
CERN
Accélérateur de science

EOS Open Storage for Science

Luca Mascetti
CERN IT Storage

luca.mascetti@cern.ch

CERN is the world's biggest laboratory for particle physics.

Our goal is to understand the most fundamental particles and laws of the universe.

Located near Geneva on either side of the Swiss French border

How do we do it?

- We build the largest machines to study the smallest particles in the universe
- We develop technology to advance the limits of what is possible
- We perform world-class research in theoretical and experimental particle physics

ACCELERATORS

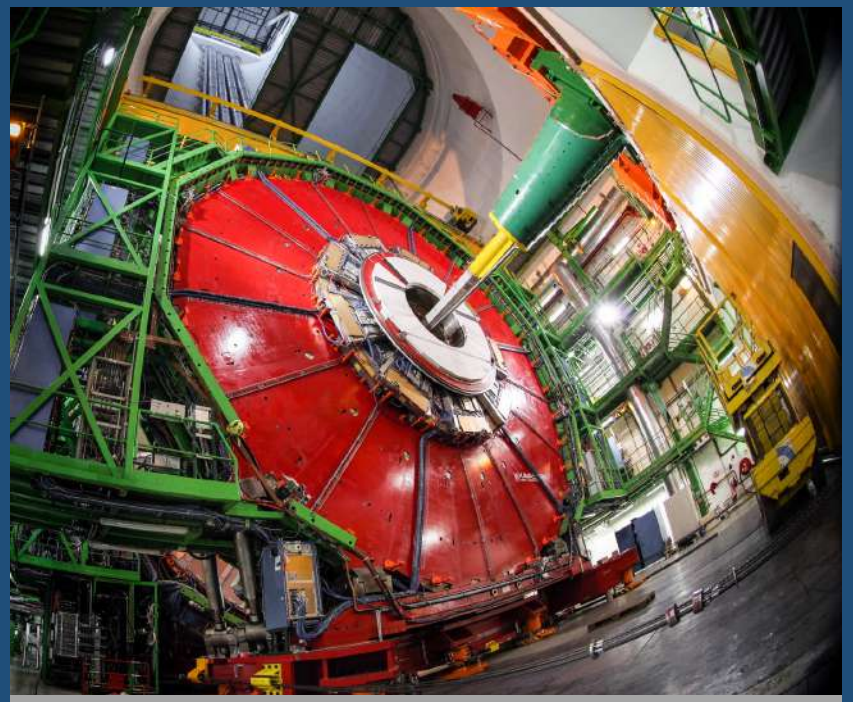

DETECTORS

COMPUTING

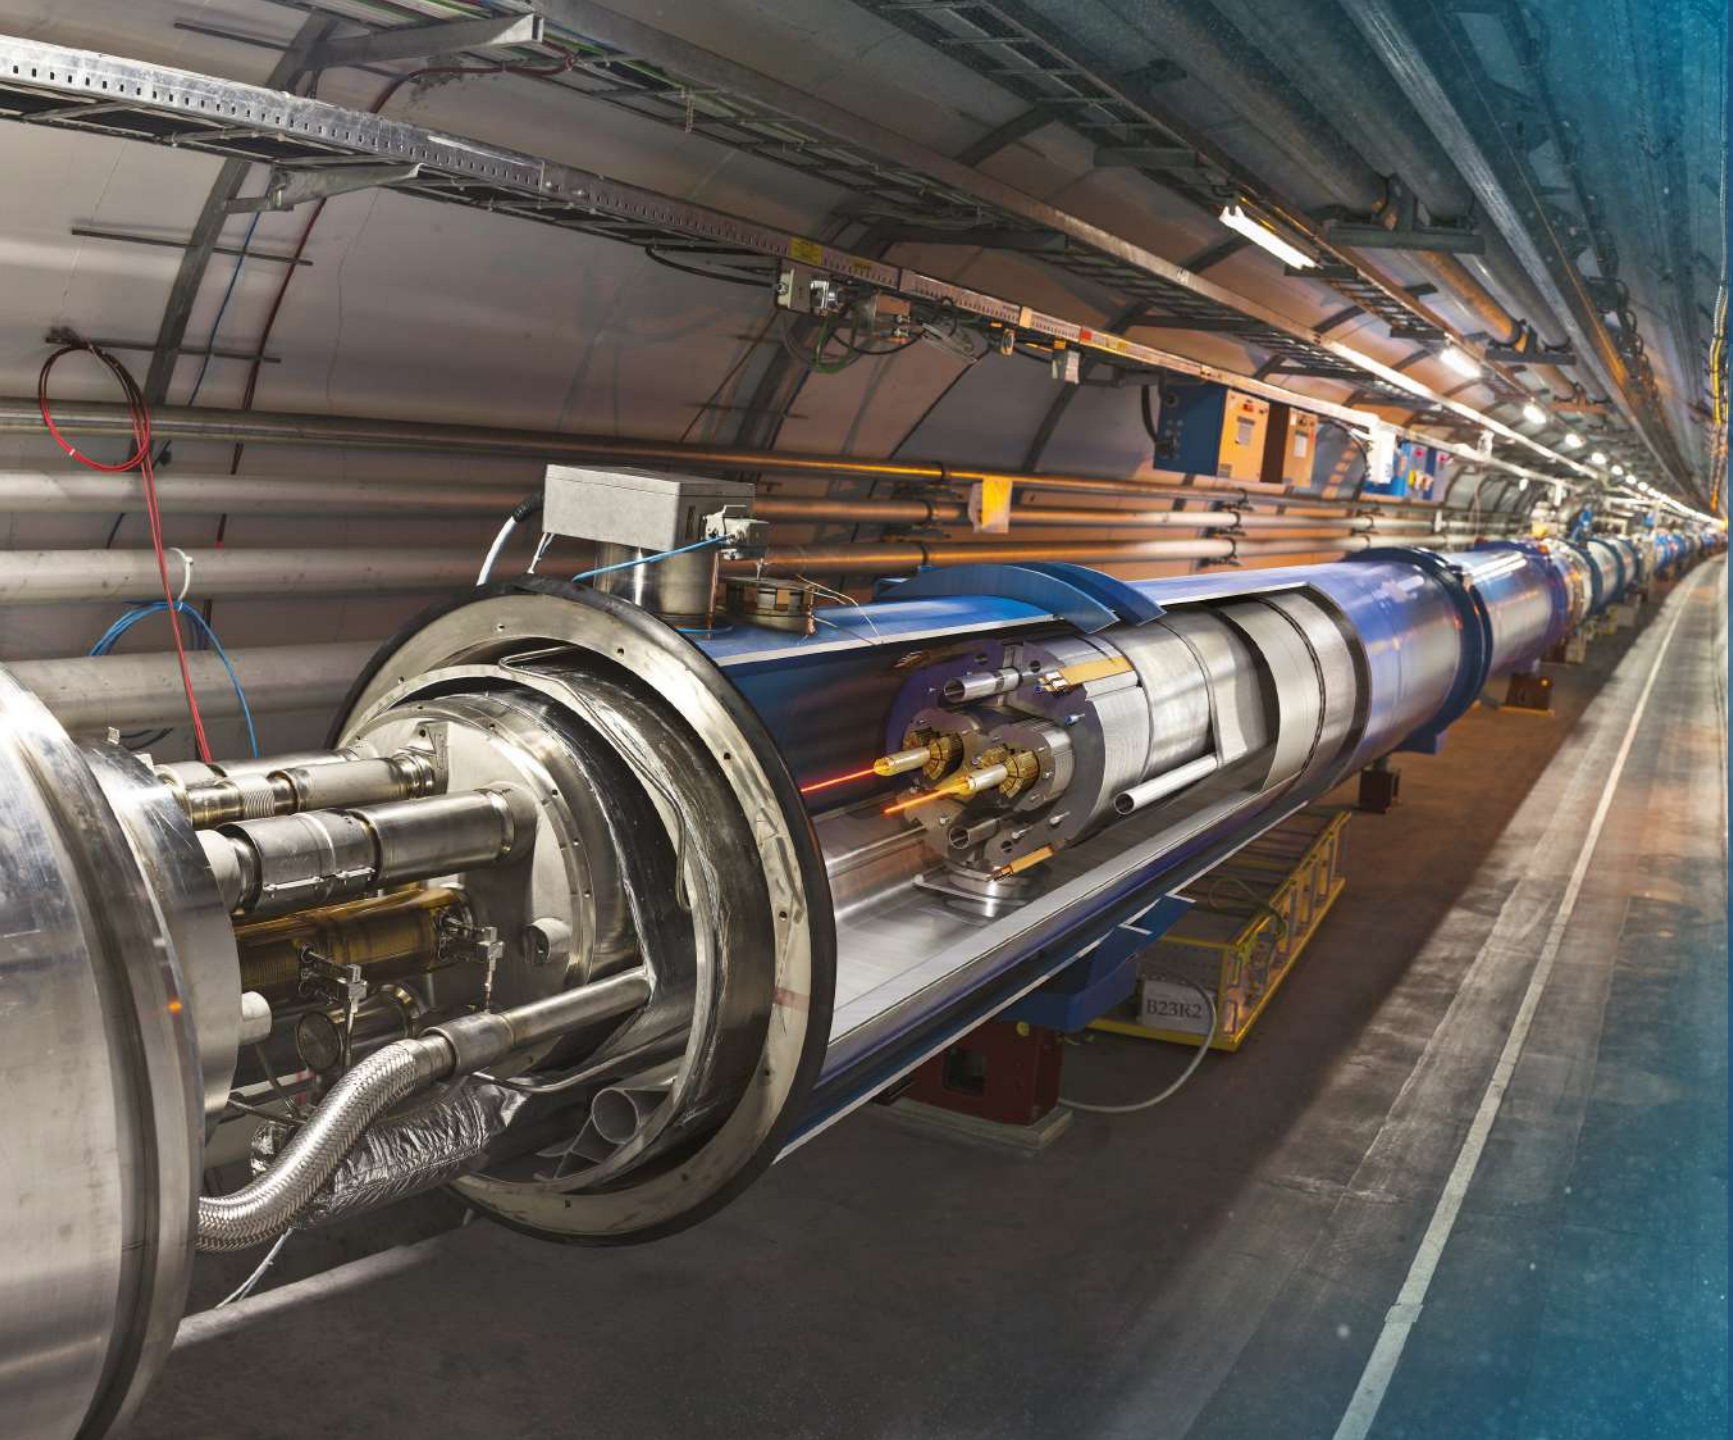

Large Hadron Collider (LHC)

- 27 km in circumference
- About 100 m underground
- Superconducting magnets steer the particles around the ring
- Particles are accelerated to close to the speed of light

The LHC detectors are analogous to 3D cameras

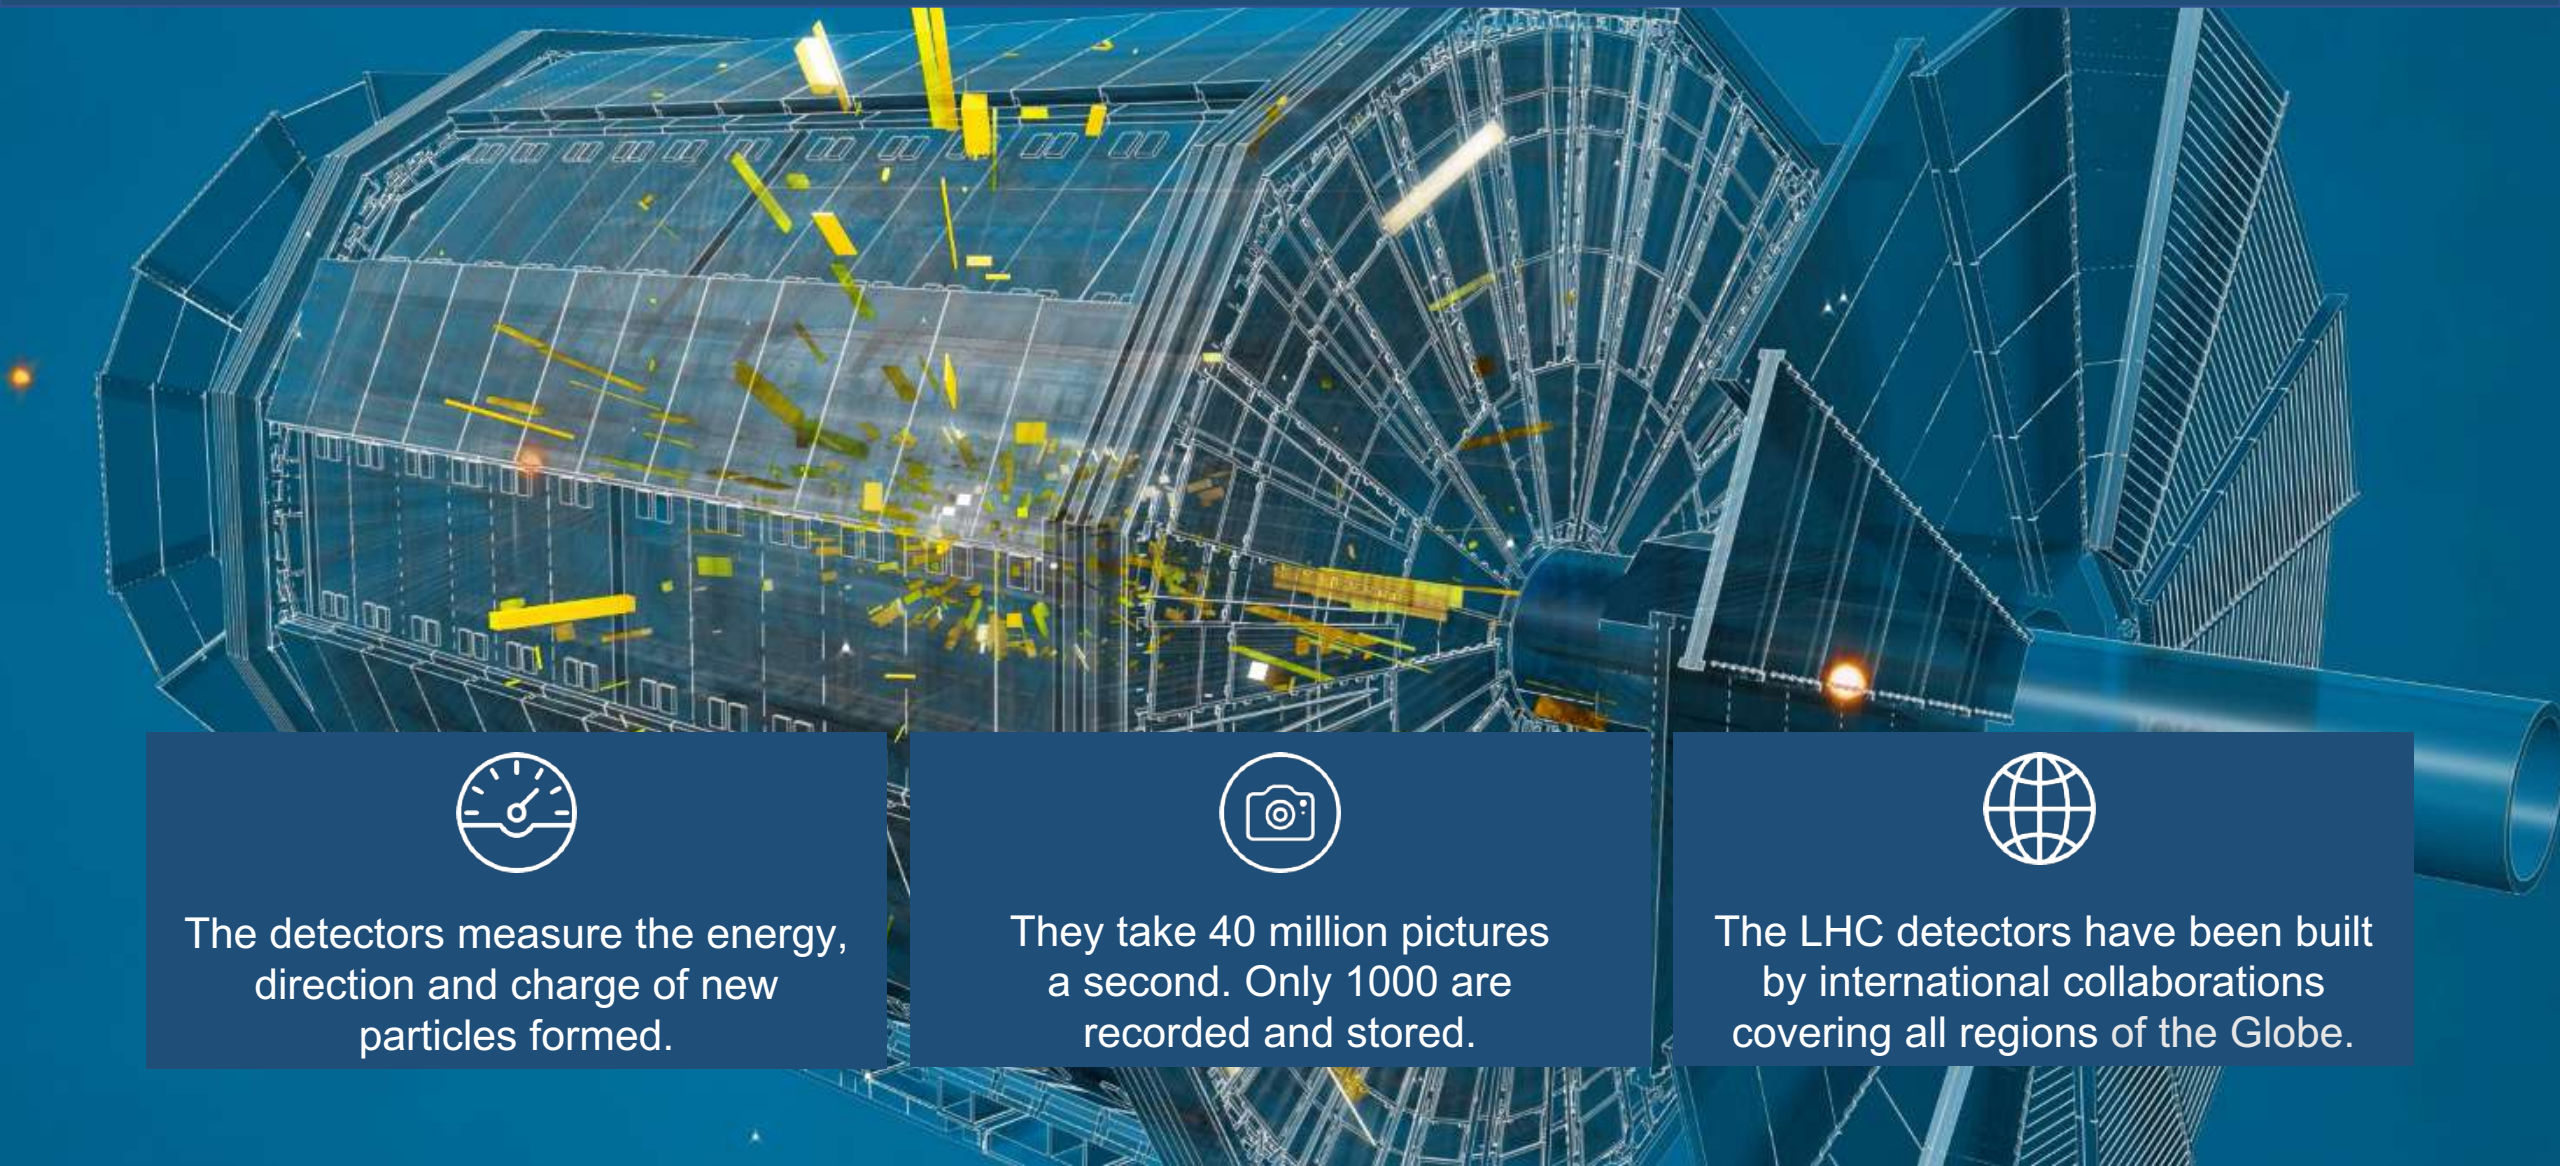

The detectors measure the energy, direction and charge of new particles formed.

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

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

The Worldwide LHC Computing Grid (WLCG)

Used to store, distribute, process and analyse data.

1 million processing cores in about 160 data centres and 42 countries.

More than 1000 Petabytes of CERN data stored world-wide.

Computing

A central role in the fulfilment of CERN's mission

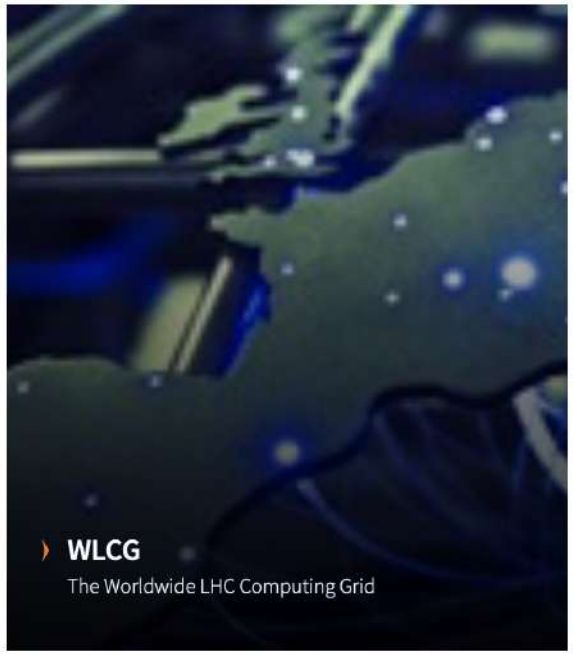

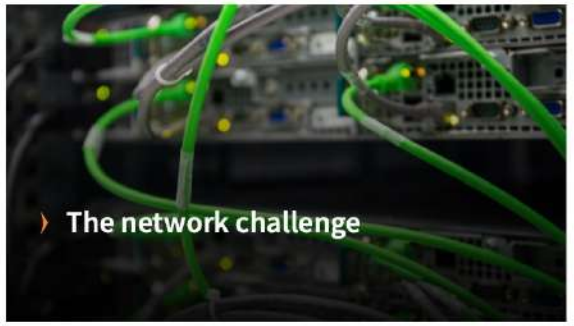

CERN openlab

CERN OPENLAB'S MISSION

- Evaluate state-of-the-art technologies in a challenging environment and improve them.
- Test in a research environment today technologies that will be used in many business sectors tomorrow.
- Train the next generation of engineers and researchers.
- Promote education and cultural exchanges.
- Communicate results and reach new audiences.
- Collaborate and exchange ideas to create knowledge and innovation.

DRIVING INNOVATION FOR 20 YEARS

SET UP 2001

I 2003

II 2006

III 2009

IV 2012

V 2015

VI 2018

2021 Seventh Phase VII

Comtrade Three-Sixty

In 2015 Comtrade 360 joined CERN openlab as associate member

What is EOS ?

Open-Source Storage designed and developed in CERN IT

Disk-based distributed filesystem Elastic, Adaptable and Scalable

Software solution for data recording, user analysis and data processing

Supports thousands of parallel clients

Multiprotocol support (native xrood, FUSE, HTTP, WebDAV, CIFS)

Offers a variety of authentication methods (KRB5, X509, SharedSecret, tokens, unix)

EOS Project History

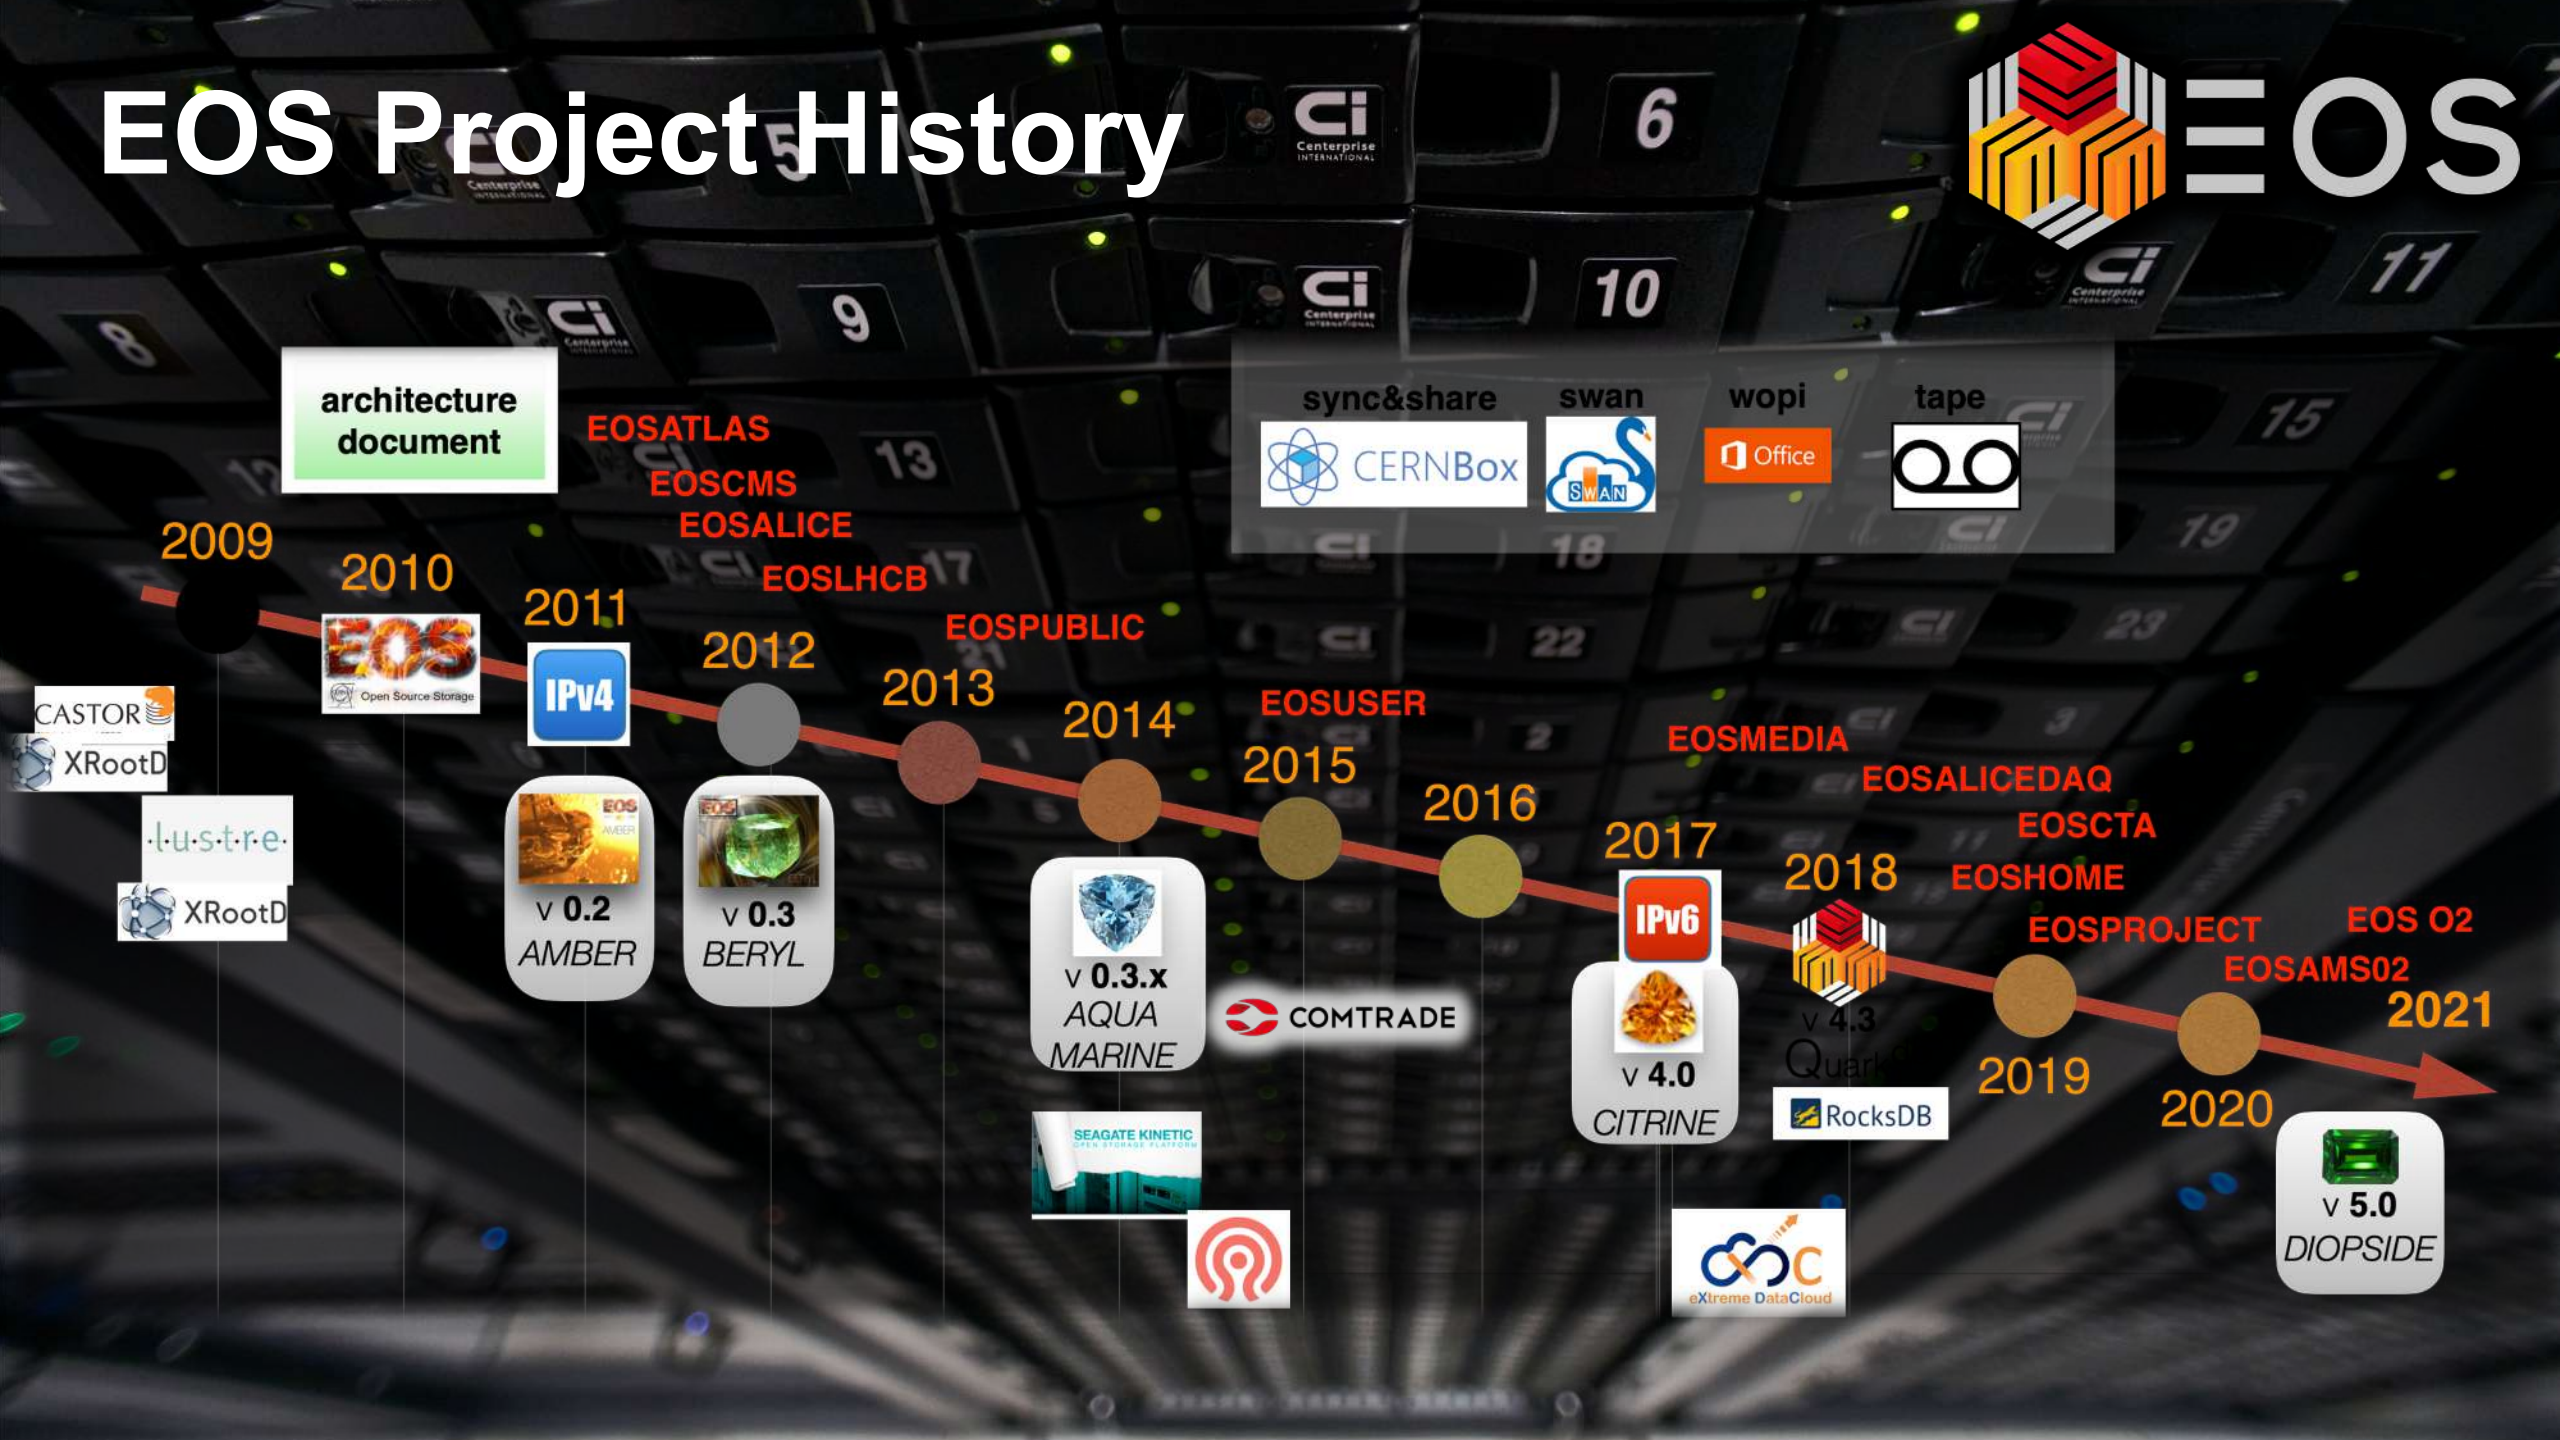

EOS @ CERN

Total Space
600 PB

Files Stored
~7 Bil

 # Storage Nodes
~1600

 # Disks
~80000

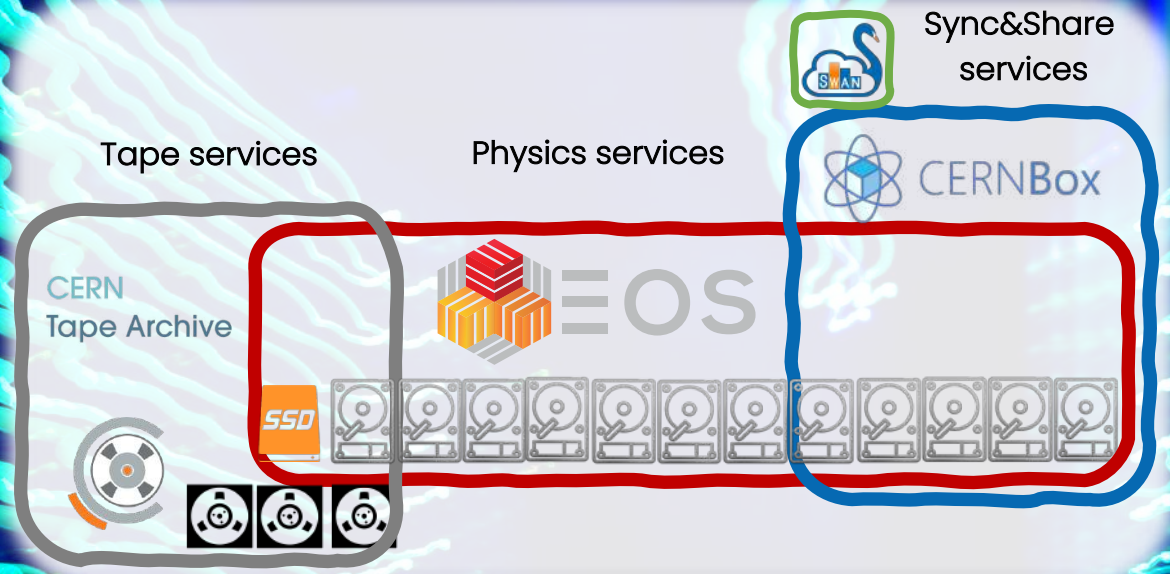

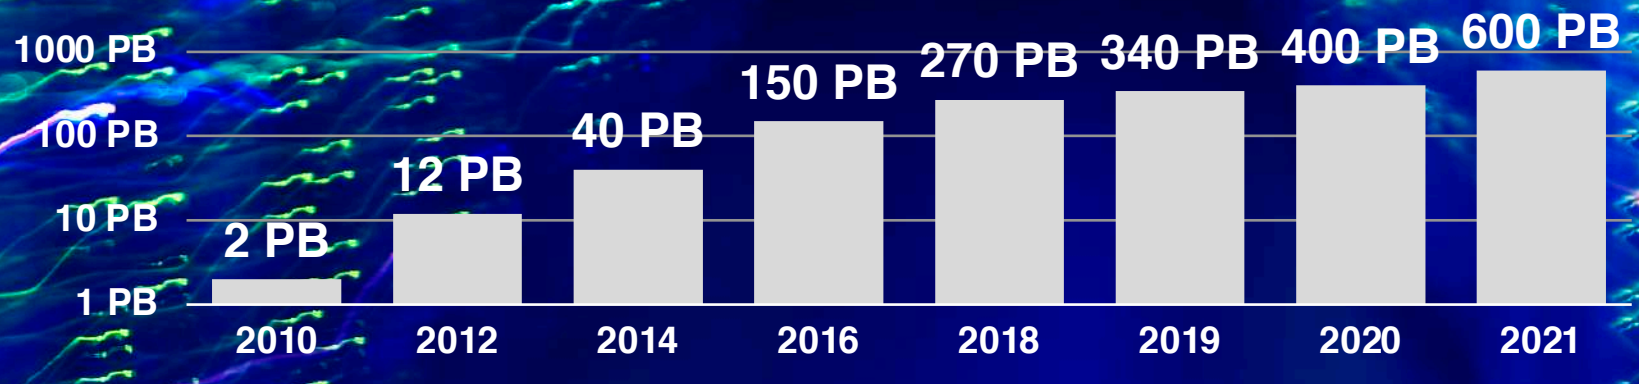

EOS Service for Physics

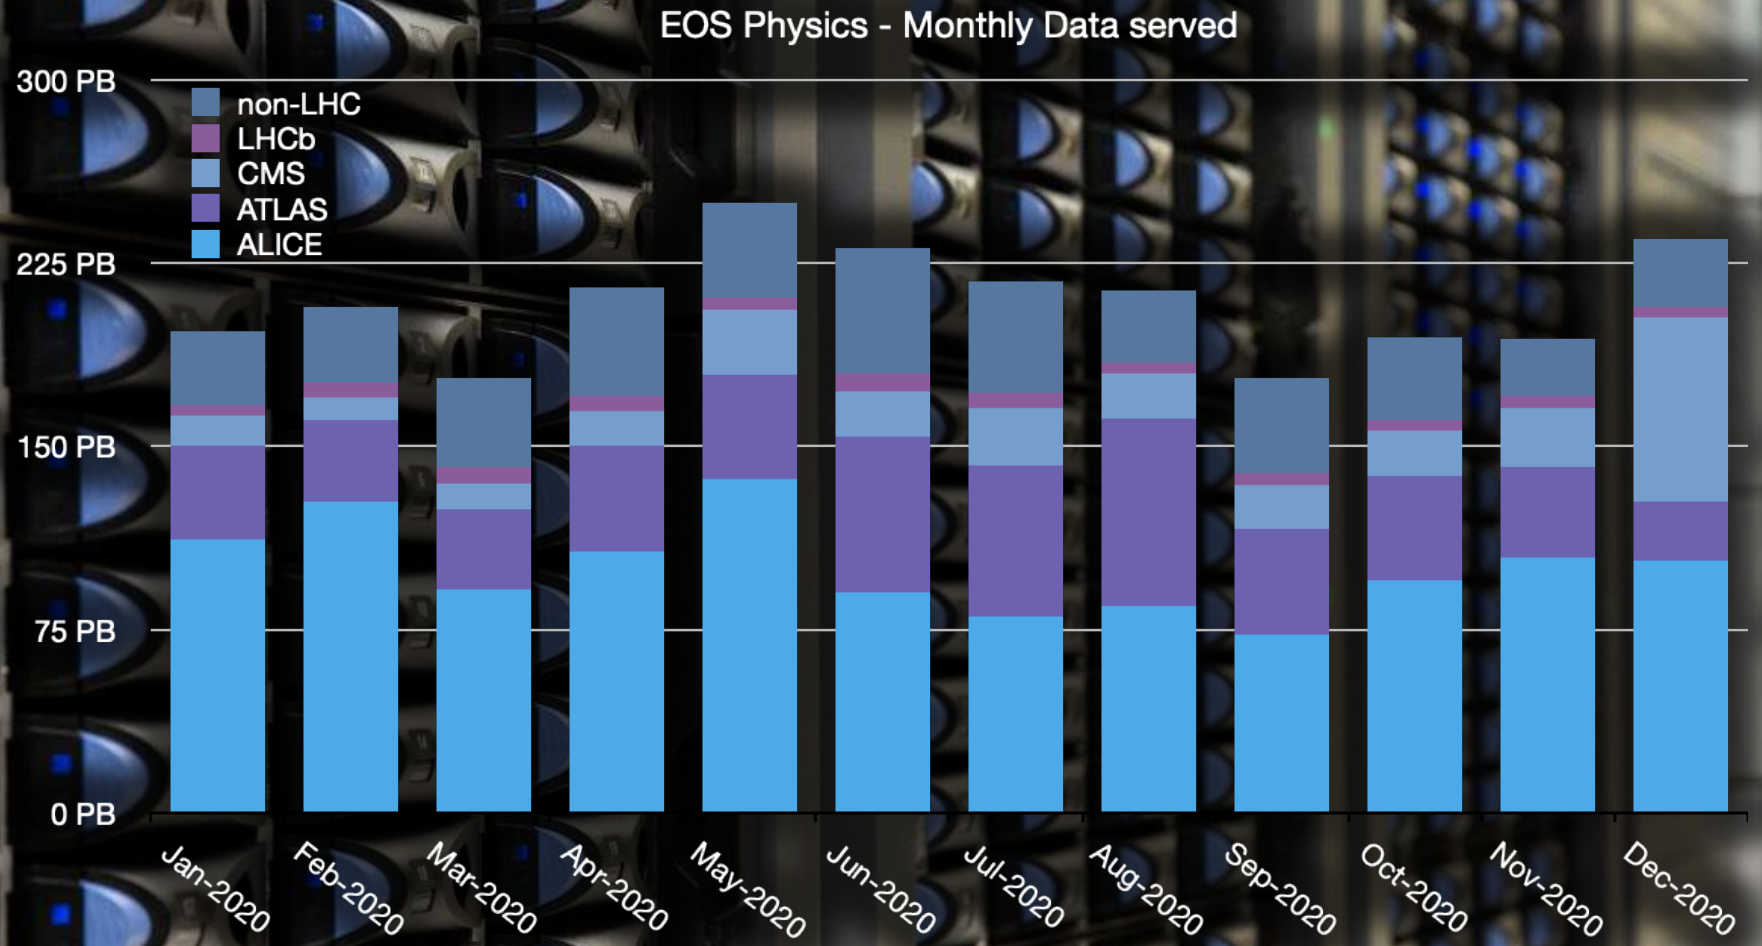

~ 2.5 Exabytes of data delivered in 2020
so far (Sep 2021) ~1.8 Exabytes delivered in 2021

CERN Physics Data Recording

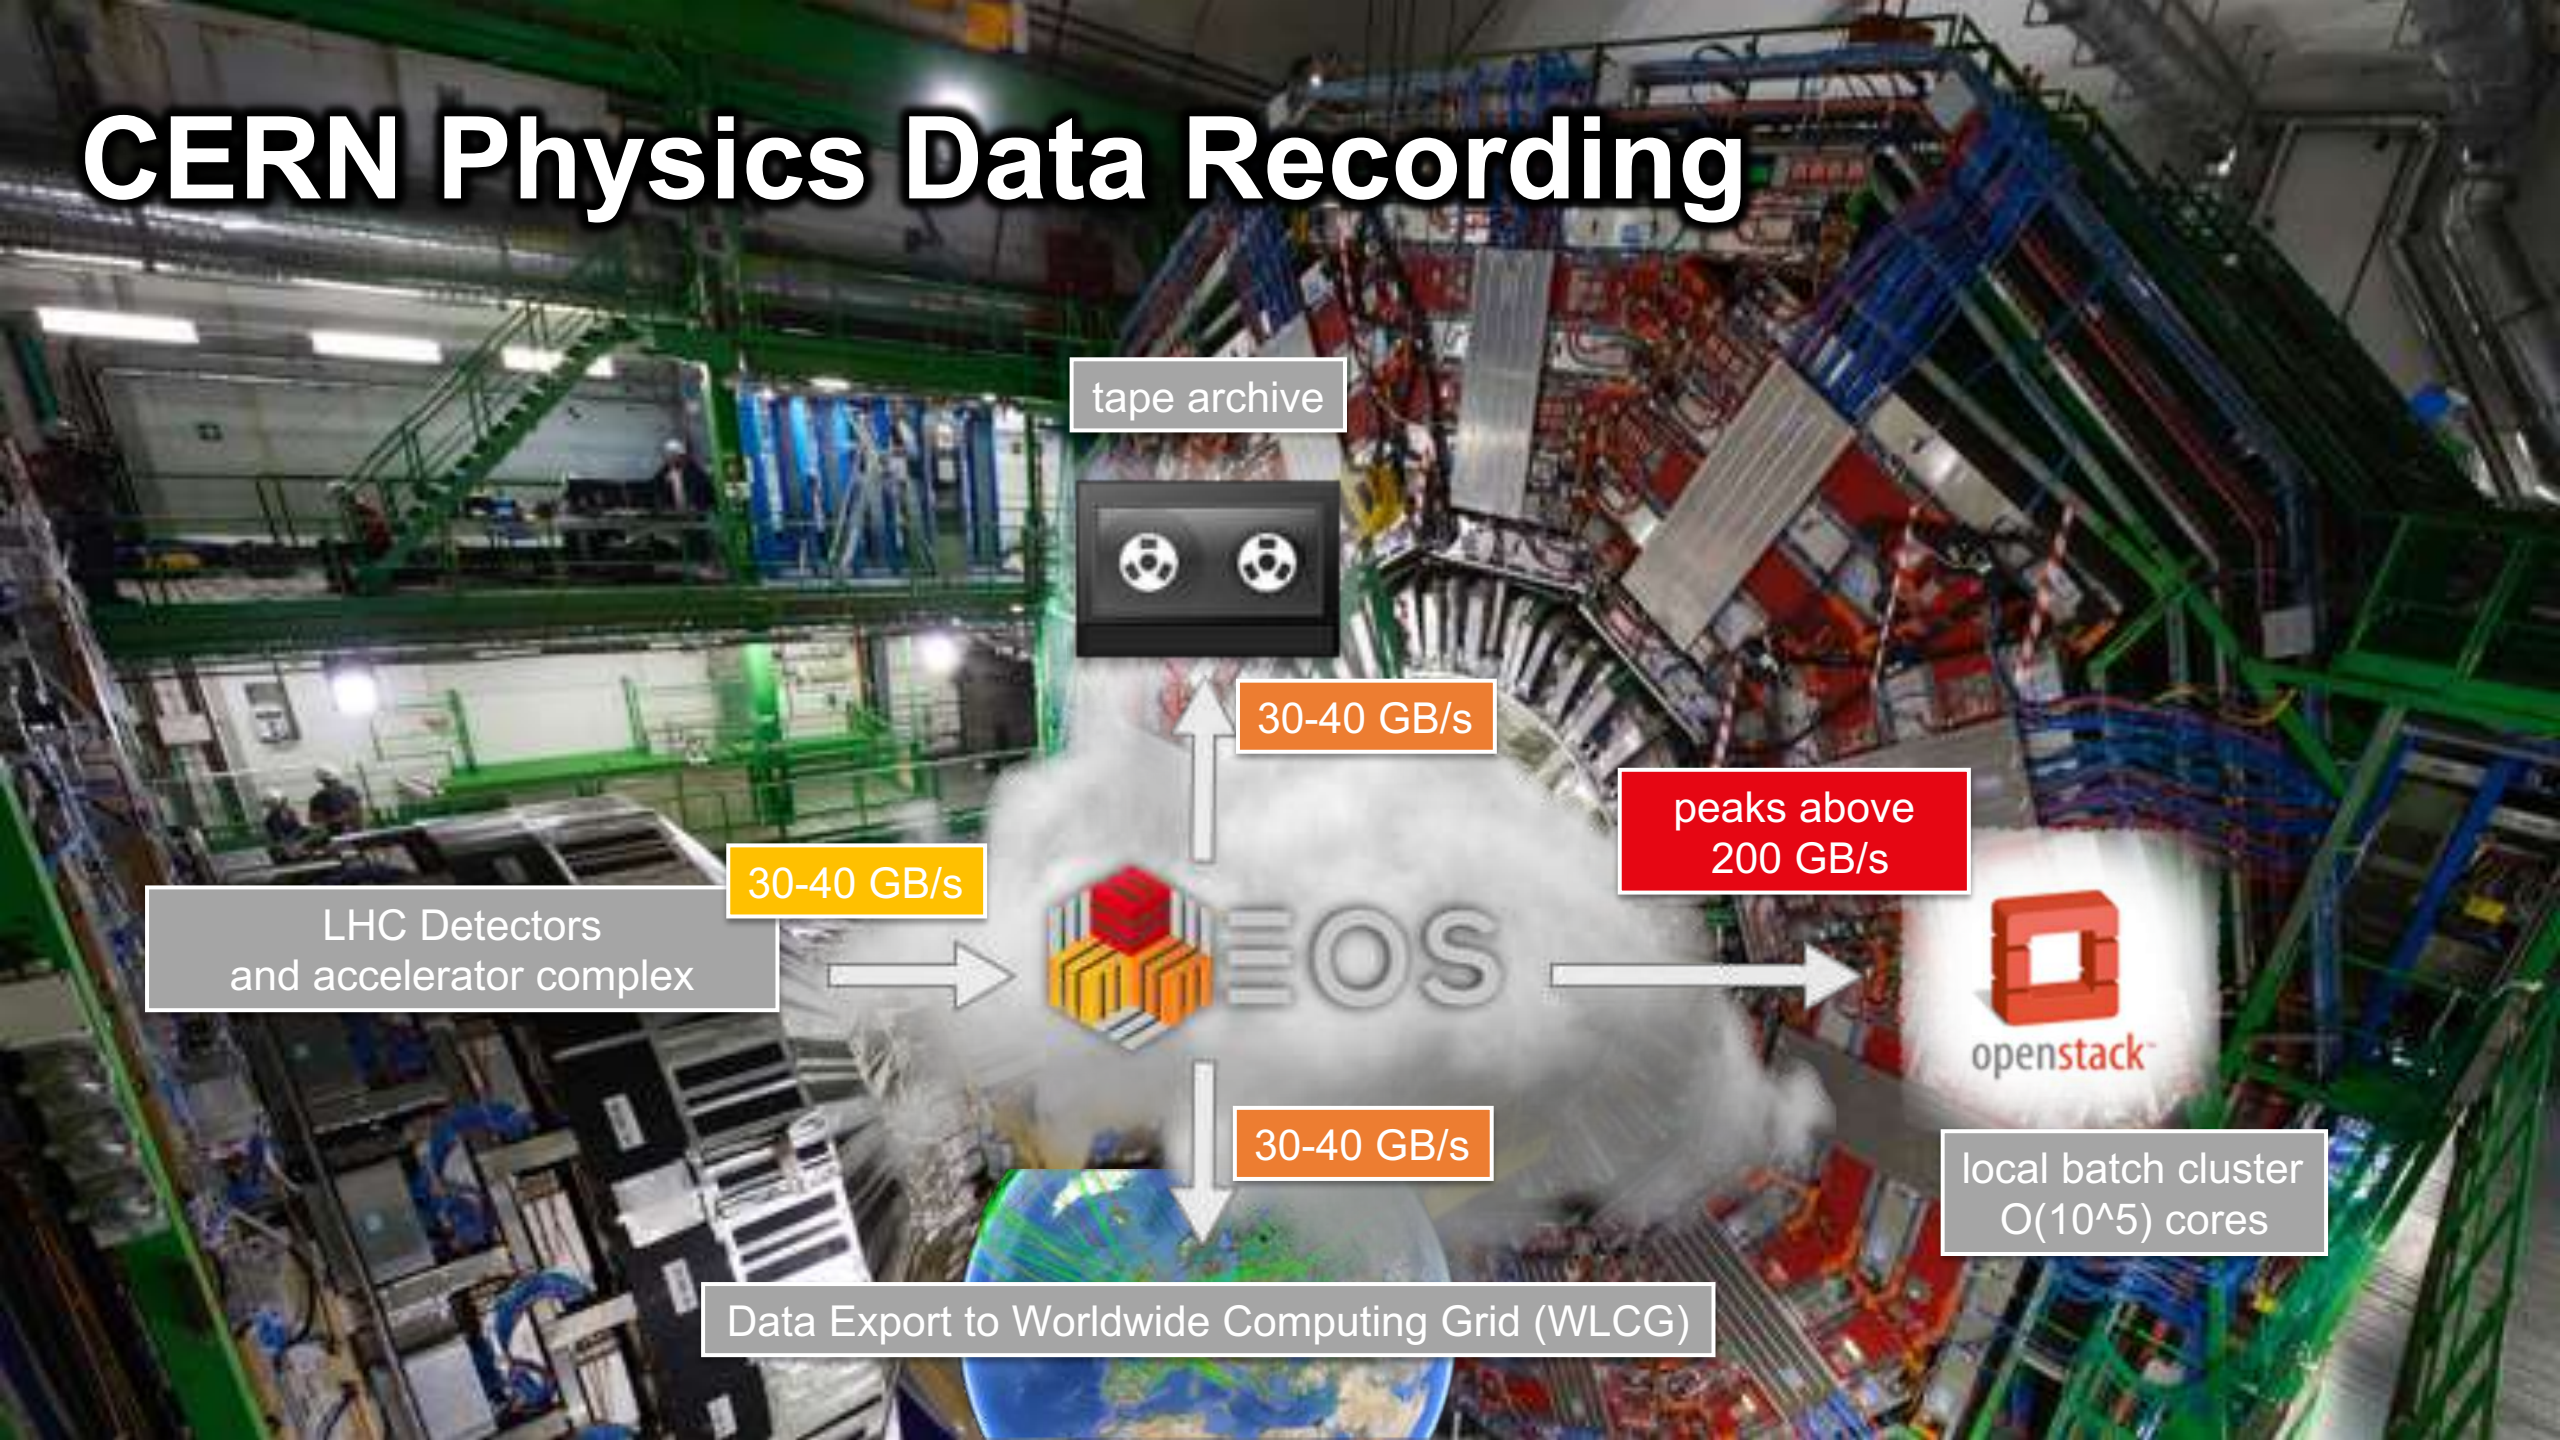

CERN Sync & Share platform

CERNBox
powered by EOS

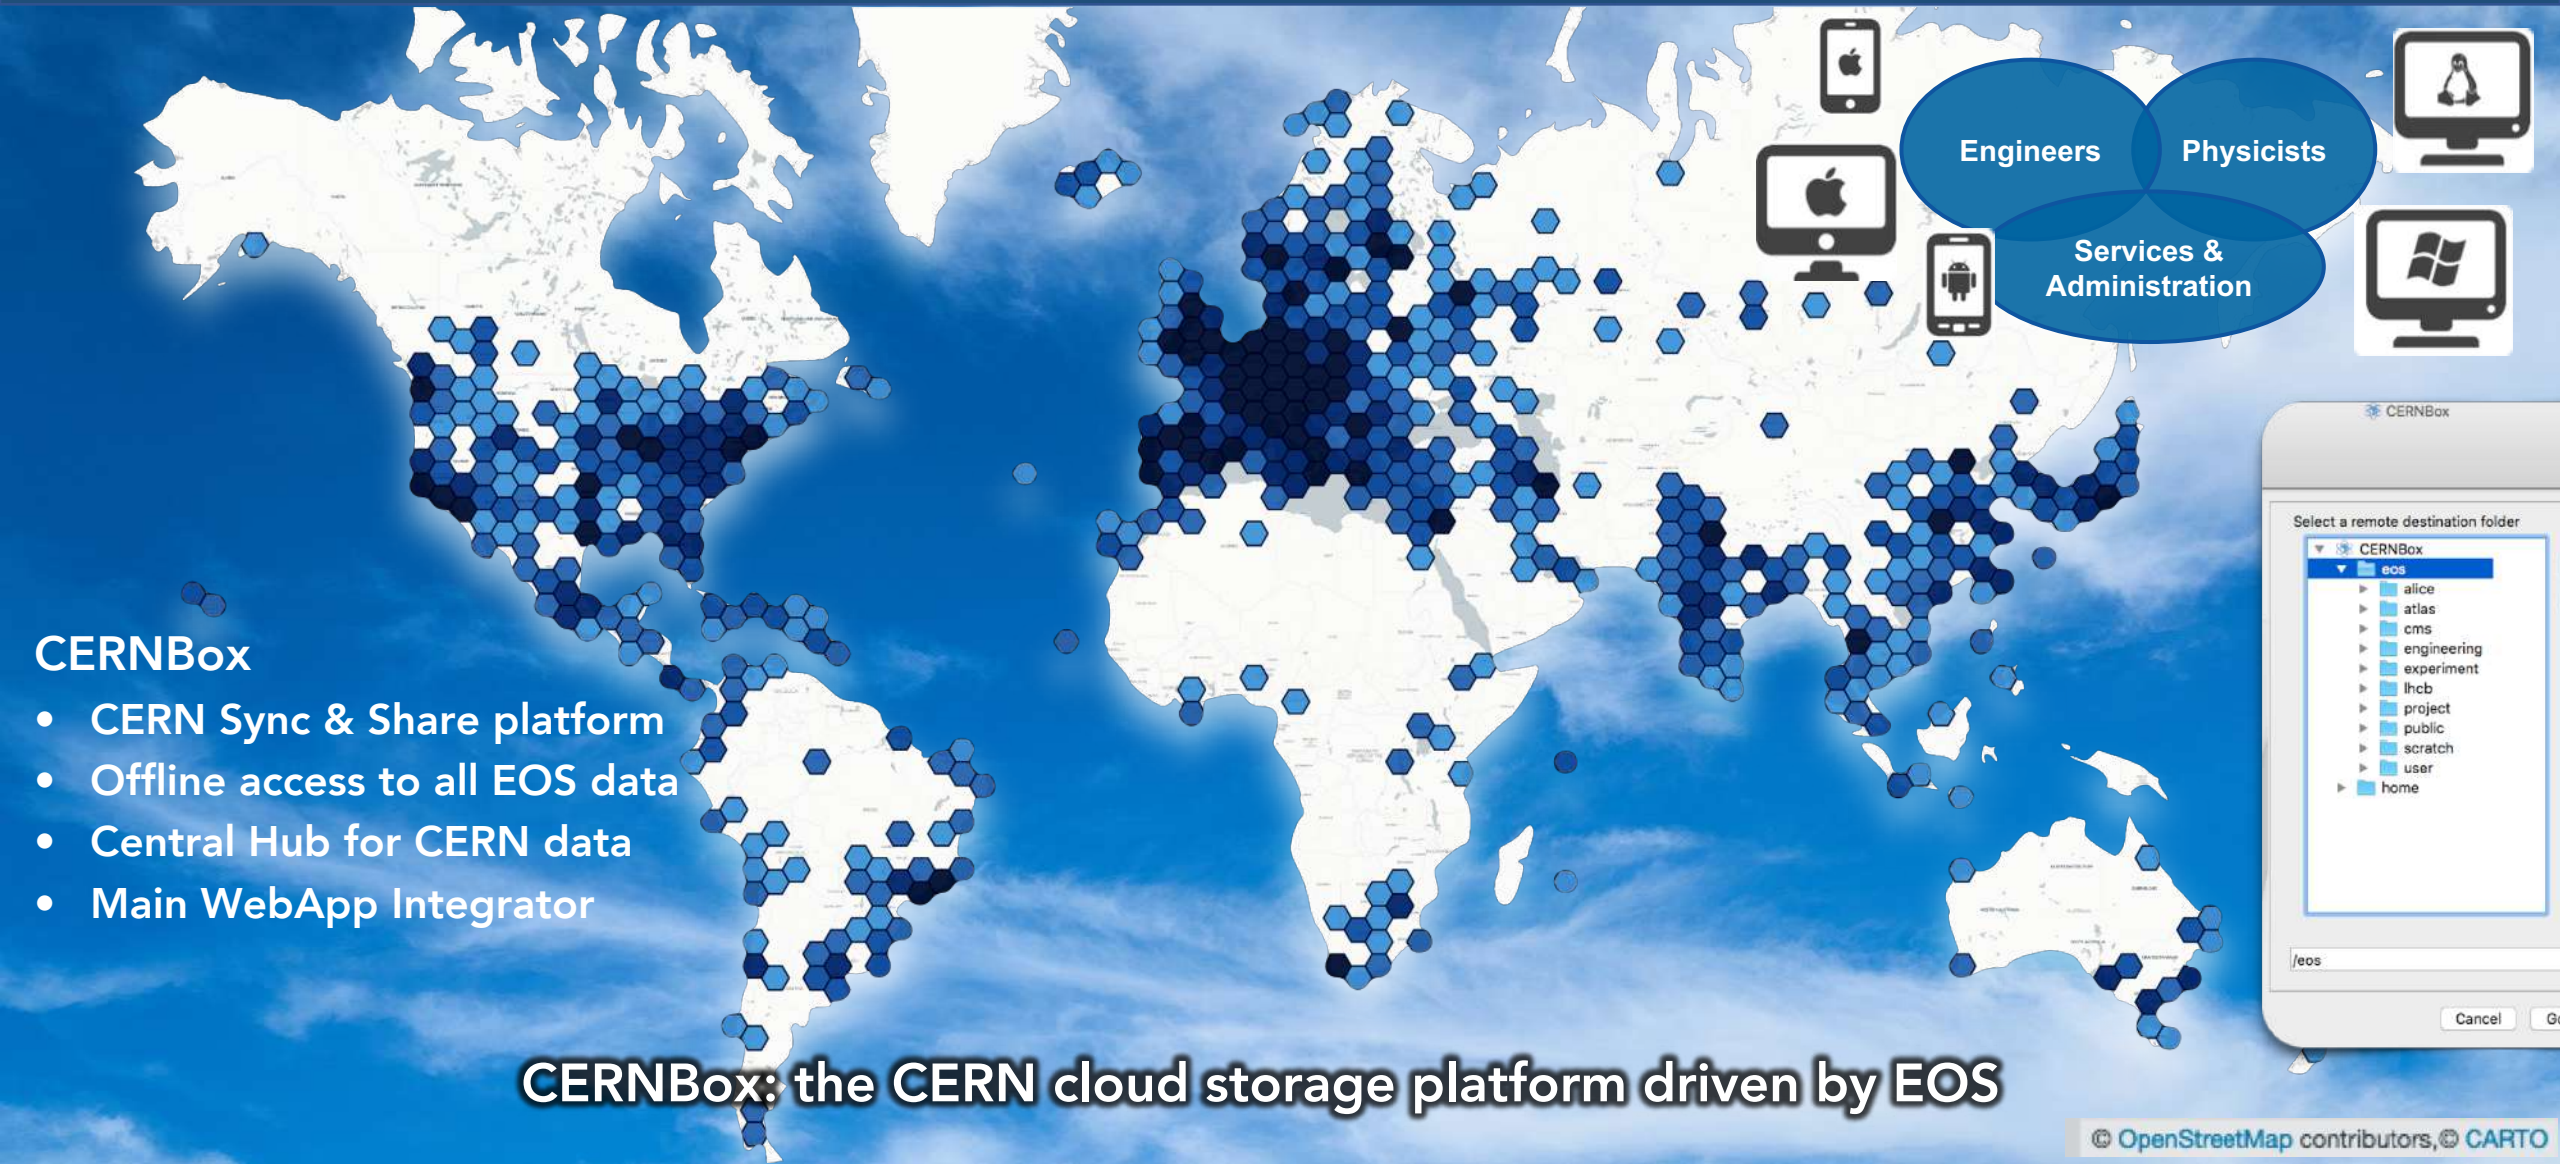

CERNBox

- CERN Sync & Share platform
- Offline access to all EOS data
- Central Hub for CERN data
- Main WebApp Integrator

CERNBox: the CERN cloud storage platform driven by EOS

Tape Storage Capabilities

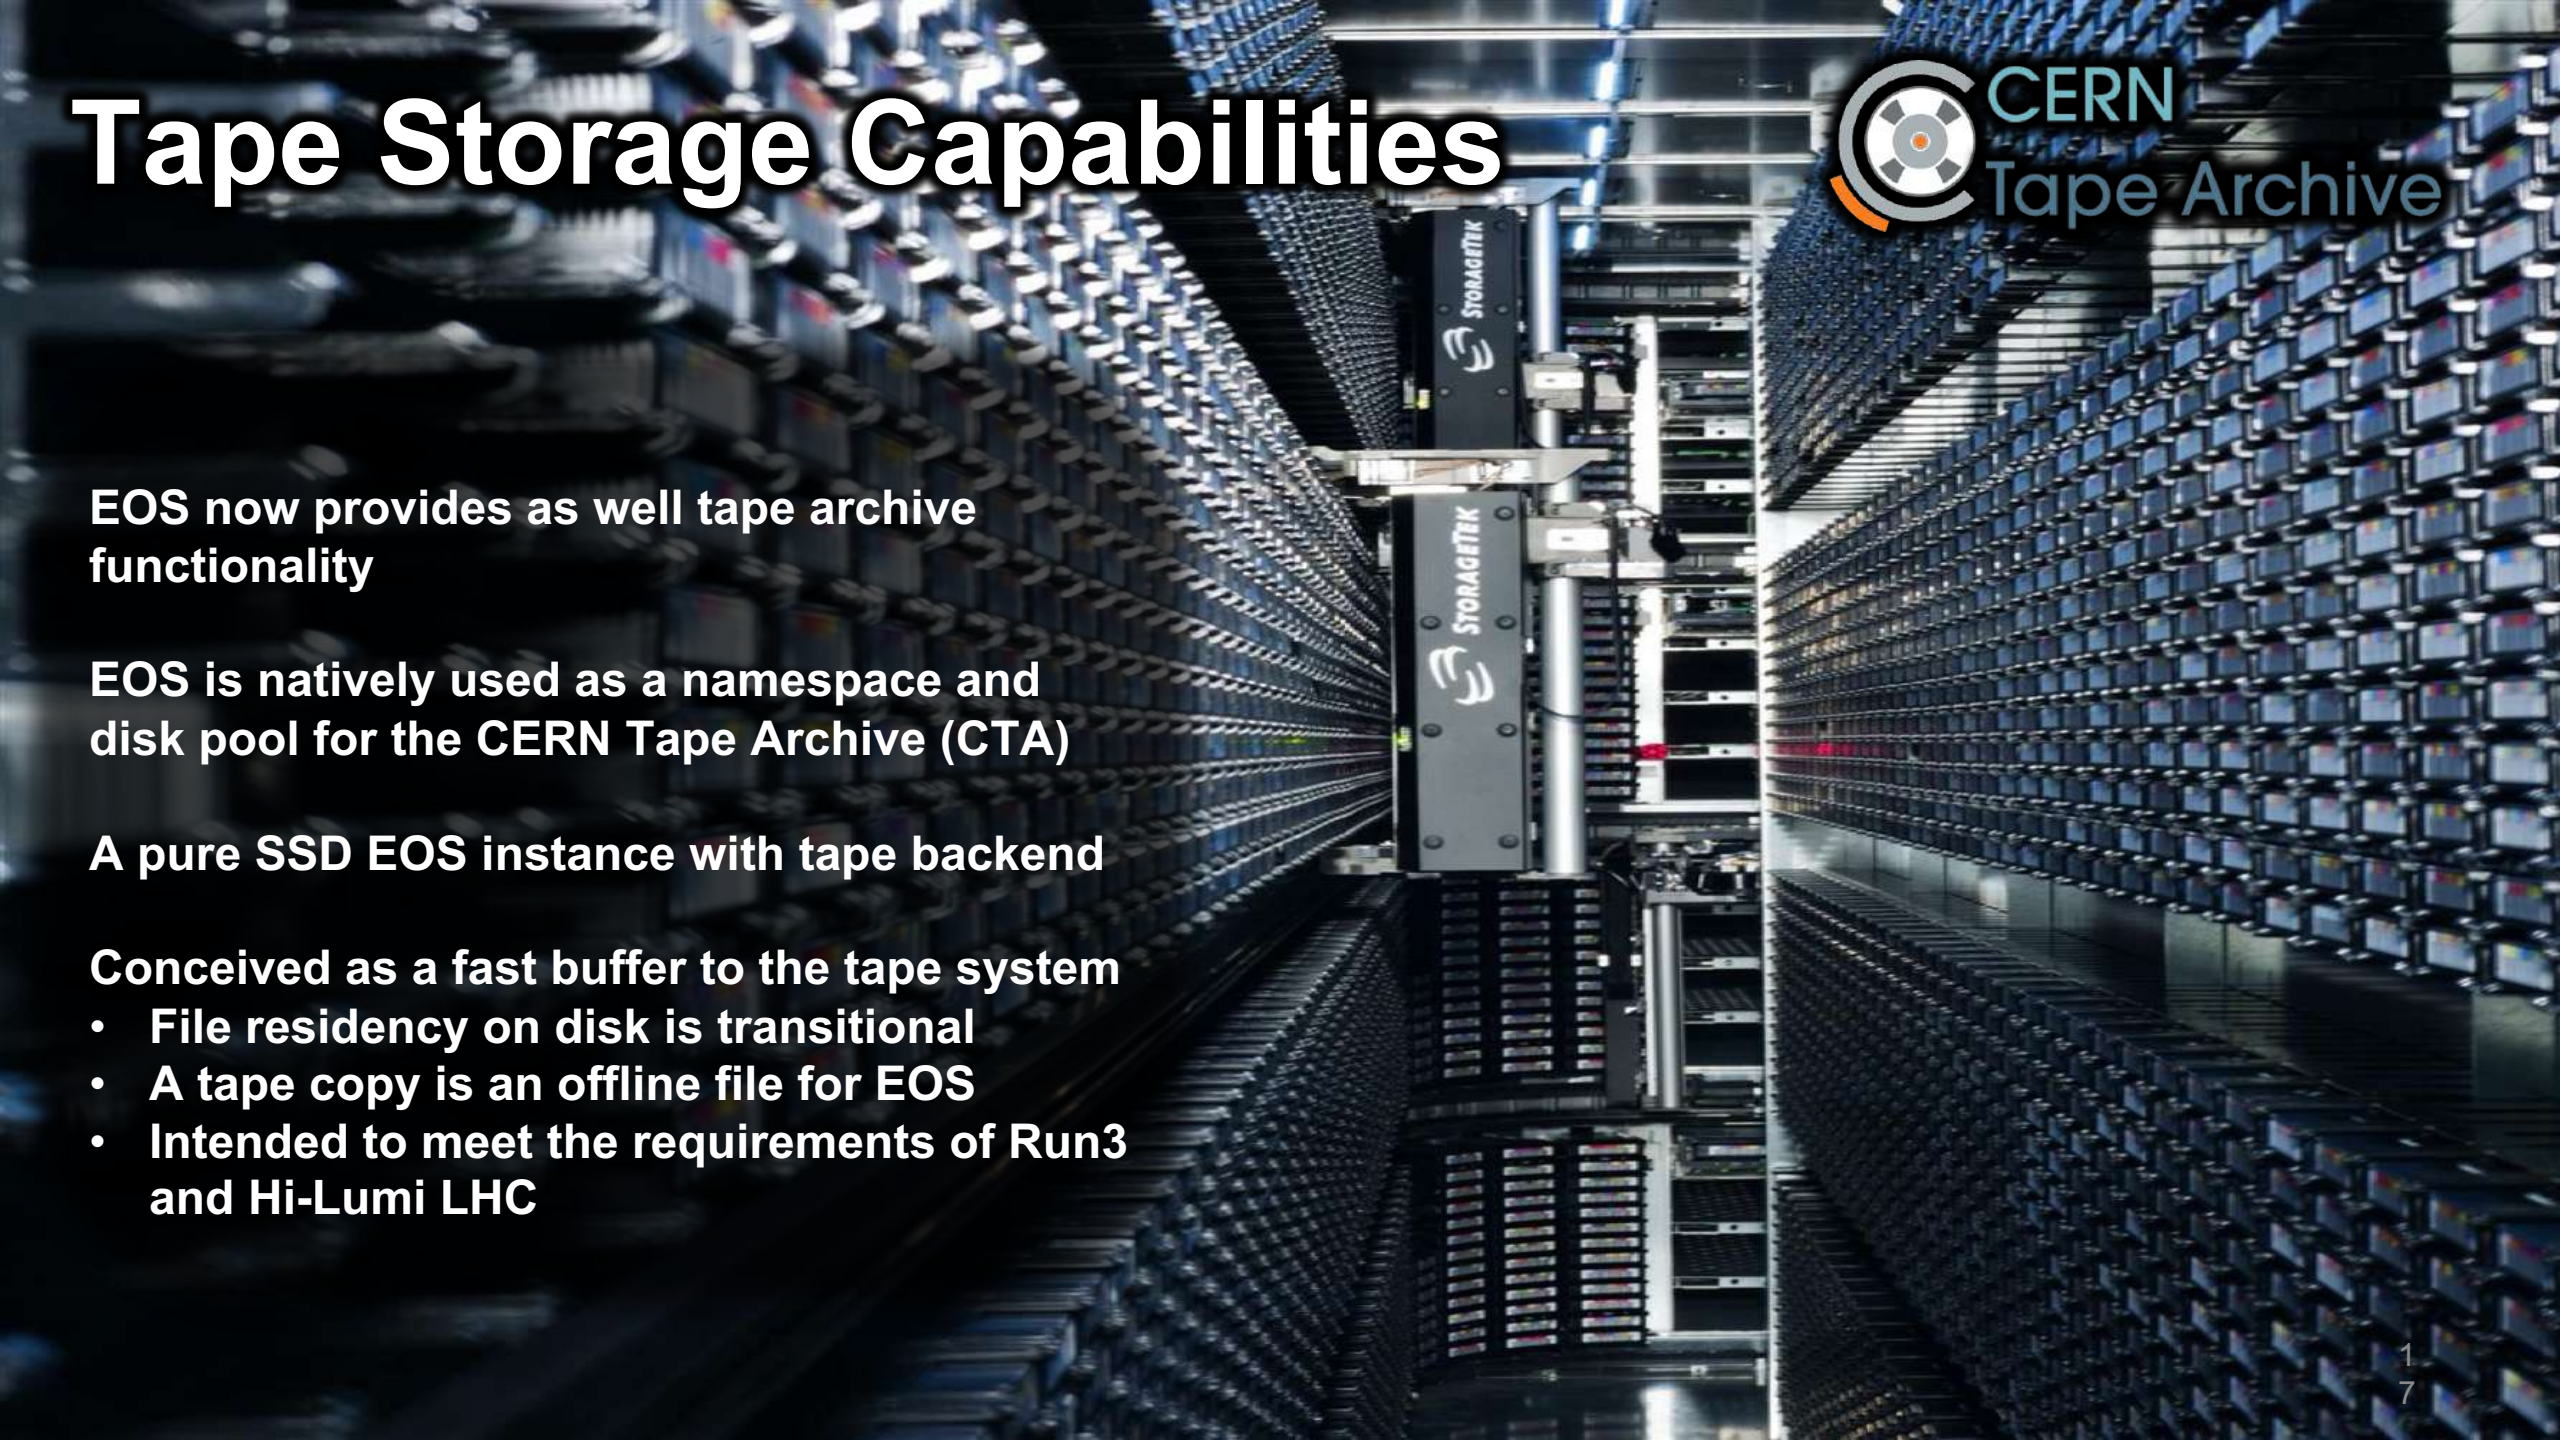

EOS now provides as well tape archive functionality

EOS is natively used as a namespace and disk pool for the CERN Tape Archive (CTA)

A pure SSD EOS instance with tape backend

Conceived as a fast buffer to the tape system

- File residency on disk is transitional
- A tape copy is an offline file for EOS
- Intended to meet the requirements of Run3 and Hi-Lumi LHC

EOS Architecture

High-available and low latency namespace
- namespace persisted on a distributed key-value store
- working entries cached in-memory

High available and reliable file storage,
based on (cheap) JBODs:
- File replication across independent nodes and disks
- Erasure coding to optimize costs and data durability

MGM : meta data server
MQ : message queue
NS : persistent namespace
FST : file storage server

EOS online access for Windows: SAMBA

SAMBA access to EOS and CERNBox
- load-balanced *ctdb-driven* cluster setup
- EOS is fuse-mounted on the gateways
- Additional shared mount to share the state

Windows domain (AD) joined in dedicated keytab mode
- Authc performed by windbind
- Authz performed by EOS

File locking supported across gateways
- A must to support Office concurrent usage notifications

Complex stack of software
- very latency sensitive

Investigations for native solution (EOS-wnc)
- synergies with Comtrade – CERN openlab collaboration

EOS in Industry: distributed set-up

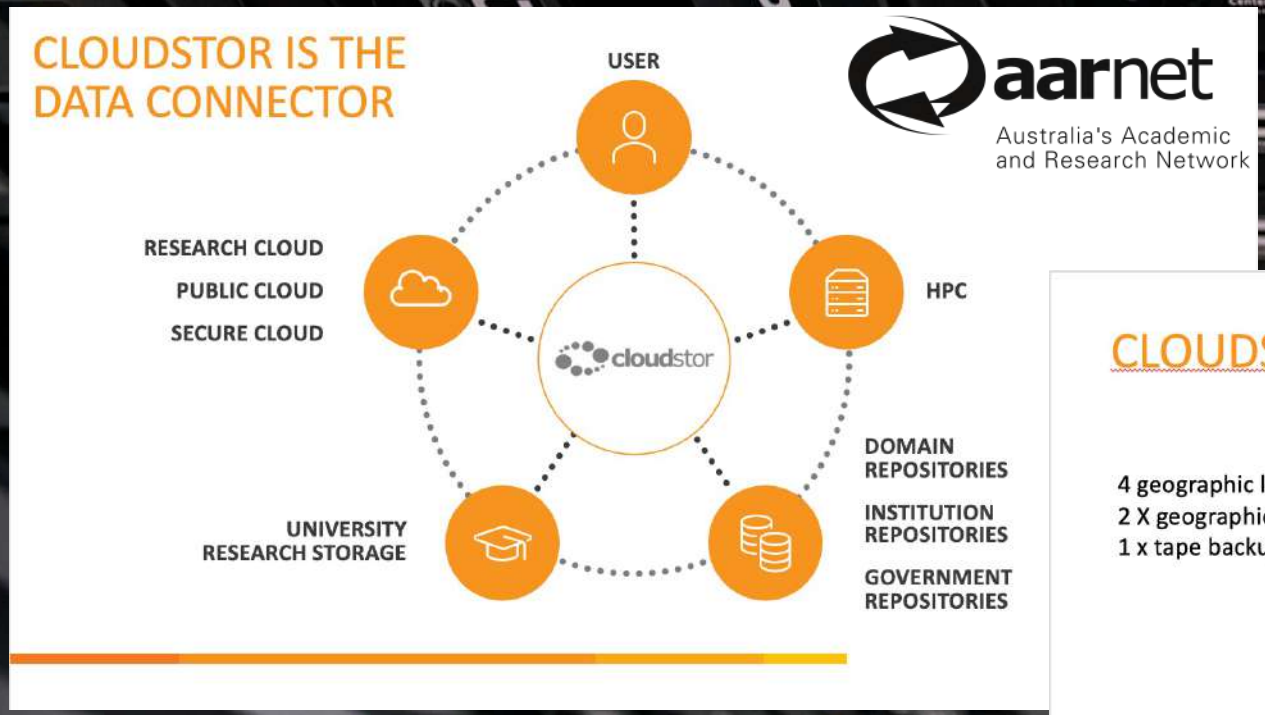

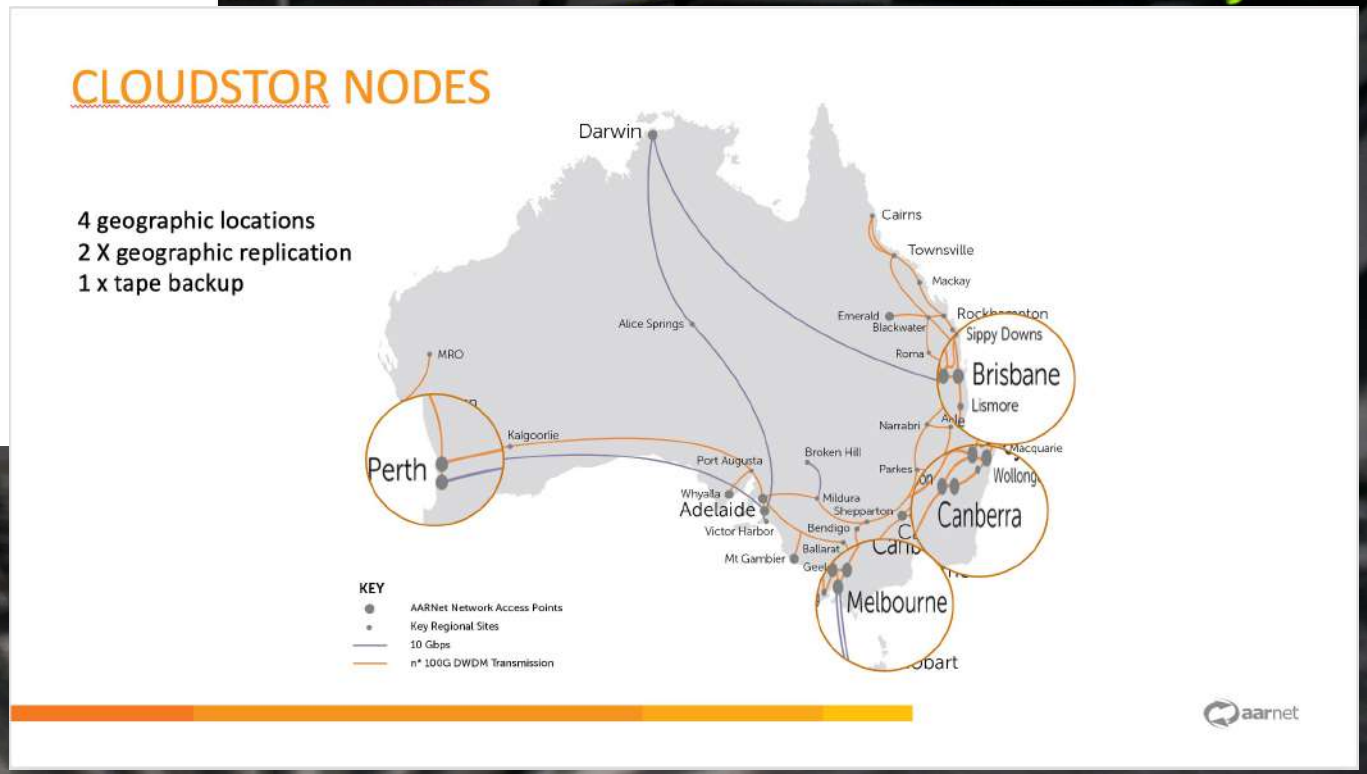

EOS in Industry: online satellite data analysis

JRC Earth Observation Data and Processing Platform (JEODPP)

JEODPP
The JRC Big Data Platform

Versatile platform bringing the users to the data

Joint Research Centre
JRC

Examples of exploratory analyses and interactive visualization

JEODPP
The JRC Big Data Platform

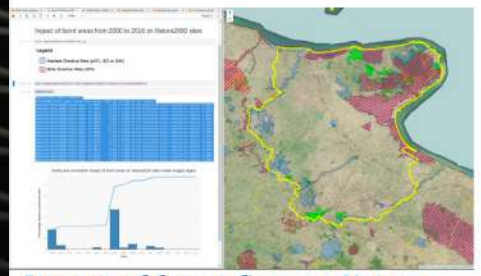

Impact of forest fires on Natura 2000 sites

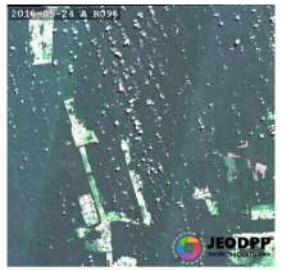

Deforestation (time lapse)

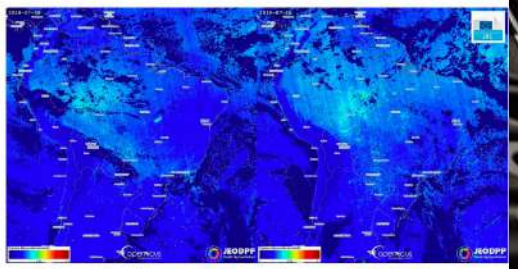

CO emissions (time lapse 08/18 vs 08/19)

Ship traffic (heat maps)

Deforestation (split map 2000 vs 2017)

Interactive combination of information layers

Joint Research Centre

European Commission


EOS Workshop 2022
7-10 March
Work in Progress


CERN
Accélérateur de science