


CERN Tape Archive

GitLab

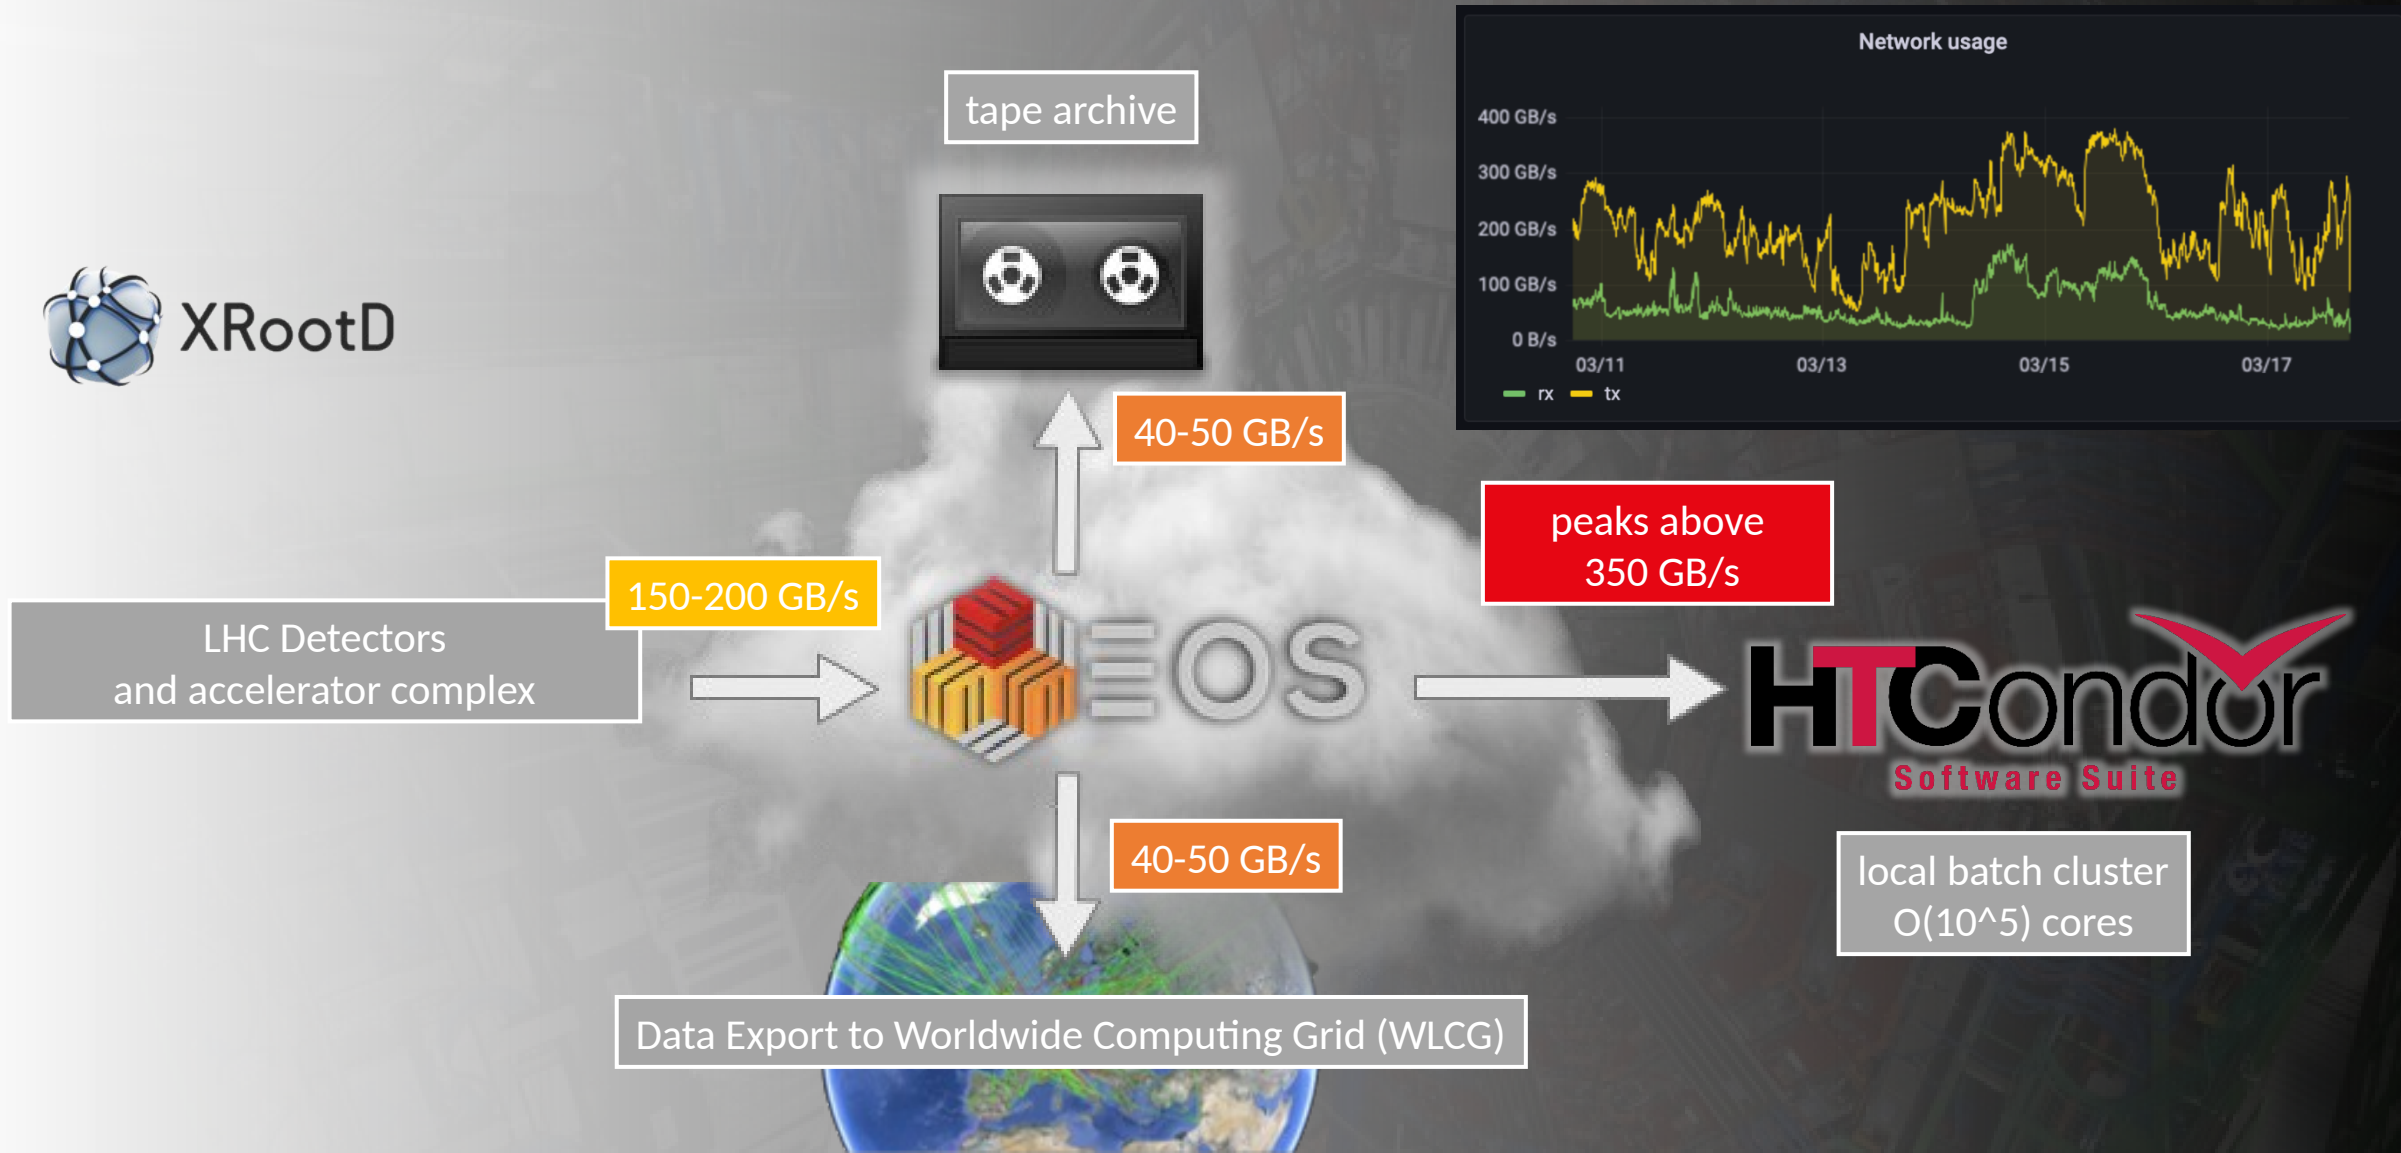
AFS FILER

HPC

CERN IT - Operated Disk Storage Capacity



EOS in action



tape archive

40-50 GB/s

150-200 GB/s

LHC Detectors
and accelerator complex



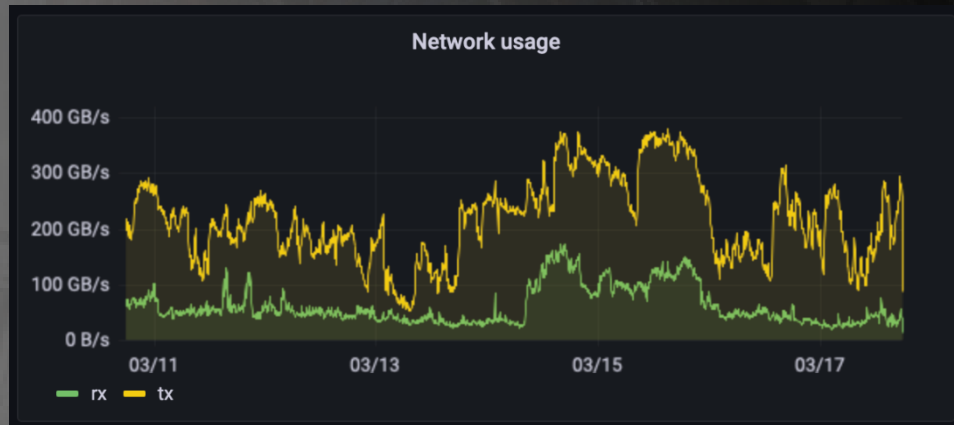
40-50 GB/s

Data Export to Worldwide Computing Grid (WLCG)

peaks above
350 GB/s



local batch cluster
O(10⁵) cores



Tape Storage (CTA)



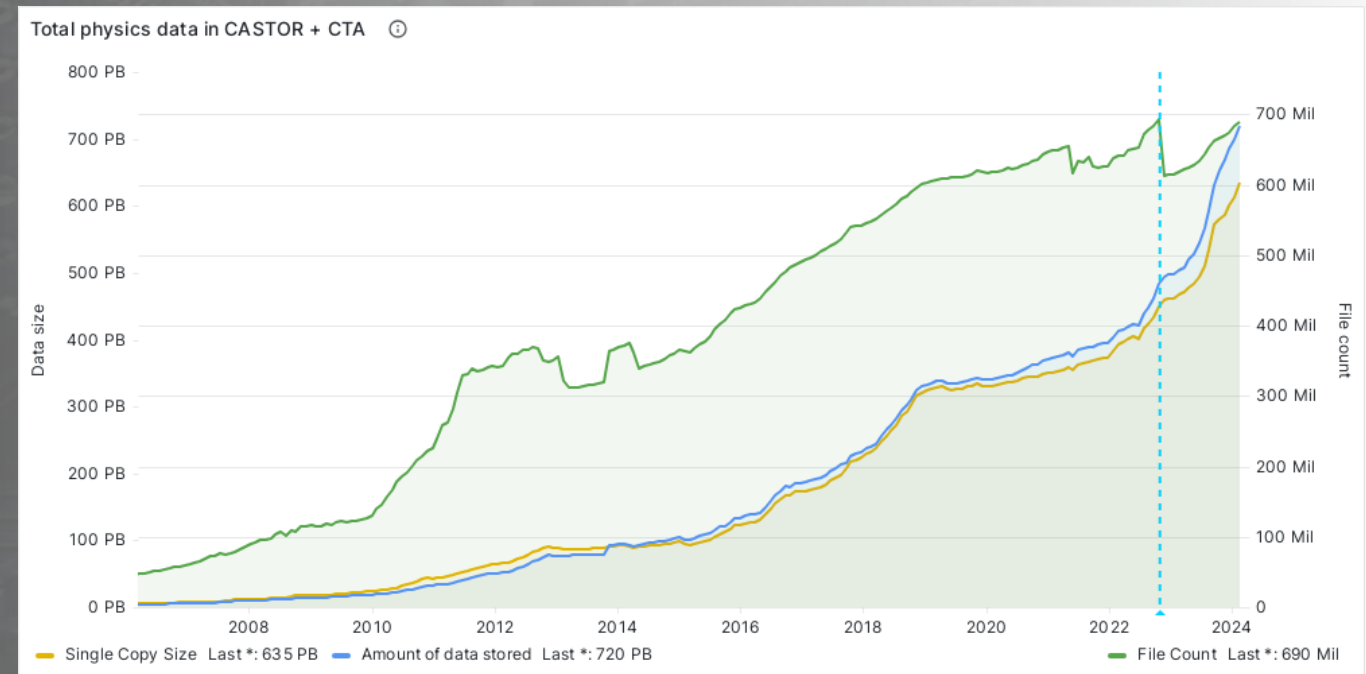
The EOS open Storage platform is natively used as a namespace and as a pure SSD disk storage for CTA

Conceived as a fast buffer to the tape system

- File residency on disk is transitional
- Basis of Run3 data archival
- Ready for Hi-Lumi LHC

Infrastructure

- Libraries:
 - 3x IBM TS4500
 - 2x Spectra Logic TFinity
- Drives:
 - ~180 Drives
 - Mixture of LTO and Enterprise
 - 20 new TS1170 drives deployed



Additionally we run a
IBM Storage Protect backup service





690M files stored, 720 PB of actual data

File Transfer Service



LAST DATA UPDATE

9.7 MB Downloaded Wednesday, 11 September 2019 14:05:12
Last transfer was on : Monday, 29 July 2019 08:00:00



Supports many Virtual Organizations

DATA TRANSFER CONSOLE


Open Source Software developed at CERN for Reliable, multiprotocol Large-scale Data Transfers

FTS distributes the majority of the LHC's data across the worldwide LHC Computing Grid

More than 1 Exabyte of Data and 1 Billion files transferred world-wide every year.

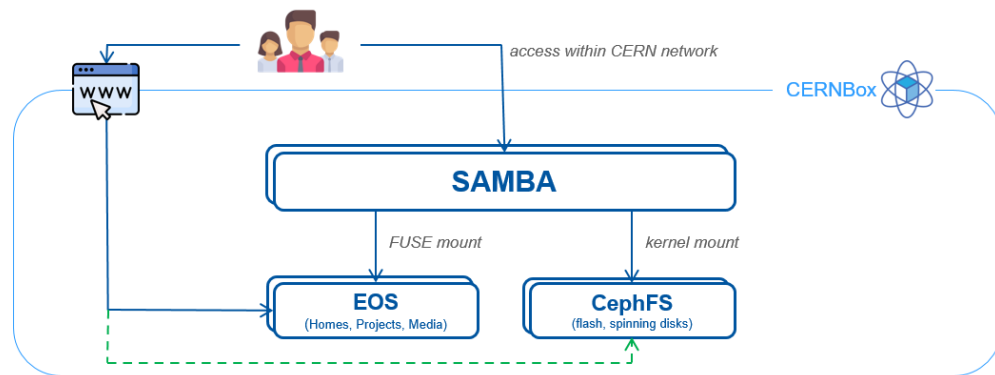
Services for Windows

500M files in DFS



SAMBA gateway in CERNBox

- Provides SAMBA access for Windows clients to CERNBox
- High-available solution implemented based on the CTDB daemon



- Windows clients are served natively by DFS
- EOS (for users and projects) is re-exported via Samba
- Ongoing migration to CERNBox
 - Most to EOS backend
 - New CephFS backend deployed for performance sensitive clients (e.g. engineering tools)

AFS



Home directory shared filesystem

Uses Ceph rbd

487 TB used and 4.4 B files (avg files size:160kB)

~50k Volumes

~7k clients



Ceph: rbd, S3 & CephFS

Ceph

Block: Openstack RBD Volumes for Virtual Machines

Objects: Backup target, native applications using S3/SWIFT

Filesystem: Openstack Manila Share, NFS-like share

Storage for IT Infrastructure:

- OpenStack, K8s/OKD, GitLab, Container registries...
- AFS, CVMFS, NFS Filers

Experiments and End-users directly interfacing:

- ATLAS Event Service, CMS Web
- The team also manages an NFS service...

Service	Managed Capacity
CephFS	14 PB
Ceph S3	33 PB
Ceph RBD	26 PB

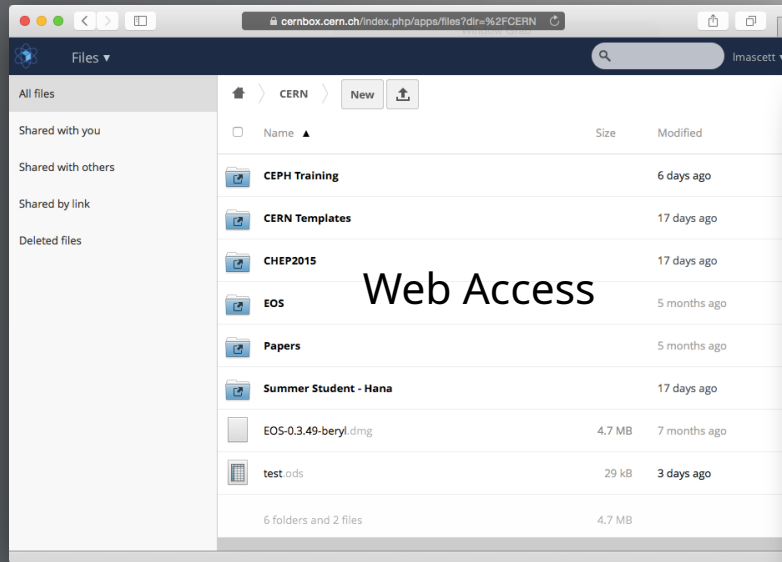
CERNBox




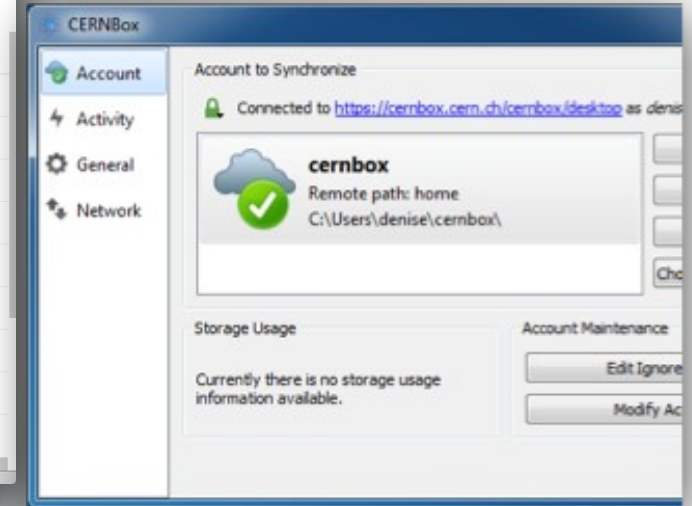
powered by



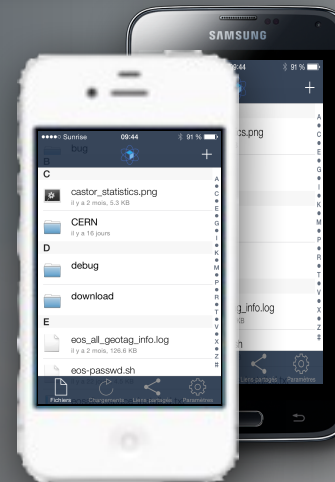
- CERN Sync & Share solution
- Offline access to all EOS data
- Central Hub for CERN data
- Main WebApp Integrator
- Supported on multiple OS/Devices



Sync Client



WebDAV




Mobile App



Directly from the storage backend (xroot, http, s3, ...)

SAMBA



Collaborations with Experimental Physics dept

• CVMFS



- Developed by the EP (Experimental Physics) dept, run by IT-SD
- Content Distribution Network
- Mounted r/o filesystem
- Loading on demand
- Used for software
- ... and now containers

• SWAN



- Turn-key data analysis platform
 - Based on Jupyter Notebooks
- Run by the “Database & Analytics” Group with IT-SD collaboration
- Support for ROOT/C++, Python, R, Octave
- Fully integrated in CERN ecosystem
 - Storage on EOS, Sharing with CERNBox
 - Software provided by CVMFS
 - Massive computations on Spark
 - Dask clusters on HTCondor



Summary

- The Storage and Data Management Group has a diverse and mature portfolio
- These projects and services are all used in solving the lab's data challenges
 - All this stuff works together
- The group represents the continuation of more than 20 years of work in storage and data management
- We are happy to showcase our software, services and collaborations!
- Enjoy the show!



Accélérateur de science