Update on EOS productisation

Comtrade 360’s EOS Windows Native Client

Luca Mascetti
CERN IT Storage

luca.mascetti@cern.ch

Gregor Molan
Comtrade 360’s AI Lab

gregor.molan@comtrade.com

COMTRADE 360: CERN openlab Associate member

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

https://openlab.cern/project/eos-productisation

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/REST, WebDAV, CIFS)

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

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 Project History

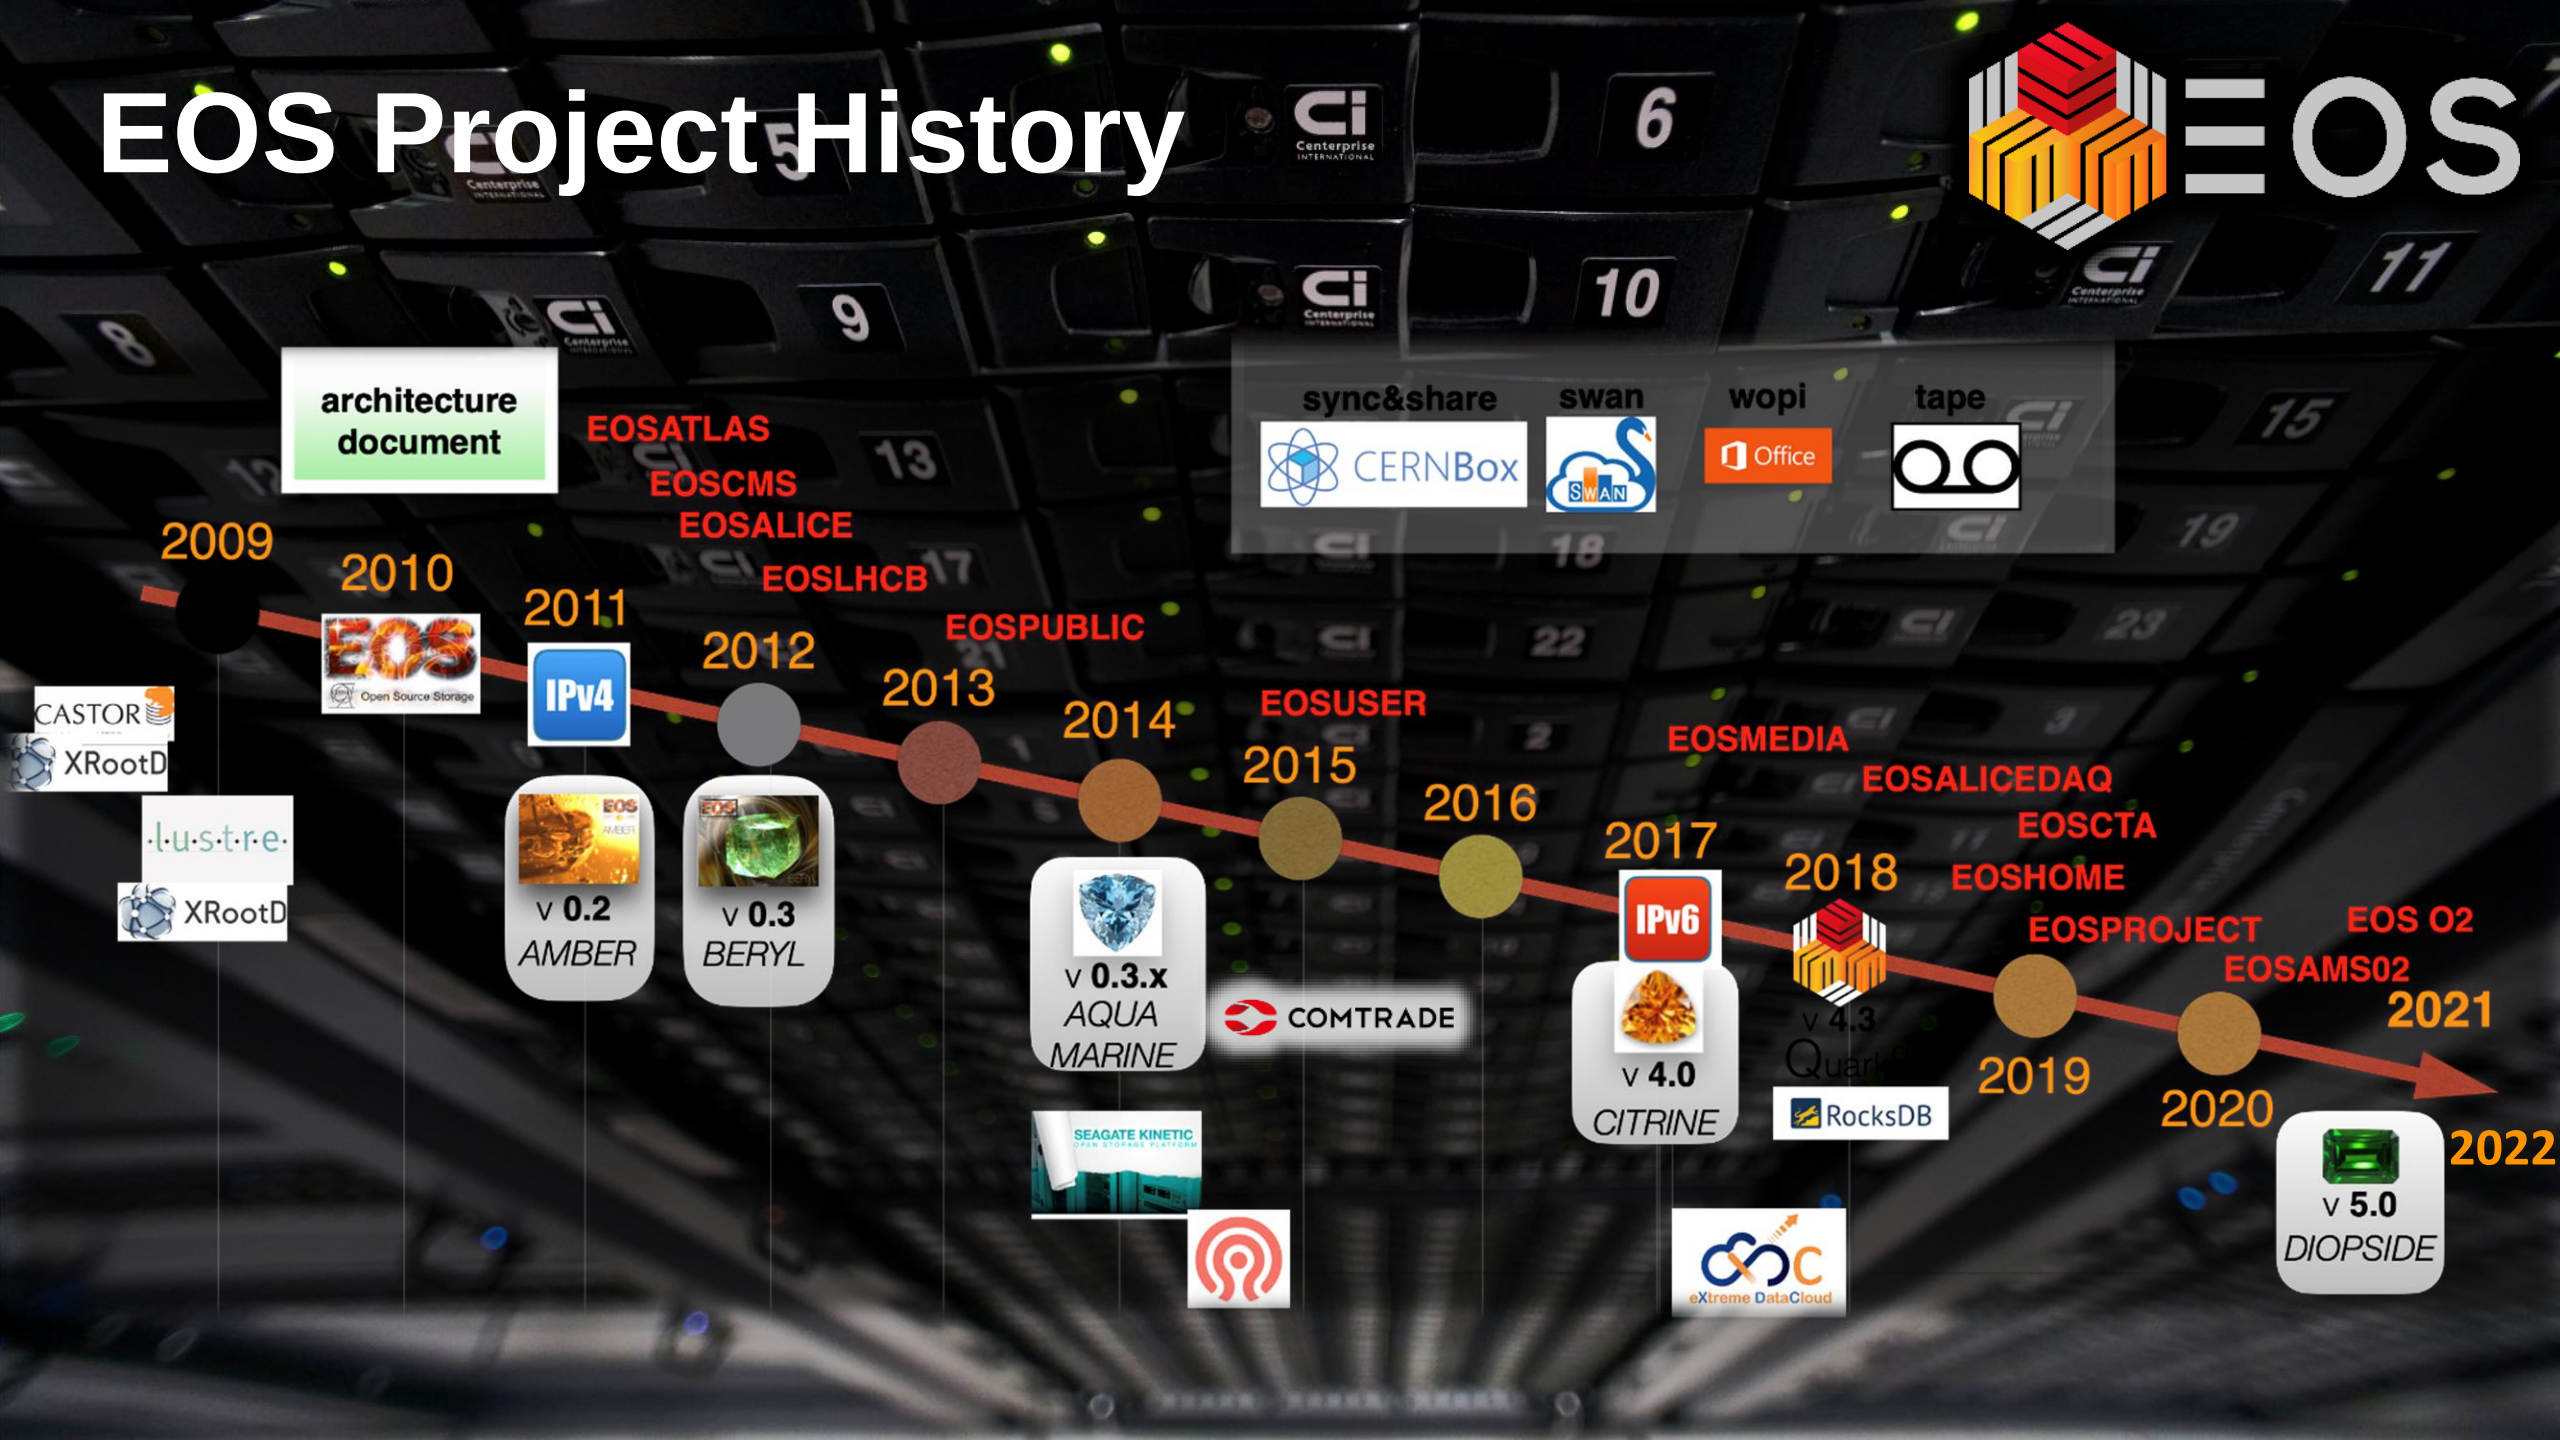

EOS @ CERN

Total Space
600 PB

Files Stored
~7 Bil

 # Storage Nodes
~1600

 # Disks
~80000

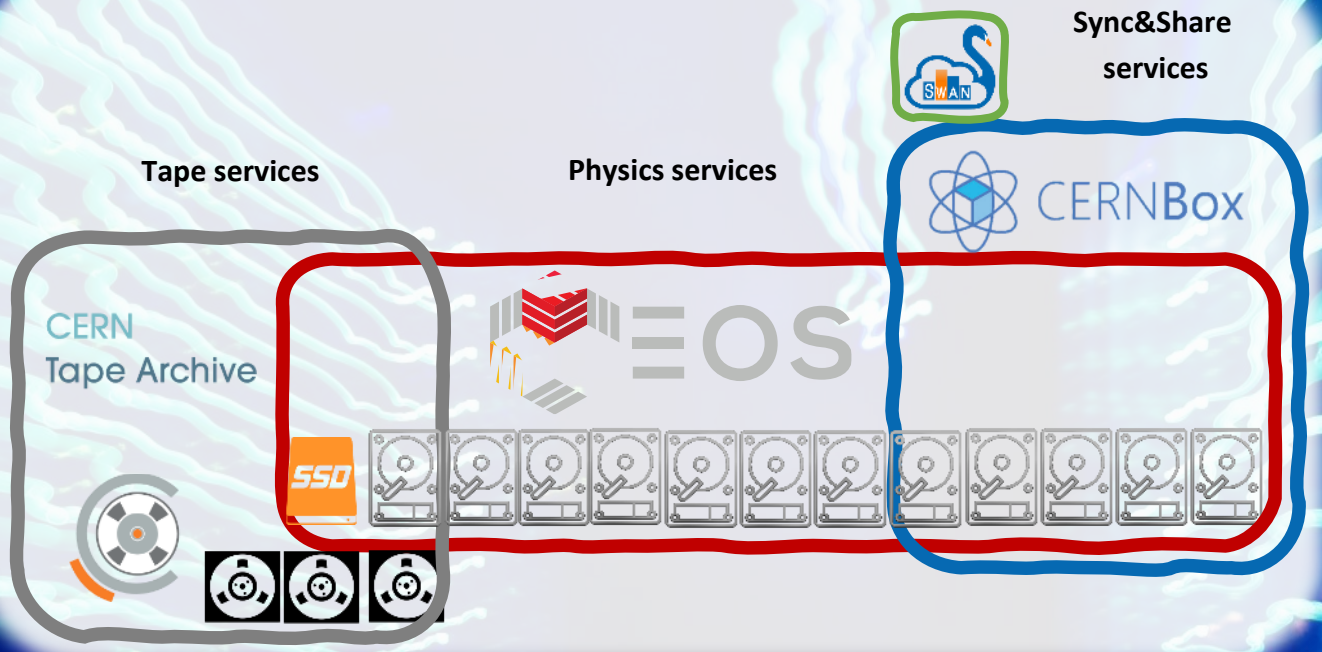

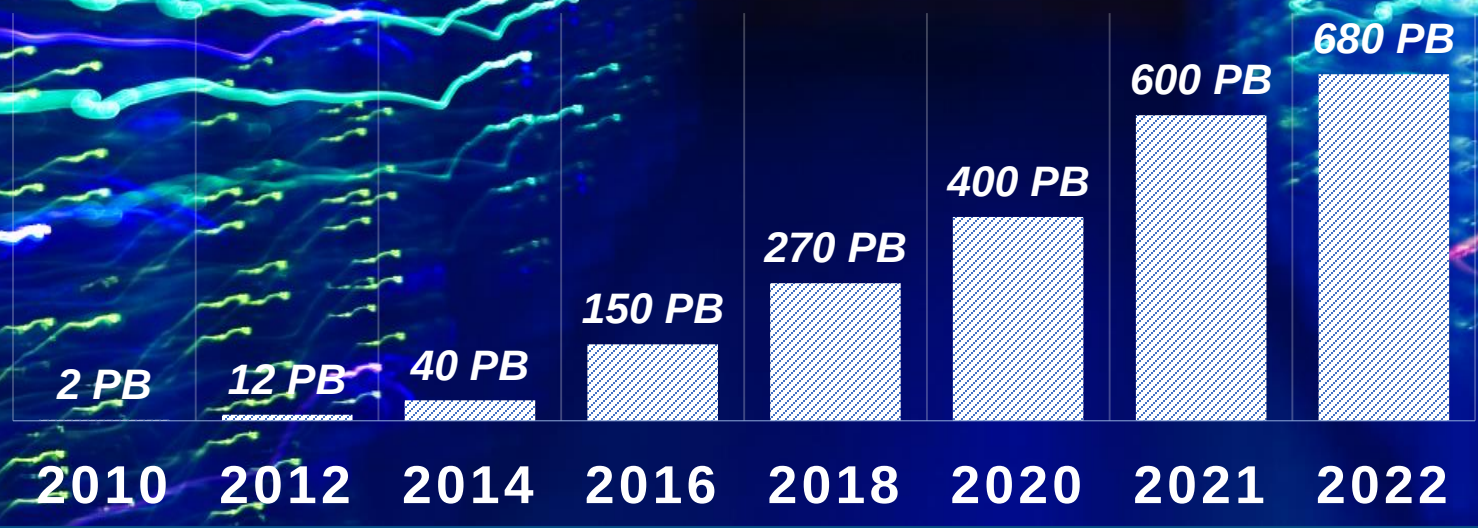

EOS Service for Physics

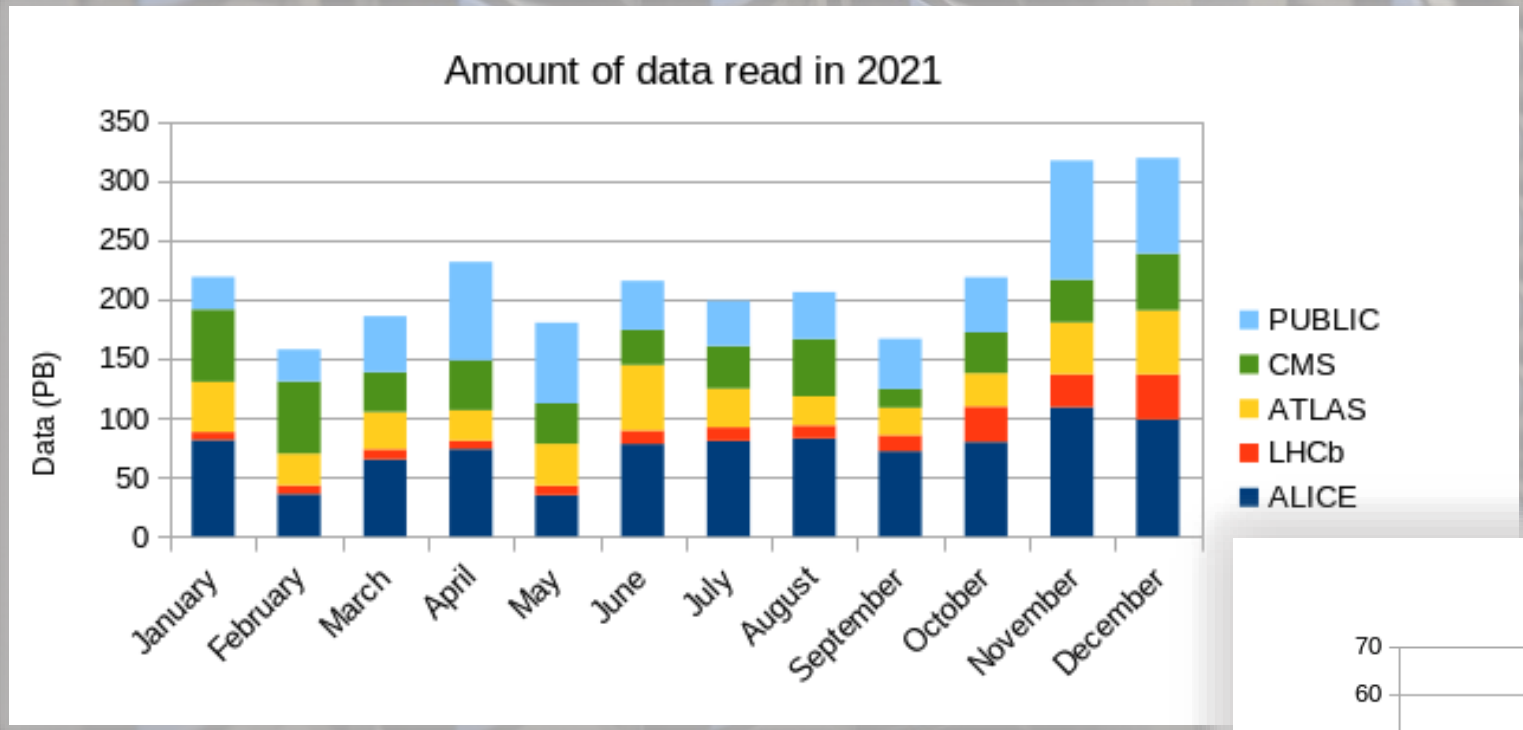

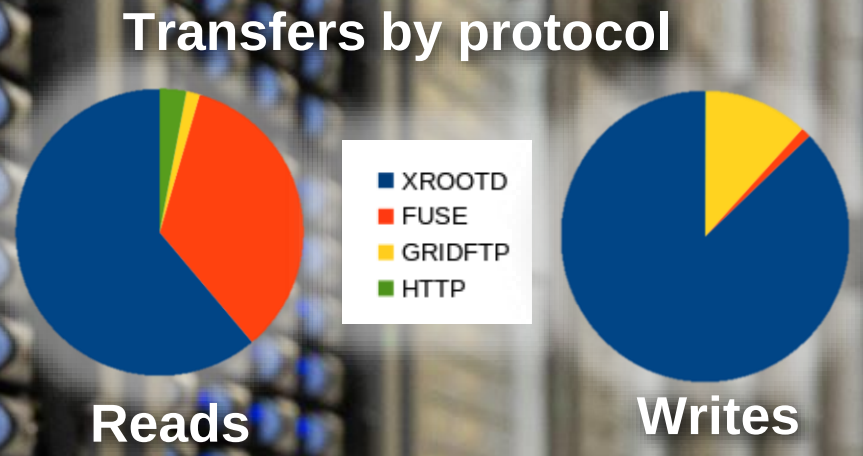

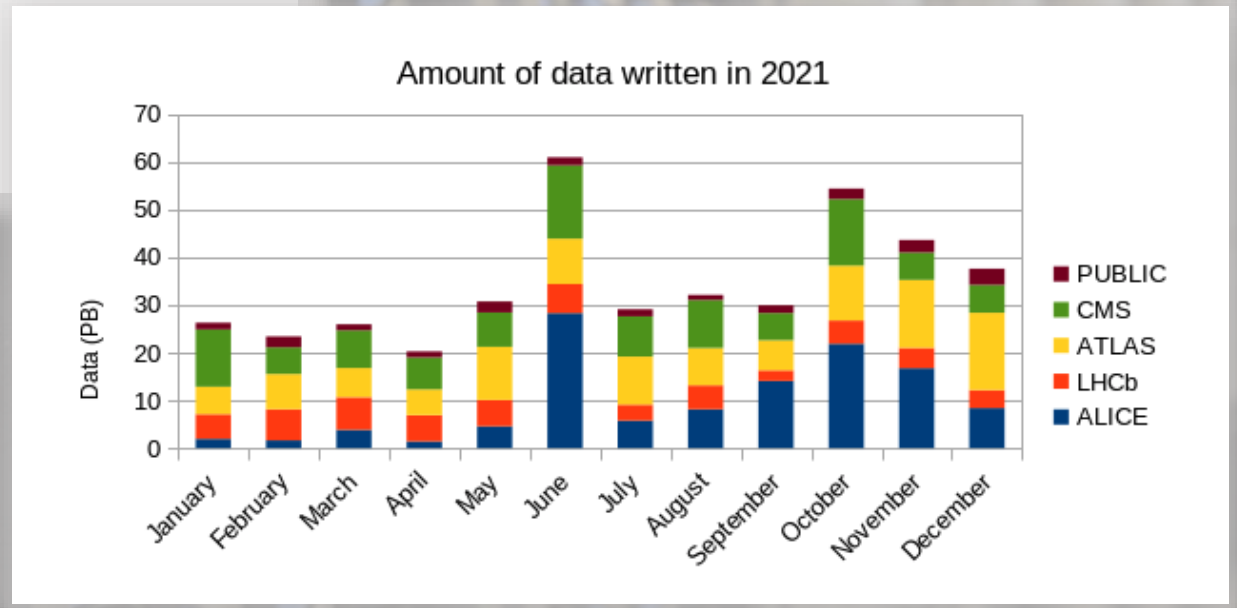

~ 2.6 ExaBytes of data delivered in 2021
~0.5 ExaBytes stored in 2021

CERN Physics Data Recording

tape archive

40-50 GB/s

150-200 GB/s

LHC Detectors
and accelerator complex

EOS

peaks above
350 GB/s

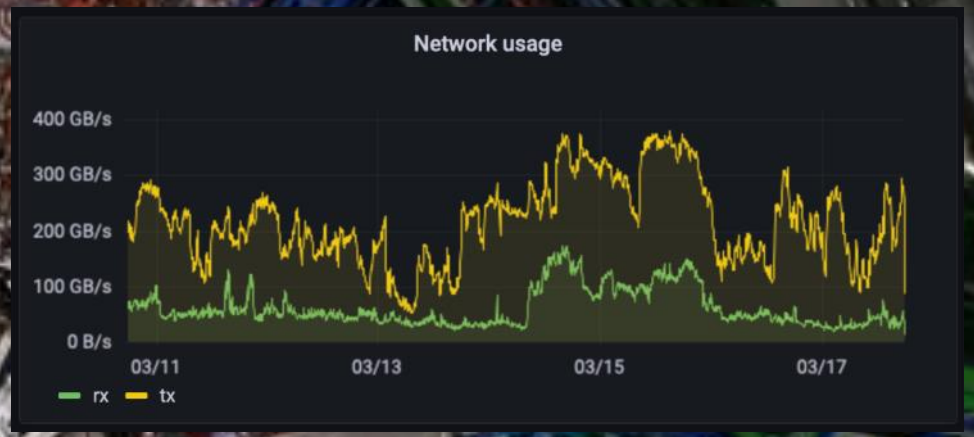

openstack

40-50 GB/s

local batch cluster
O(10^5) cores

Data Export to Worldwide Computing Grid (WLCG)

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

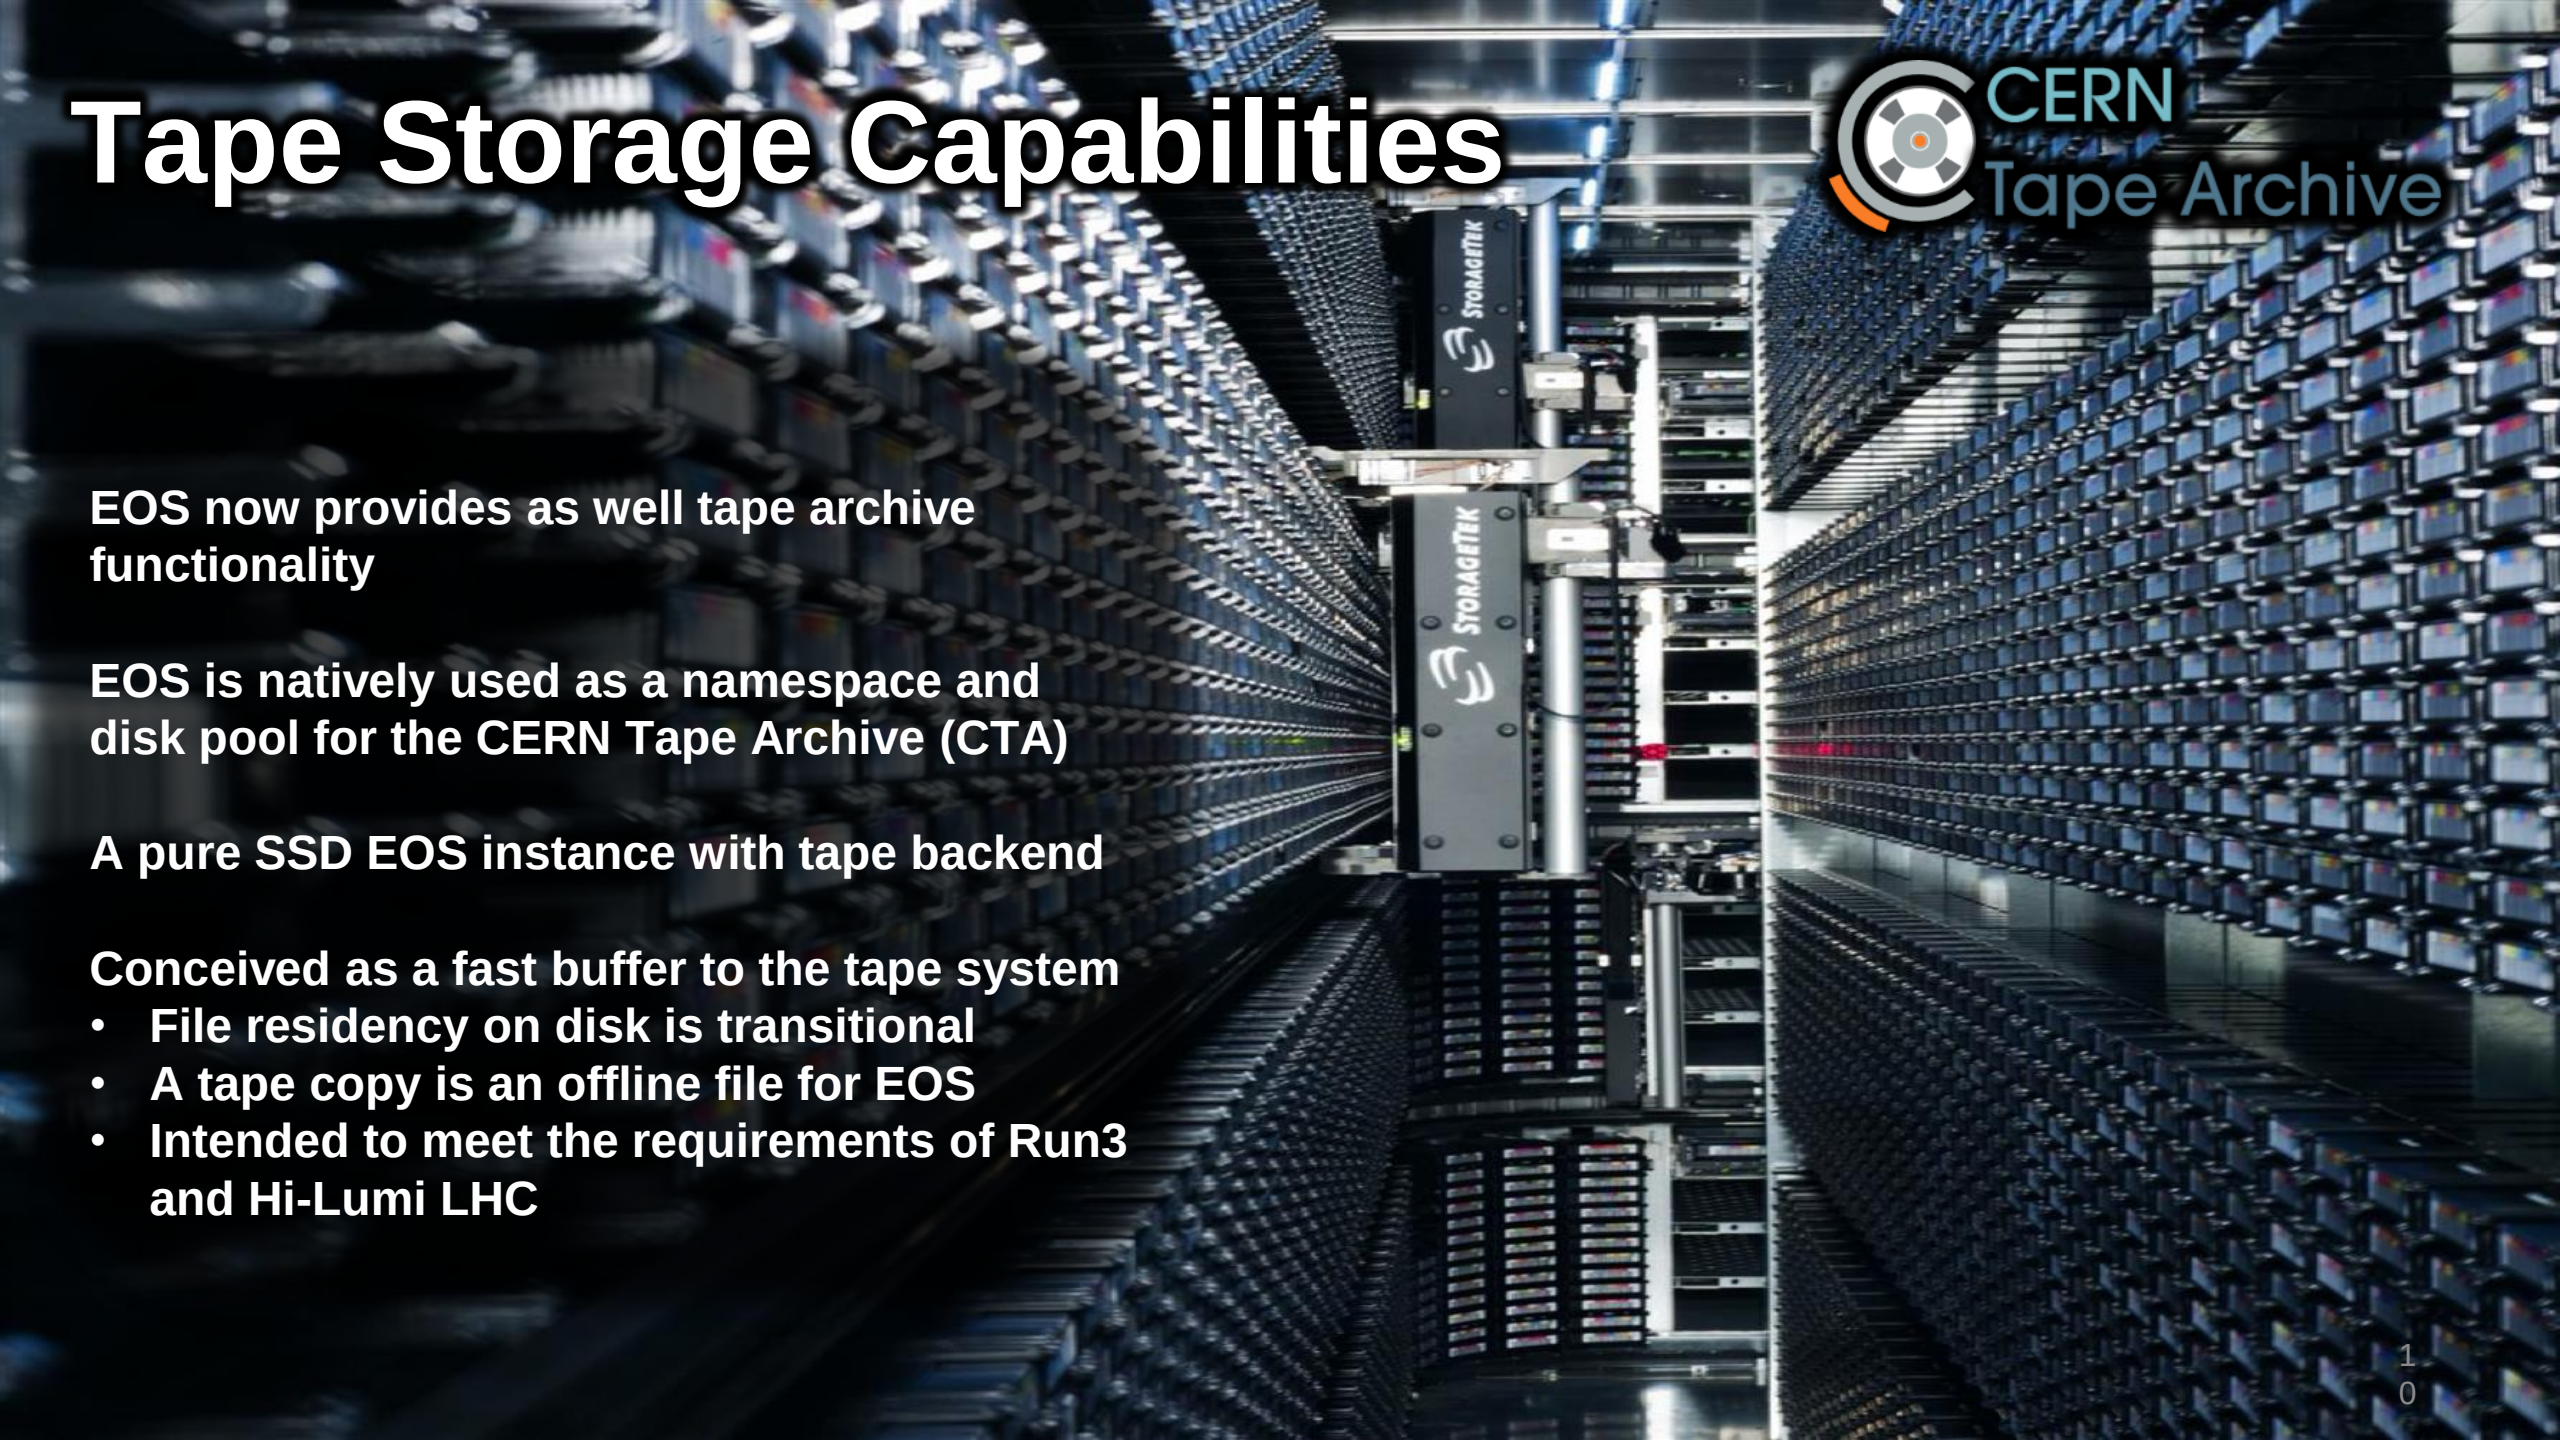

EOS now provides as well tape archive functionality

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

The timeline

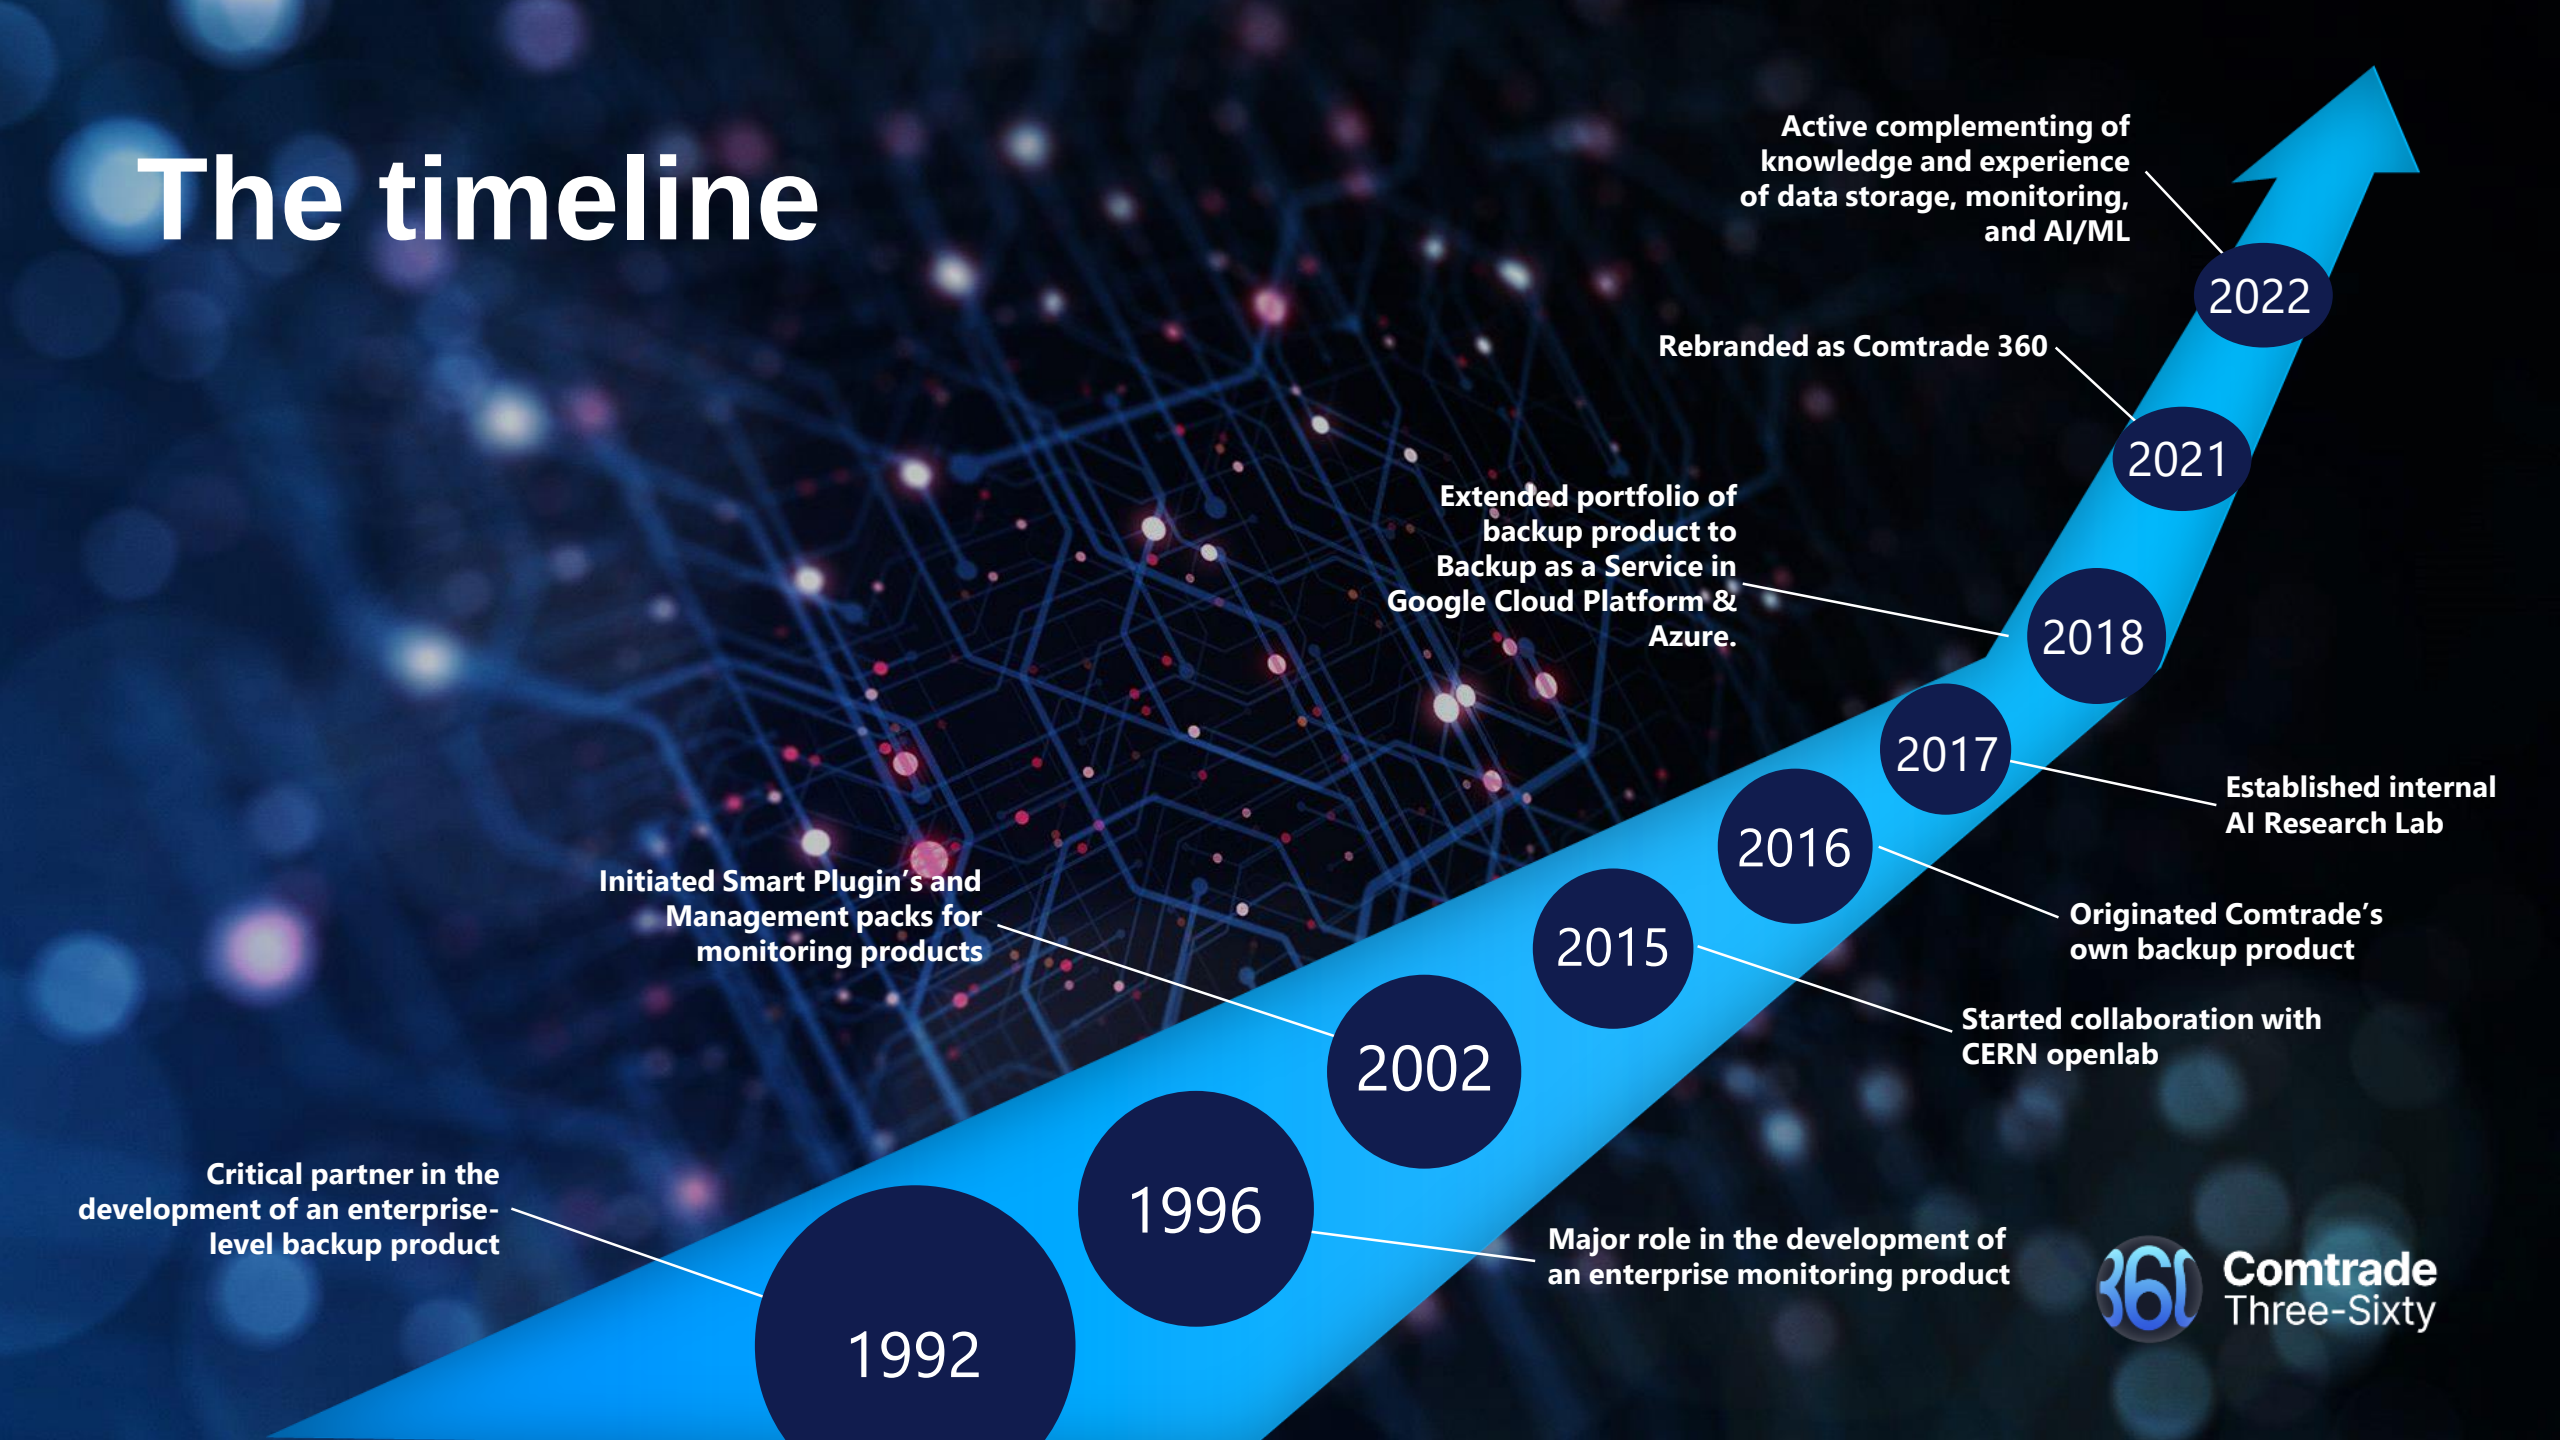

Comtrade Three-Sixty

Comtrade 360

Result of proven track record:
One of the most reliable global technology solutions providers

30
years of experience

3
R&D campuses

300+
active engineers

EOS-wnc: CERN EOS on Windows

- Enable Windows for broader use
- Access to CERN technology
  - Open-source large-scale tech. to store data from experiments in CERN
  - Native usage on Windows of a shared filesystem developed for Linux
- The bridge from Windows to physics experiments storage
- Provides high availability, low latency, high reliability, high speed, access to storage with unlimited extendibility

Comtrade Three-Sixty

The bridge: Windows through Samba

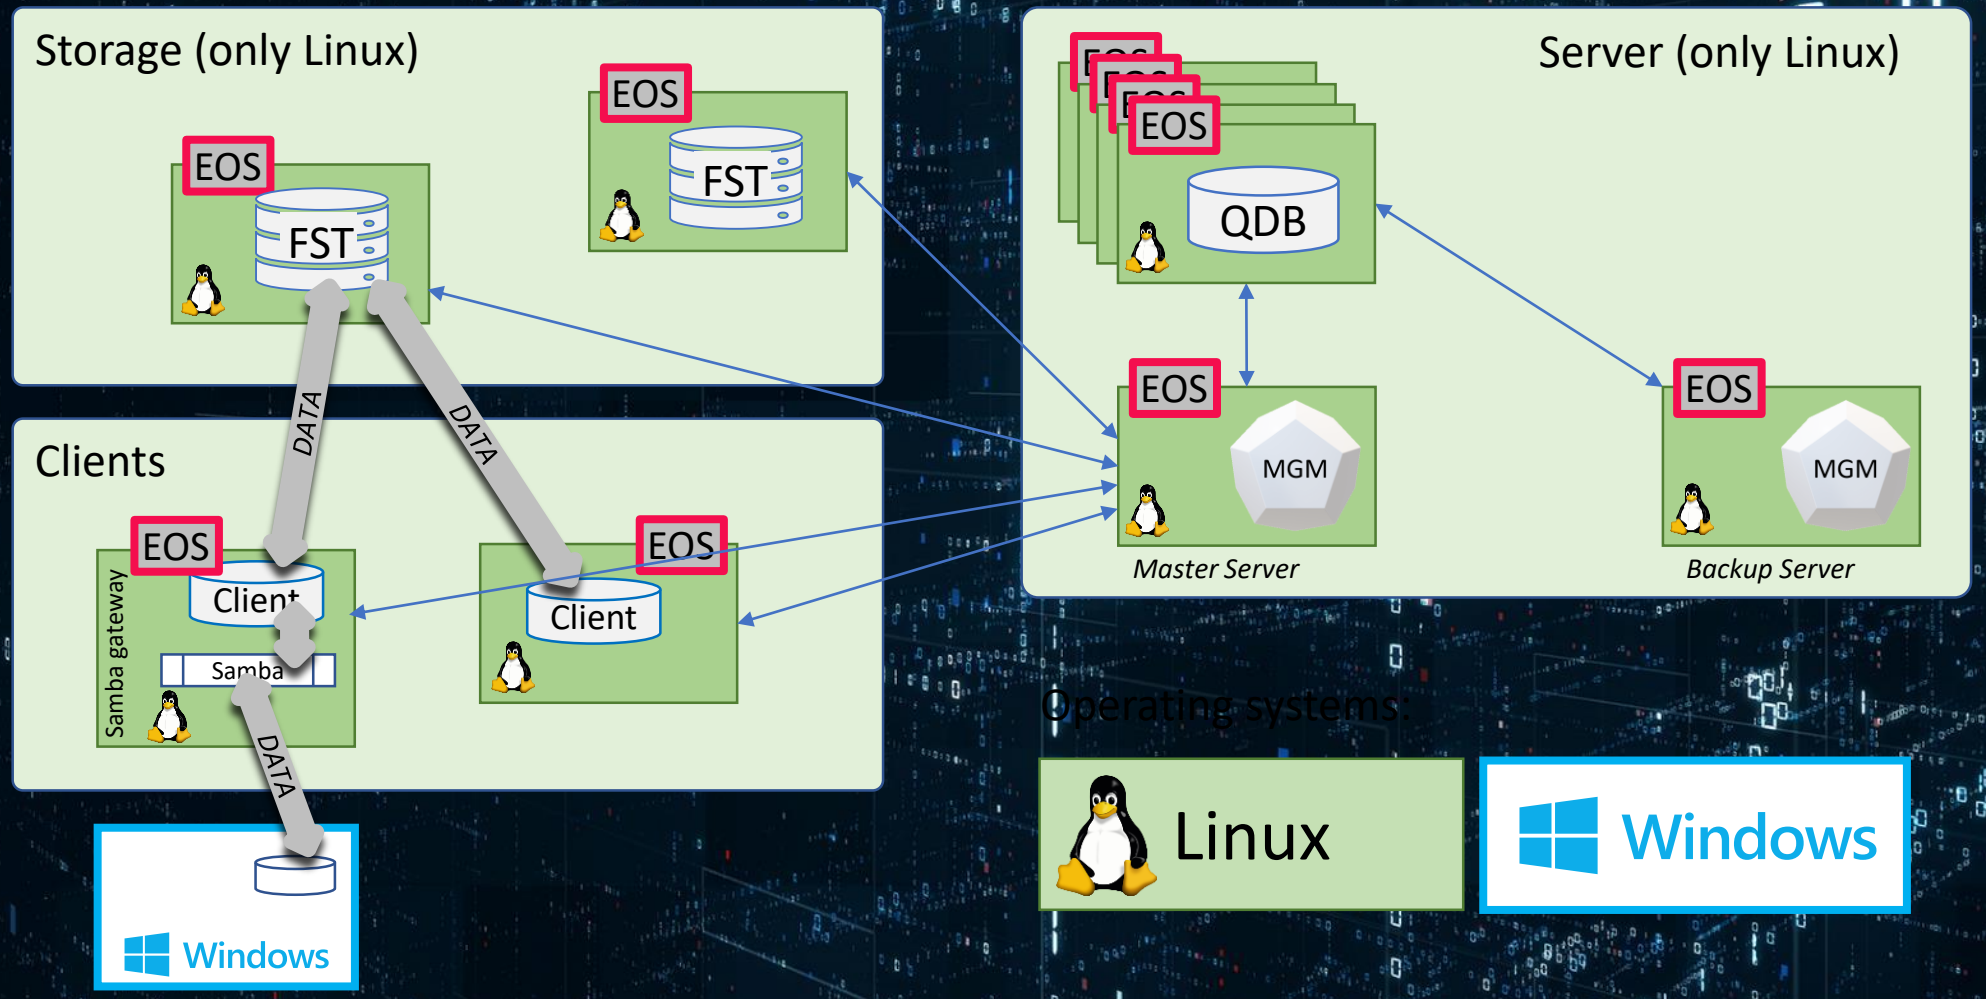

Operating systems:

Linux

Windows

The bridge: Windows drive

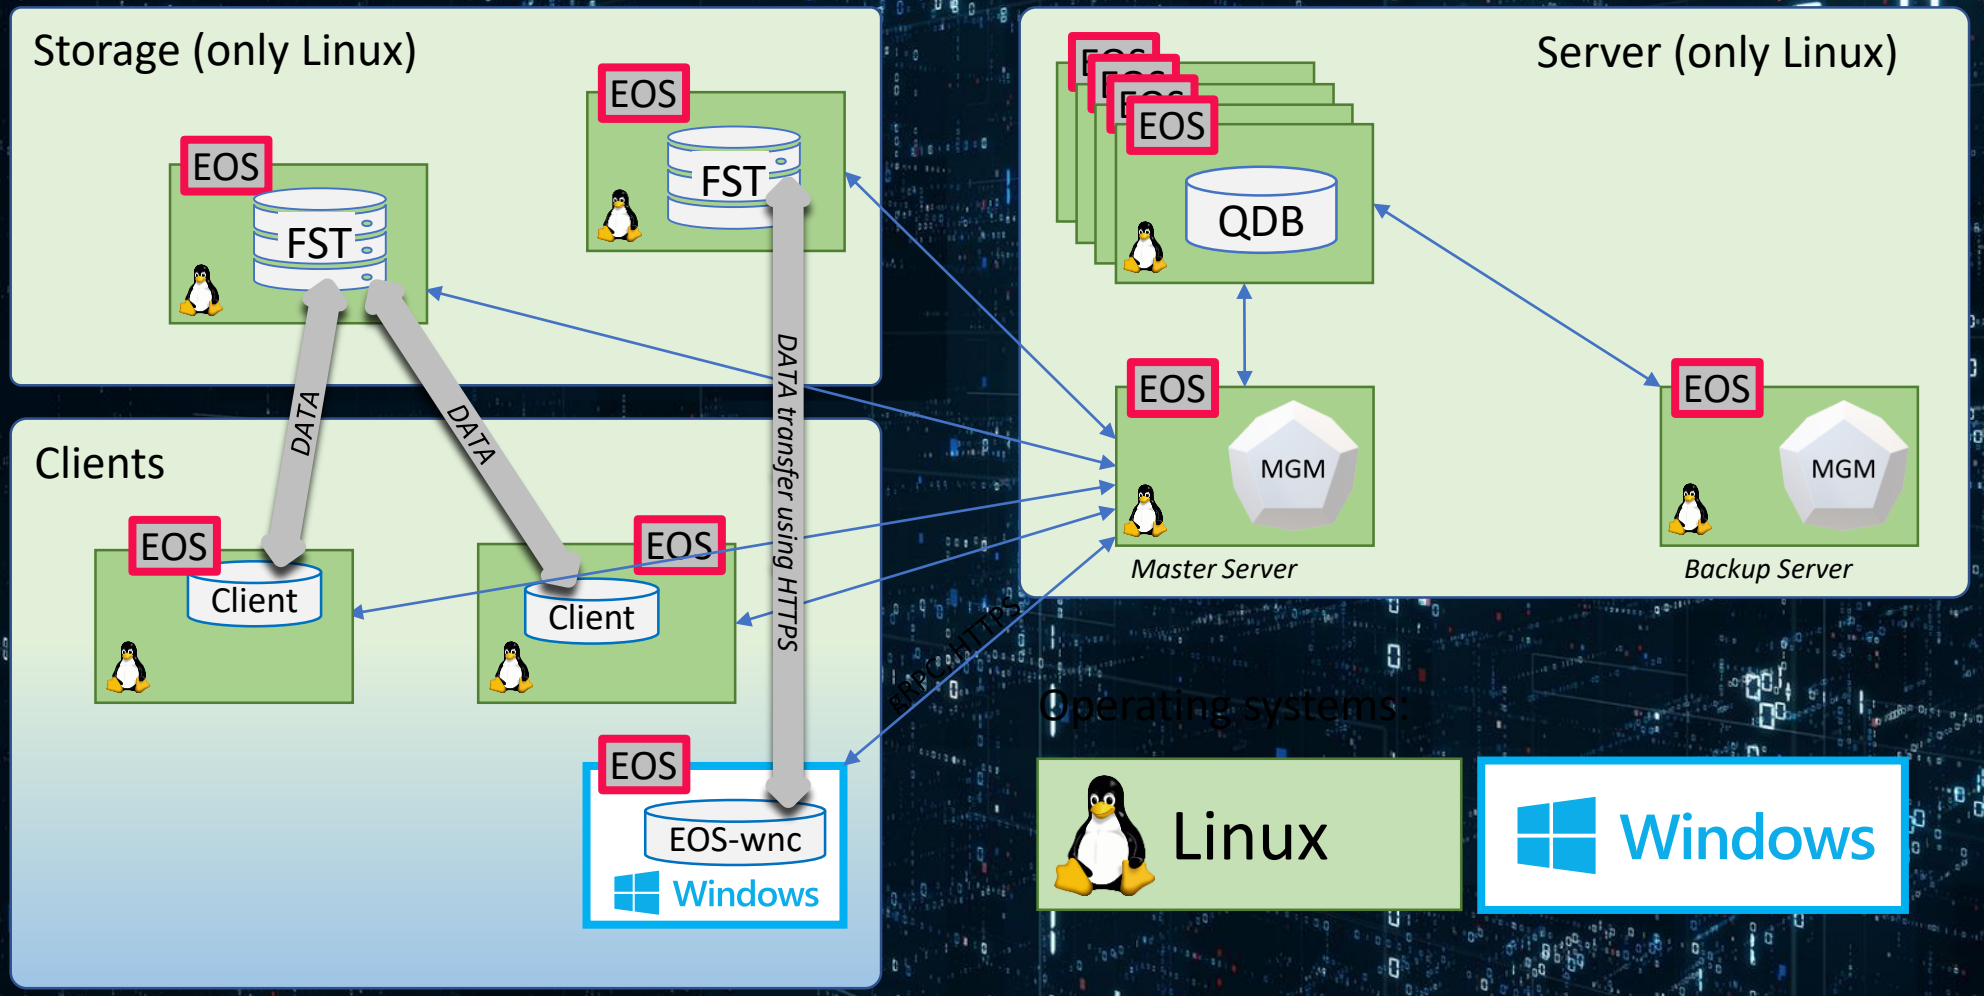

EOS CLI on Windows

- Not a port using Windows Subsystem for Linux (WSL)
- Completely new Windows EOS client
- New solution for networking issues on Windows
- New solution for security issues on Windows
- Technologies for EOS-wnc
  - Protocol Buffers
  - gRPC
  - cURL

Comtrade Three-Sixty

PS C:\Program Files (x86)\Comtrade\EOS-wnc>

Windows native access to EOS

Investigated Windows device driver development

- Windows Driver Frameworks (WDF)
- User-Mode Driver Framework (UMDF)
- **Kernel-Mode Driver Framework (KMDF) (Preferred)**

Decision: Thin driver for EOS-wnc

- Transfers commands
- Prevents complex bugs
- Device driver bugs are extremely expensive

Comtrade Three-Sixty

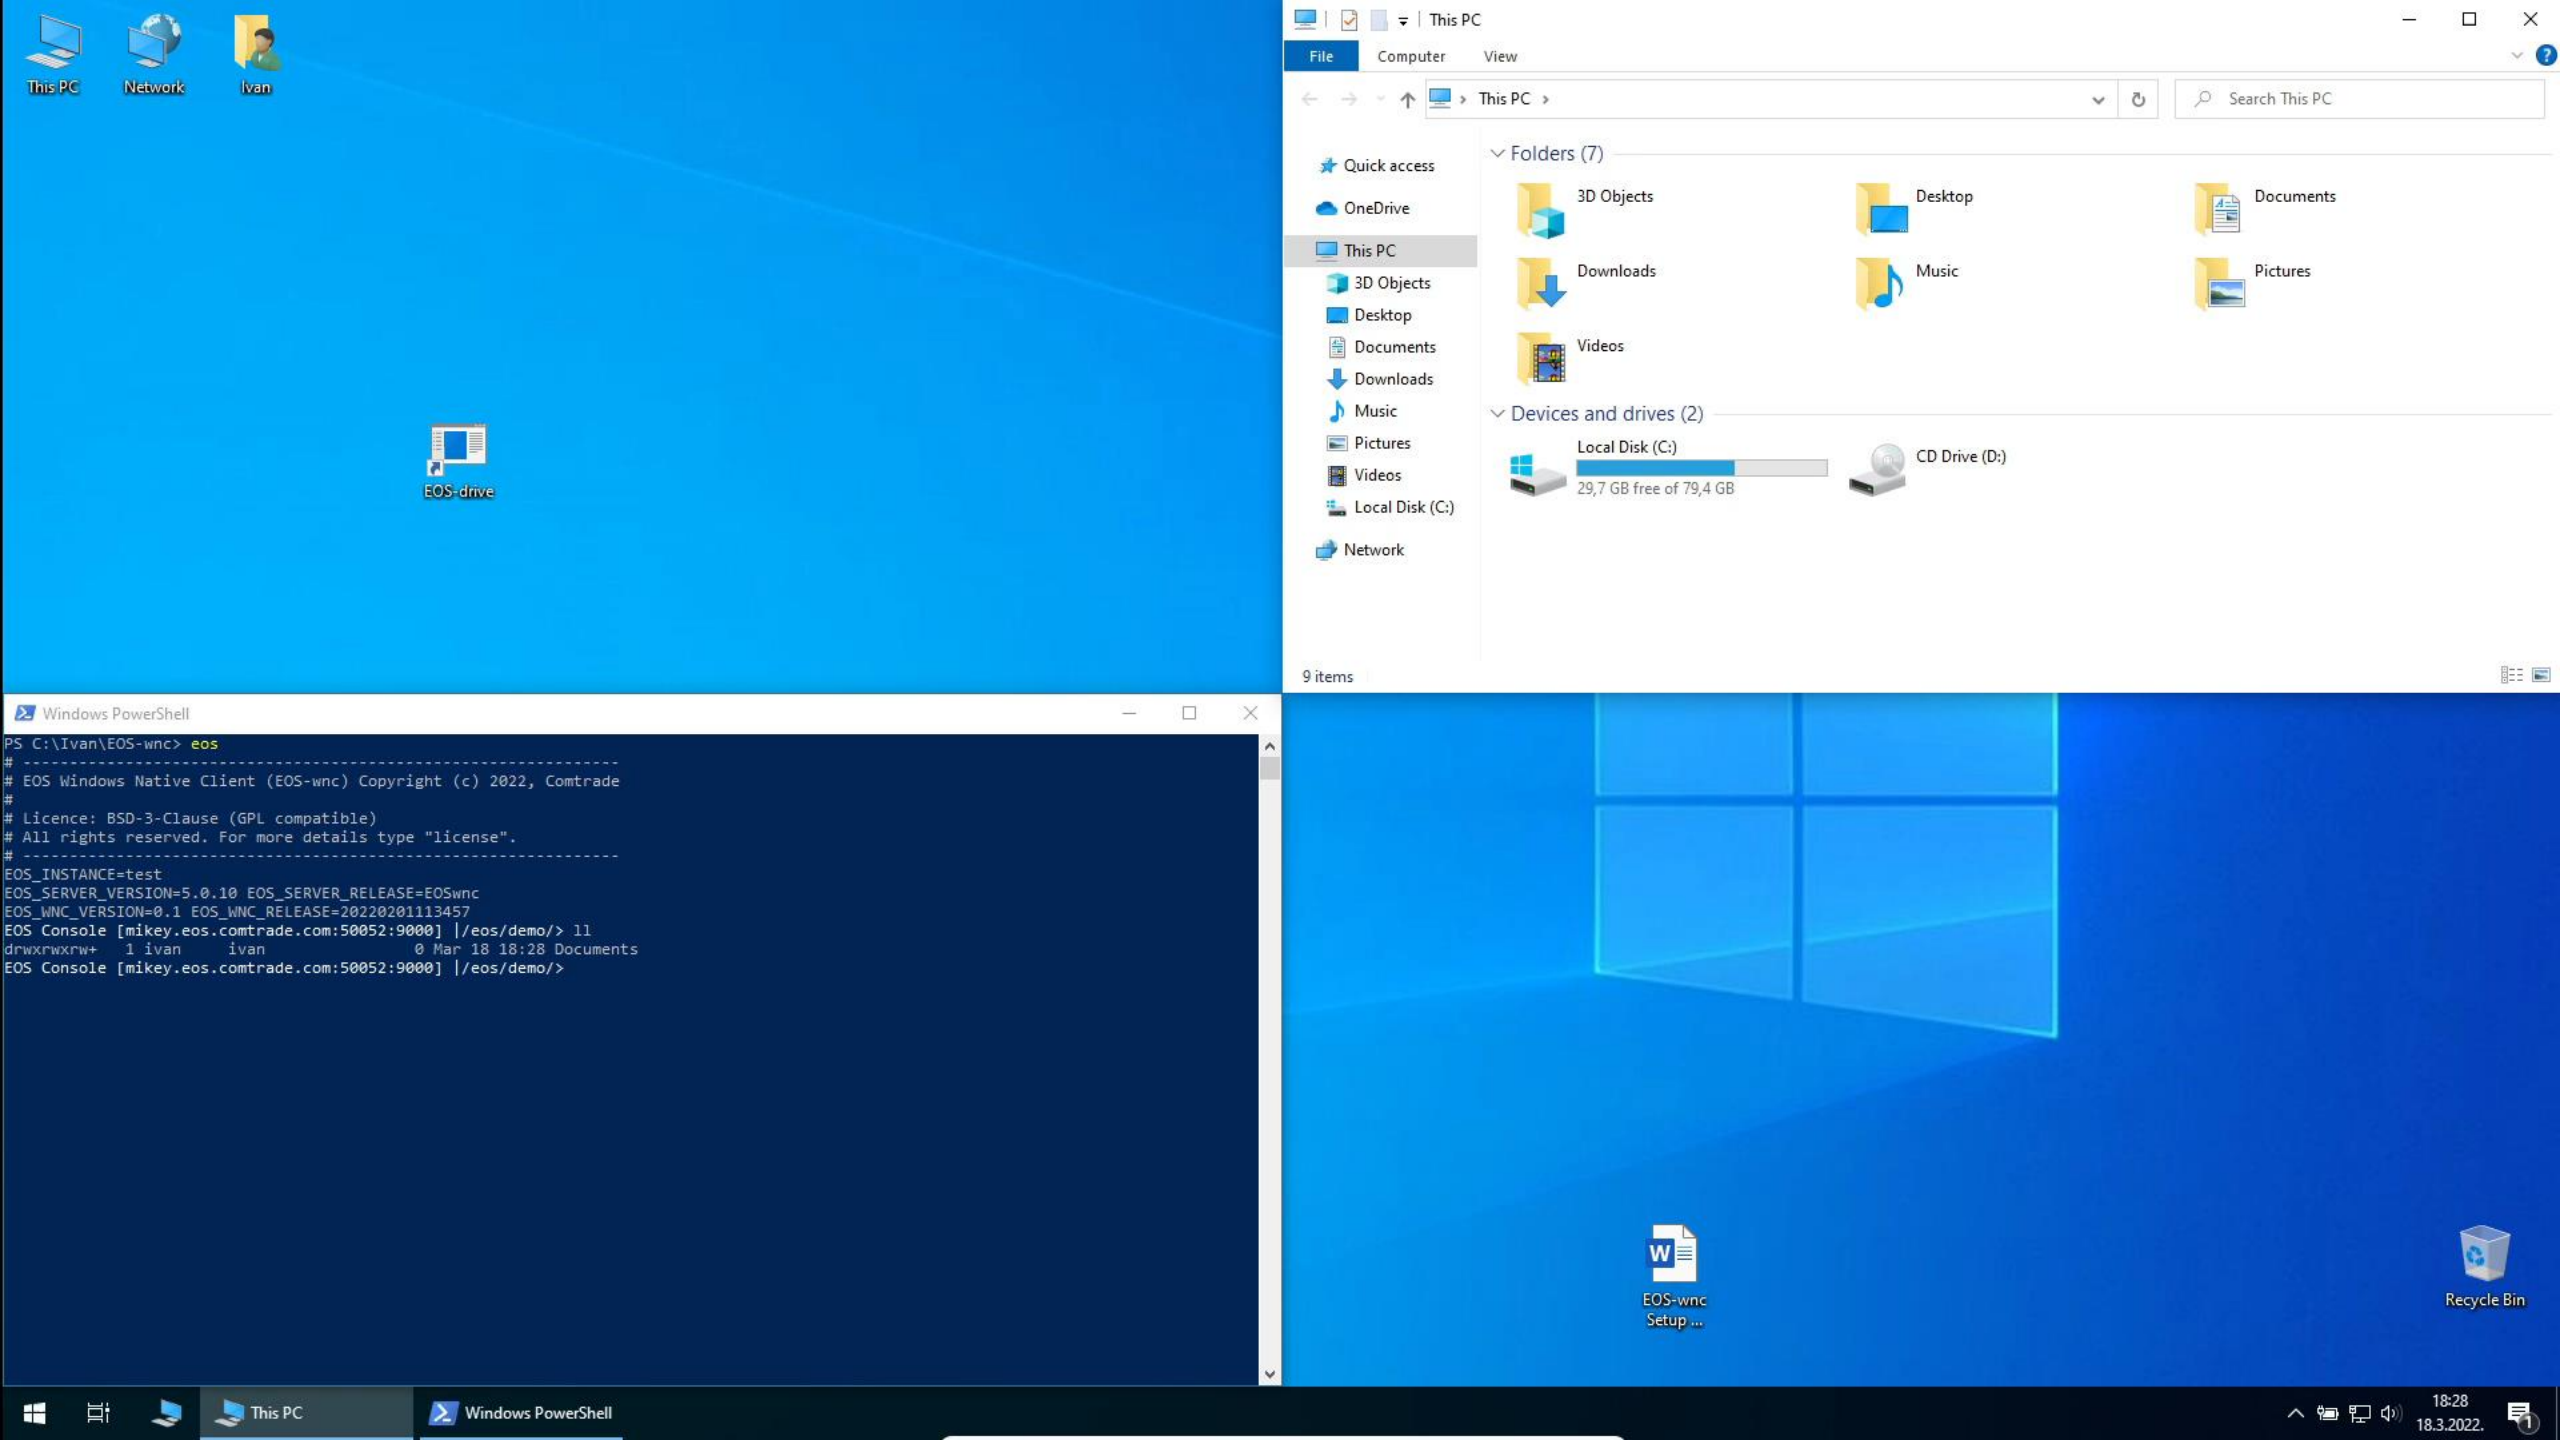
This PC
Network
Ivan
EOS-drive
This PC
File
Computer
View
This PC
Search This PC
Quick access
OneDrive
This PC
3D Objects
Desktop
Documents
Downloads
Music
Pictures
Videos
Local Disk (C:)
Network
Folders (7)
3D Objects
Desktop
Documents
Downloads
Music
Pictures
Videos
Devices and drives (2)
Local Disk (C:)
29,7 GB free of 79,4 GB
CD Drive (D:)
9 items
Windows PowerShell
PS C:\Ivan\EOS-wnc> eos
--
EOS Windows Native Client (EOS-wnc) Copyright (c) 2022, Comtrade
#
# Licence: BSD-3-Clause (GPL compatible)
# All rights reserved. For more details type "license".
# --
EOS_INSTANCE=test
EOS_SERVER_VERSION=5.0.10 EOS_SERVER_RELEASE=EOSwnc
EOS_WNC_VERSION=0.1 EOS_WNC_RELEASE=20220201113457
EOS Console [mikey.eos.comtrade.com:50052:9000] |/eos/demo/> ll
drwxrwxrw+   1 ivan   ivan   0 Mar 18 18:28 Documents
EOS Console [mikey.eos.comtrade.com:50052:9000] |/eos/demo/>
EOS-wnc Setup ...
Recycle Bin
This PC
Windows PowerShell
18:28
18.3.2022.

Comtrade 360: future plans for EOS

- Testing of EOS-wnc at CERN Preproduction (*ongoing, need to complete*)
- EOS cluster
  - Business model for “open-source support”
- EOS-wnc
  - Business plan for “proprietary software”
- Comparison between Distributed File Systems (*started*)
  - CephFS
  - Hadoop
  - Lustre
  - CERN EOS
- Update Comtrade’s professional documentation for EOS

Comtrade Three-Sixty

Thanks for the attention

- Gregor Molan
- Luca Mascetti
- Manuel Rodrigo Batalha Reis
- Elvin Alin Sindrilaru
- Andreas Joachim Peters
- Branko Blagojević
- Ivan Arizanović
- Svetlana Milenković
- Veselin Jevrosimović
- Alexis Lope-Bello

Comtrade Three-Sixty