



Storage Operations at CERN: EOS and CERNBox

LHC 27 km

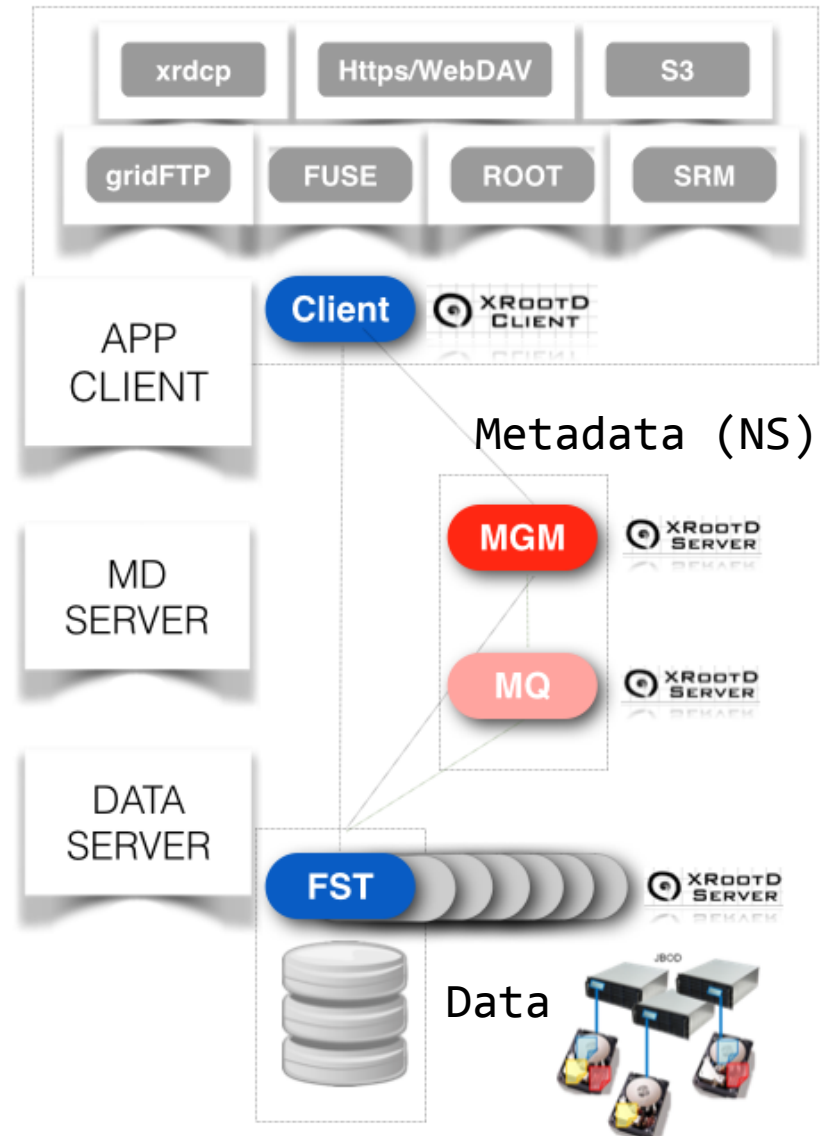
Luca Mascetti
CERN/IT-Storage

Outline

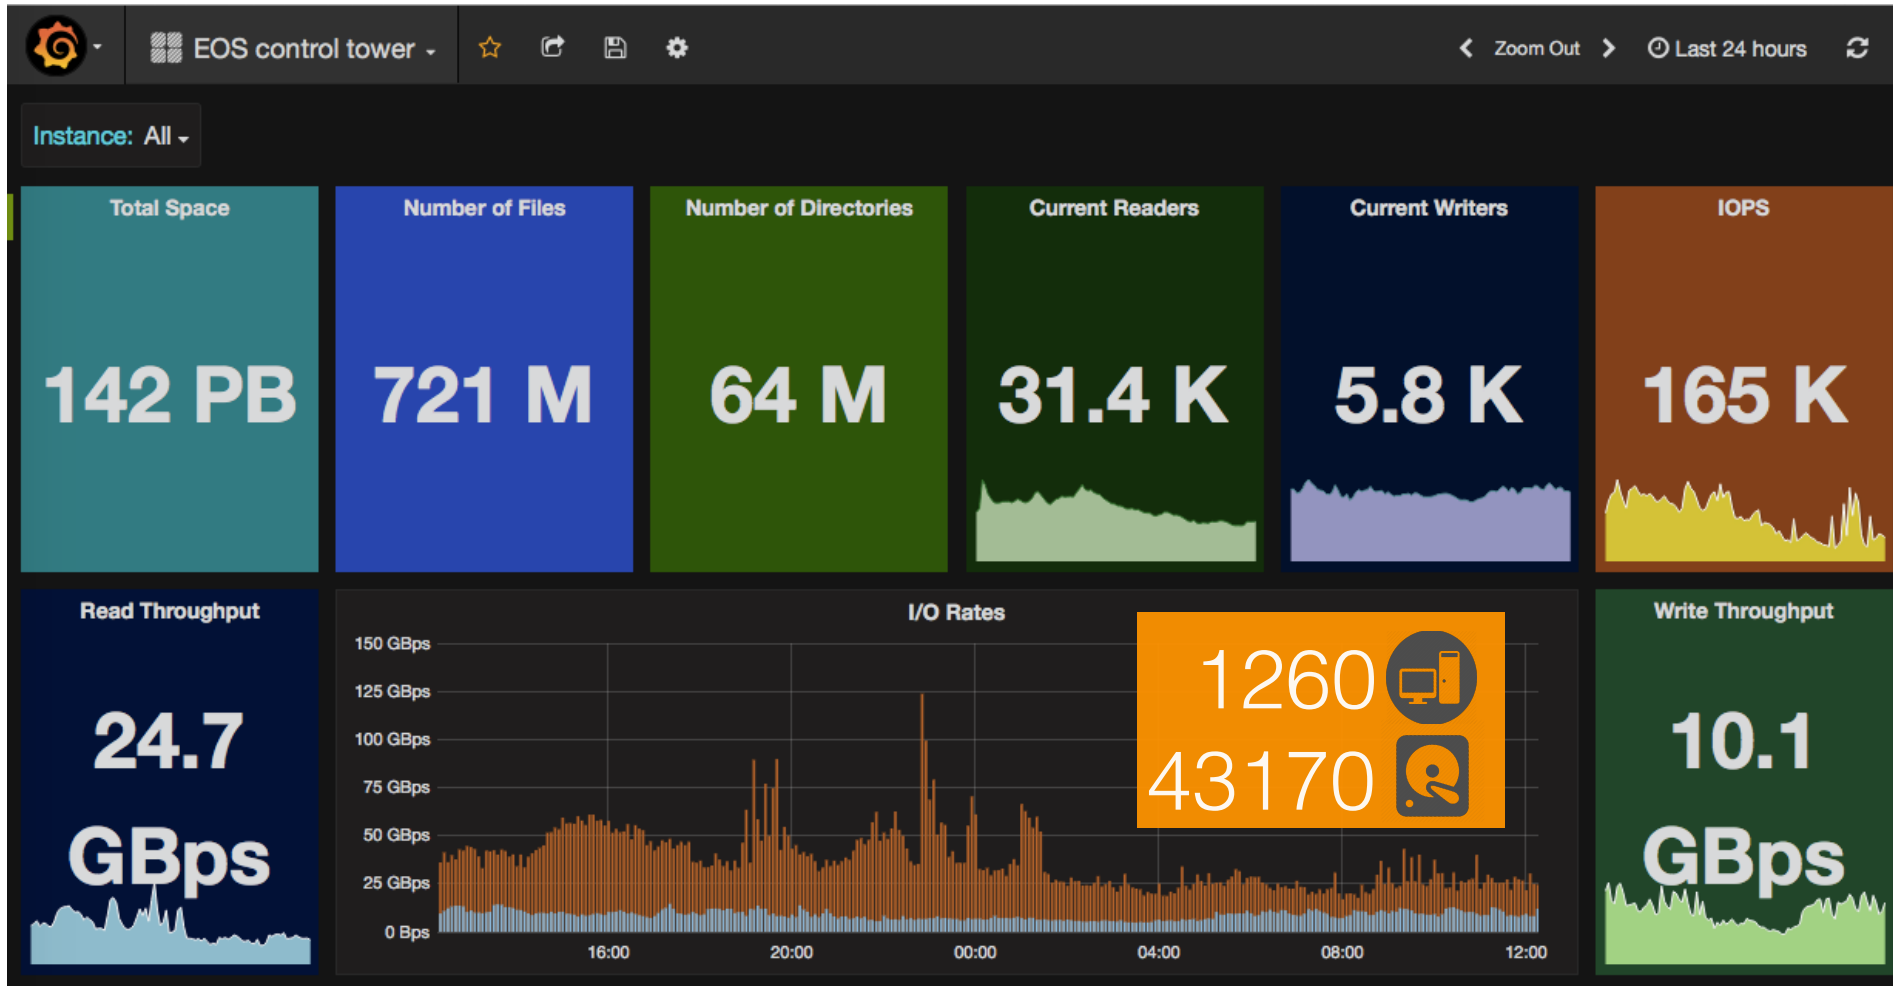
- **EOS Architecture**
- **EOS Service at CERN**
- **LHC Data Taking (Run 2)**
- **R&D multi-site instance**
- **CERNBox Architecture**
- **CERNBox Service at CERN**
- **Future HOME Directory**
- **Summary and Outlook**

EOS

- Project started in 2010
- Licence free
- Simple and scalable solution
- Easy to operate
- In-memory namespace
- Secure access (krb5, gsi)
- Quotas (user/group)
- Network RAID (RAIN)
- Tuneable QoS
- Dev&Ops in CERN/IT-ST




EOS Service at CERN



...and we just received **60 PB** to be added to the system!


Commodity Hardware

- Uniform storage for all our services
- Profiting from economy of scale
- System Unit:
 - 8 physical cores (16 virtual)
 - 64/128GB of RAM
 - 2x disk-tray of 24x 6TB HDDs
 - almost **290 TB** raw building block unit



EOS Service at CERN

EOS



140 PB

+60 PB



EOS 38 PB

EOS 24 PB



EOS 45 PB

EOS 14 PB



EOS 18 PB

EOS 1 PB



Public
AMS
COMPASS
NA61/NA62
others

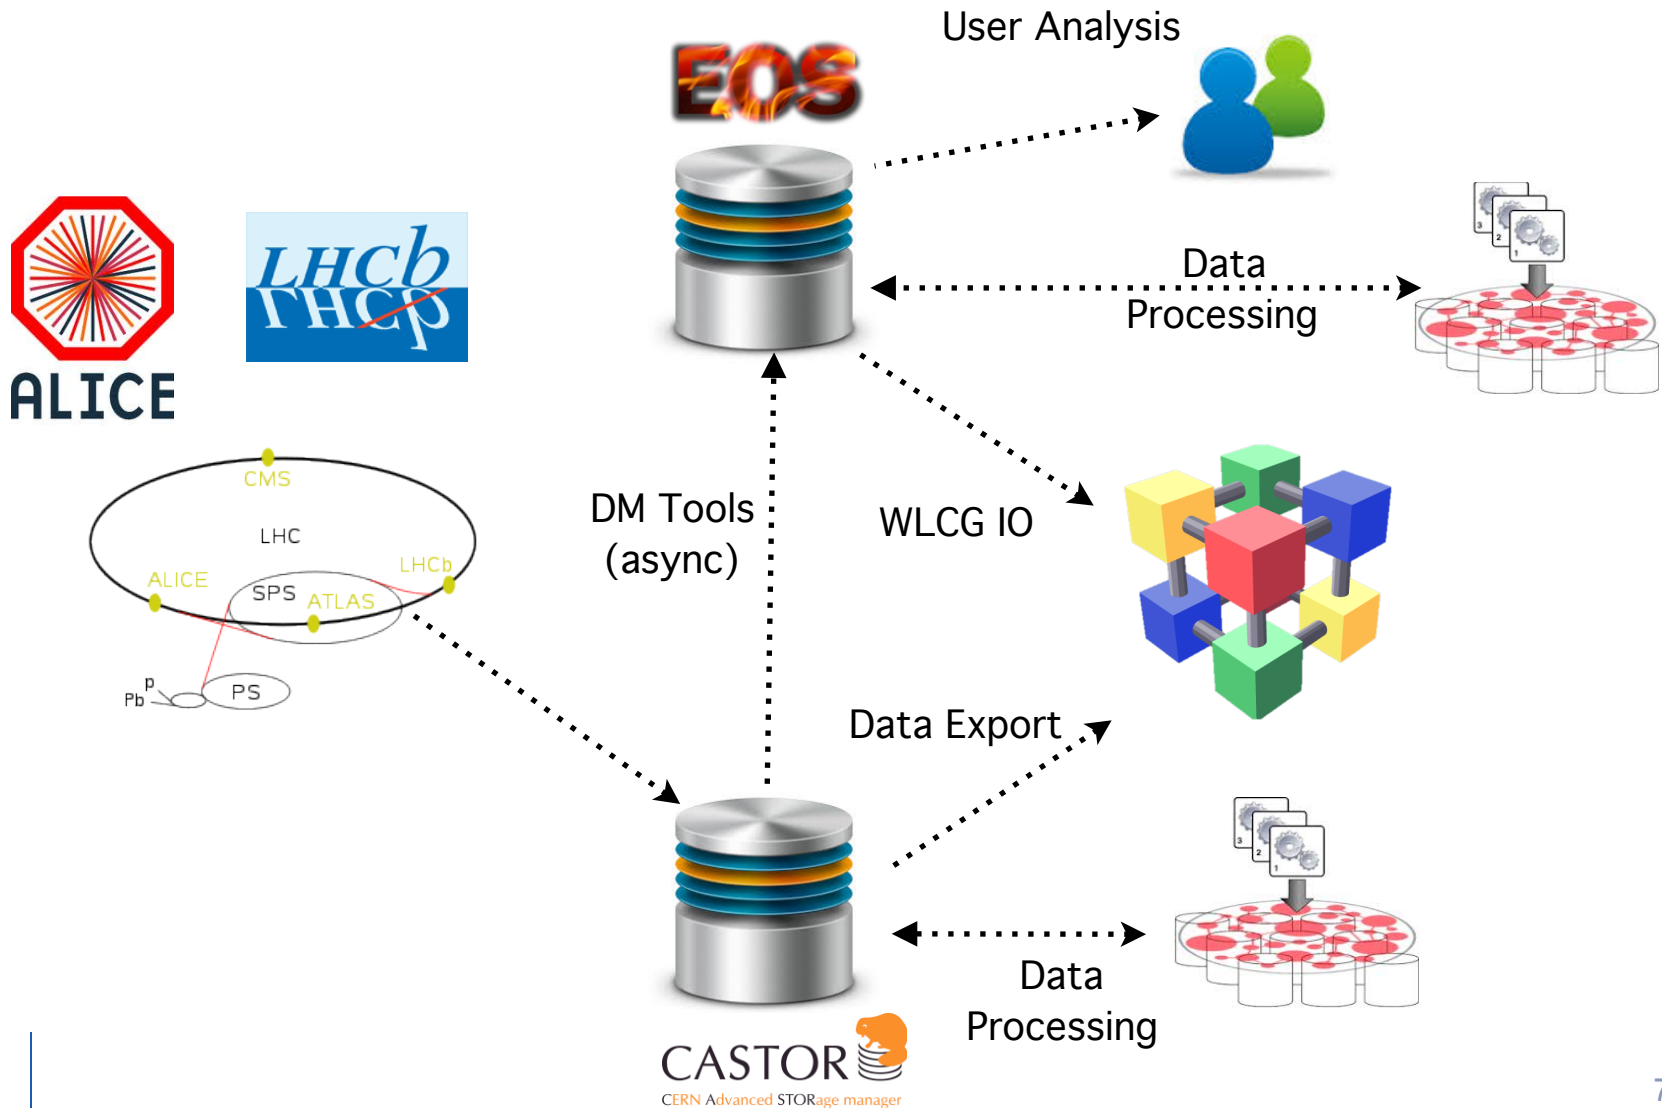
CERNBox



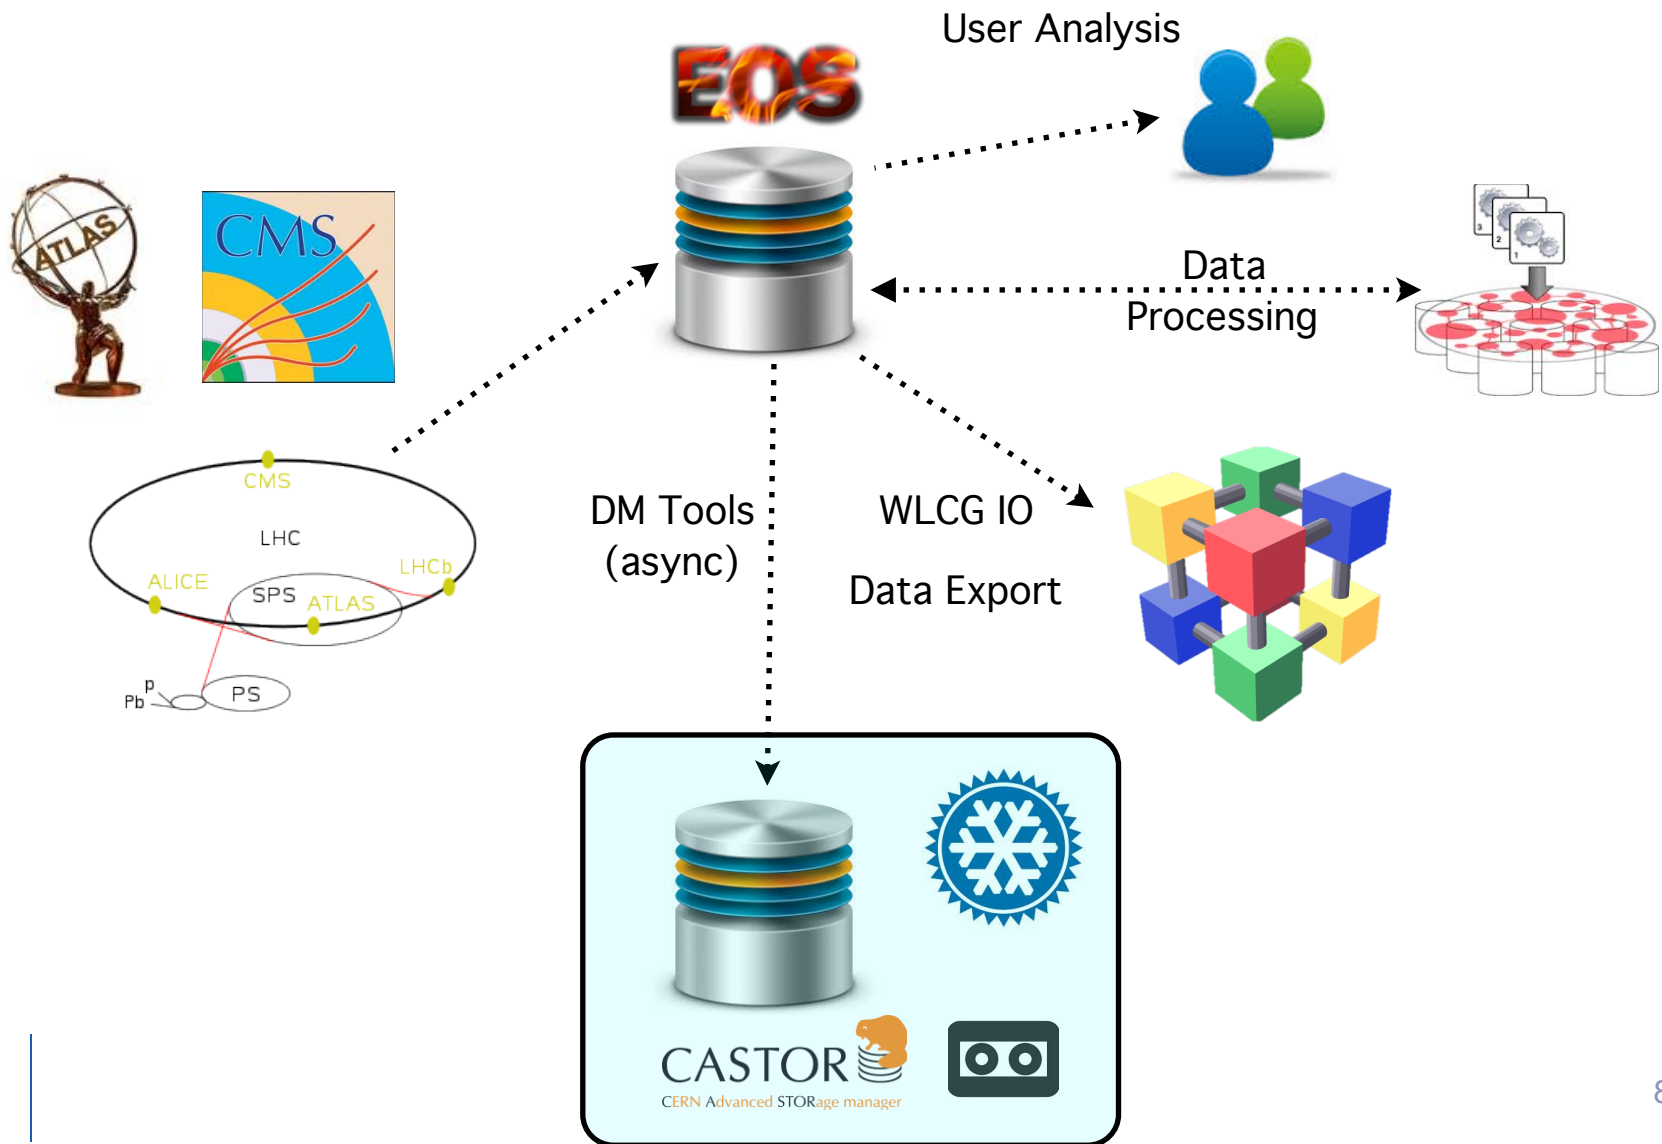
One instance per LHC Experiment



LHC Data Taking (Run 2)




LHC Data Taking (Run 2)



Data Rates during Run2


ALICE NETWORK USAGE (B/S)

EOS Out EOS In per 5m | (247784 hits)




ATLAS NETWORK USAGE (B/S)

EOS Out EOS In per 5m | (419480 hits)




PUBLIC NETWORK USAGE (B/S)

EOS Out EOS In per 5m | (153298 hits)




CMS NETWORK USAGE (B/S)

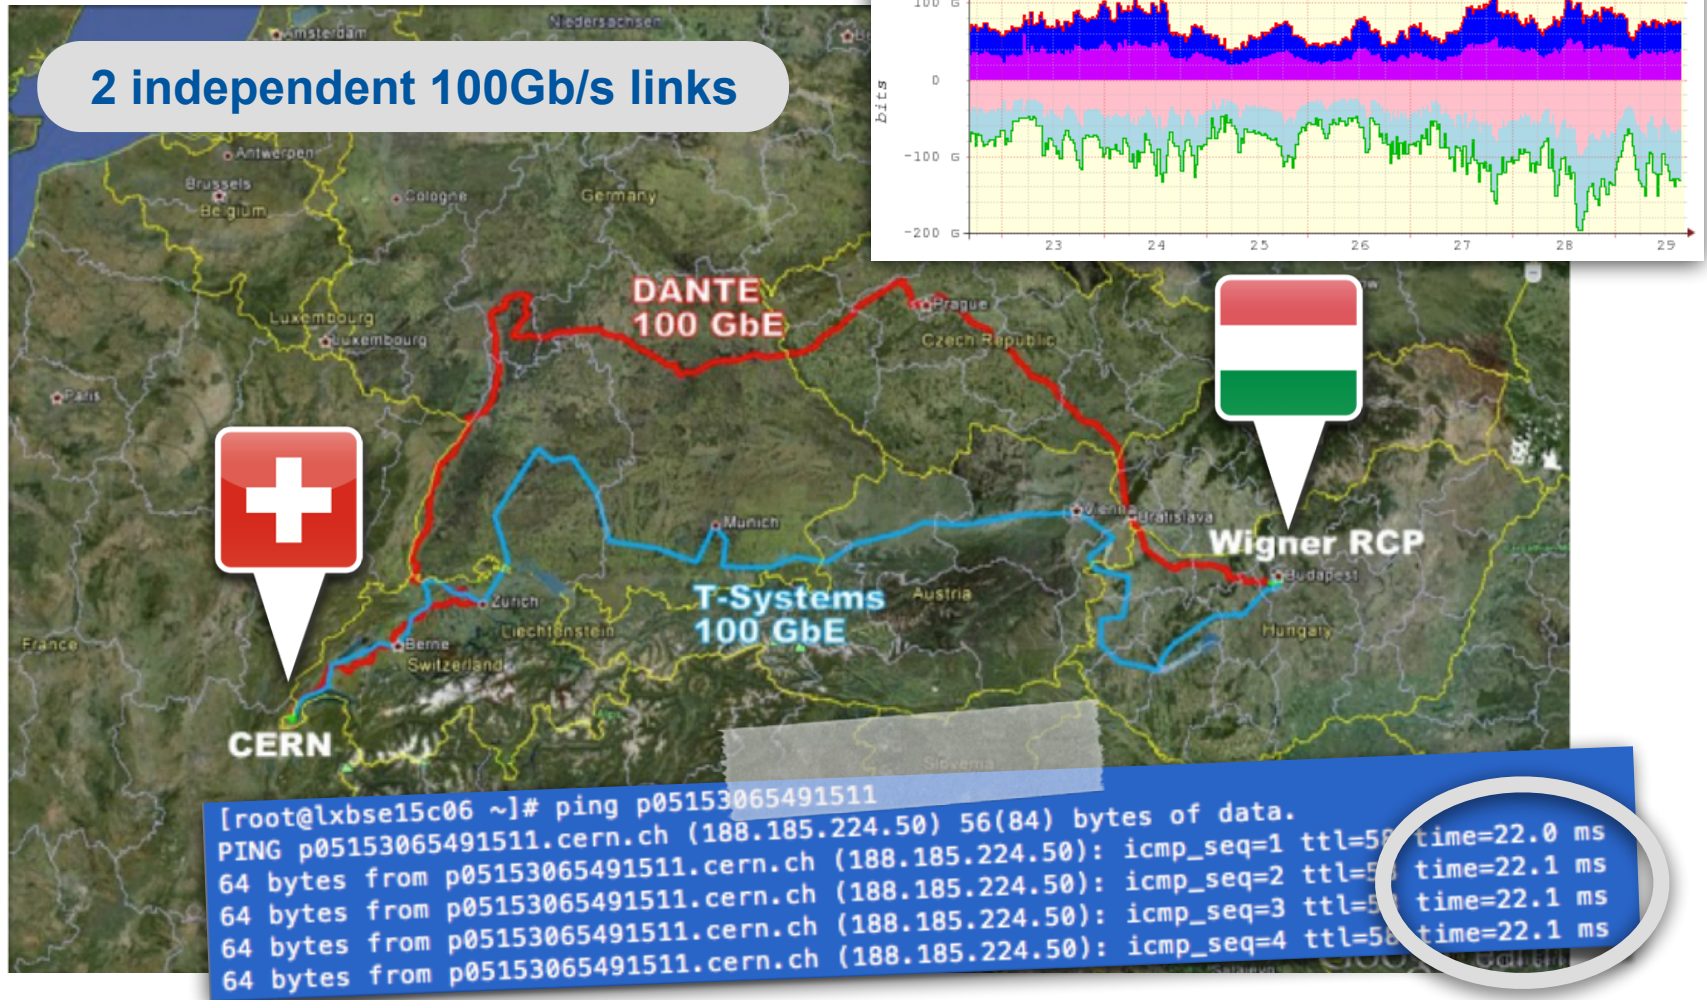
EOS Out EOS In per 5m | (397534 hits)



CERN-EOS in the ALICE Grid

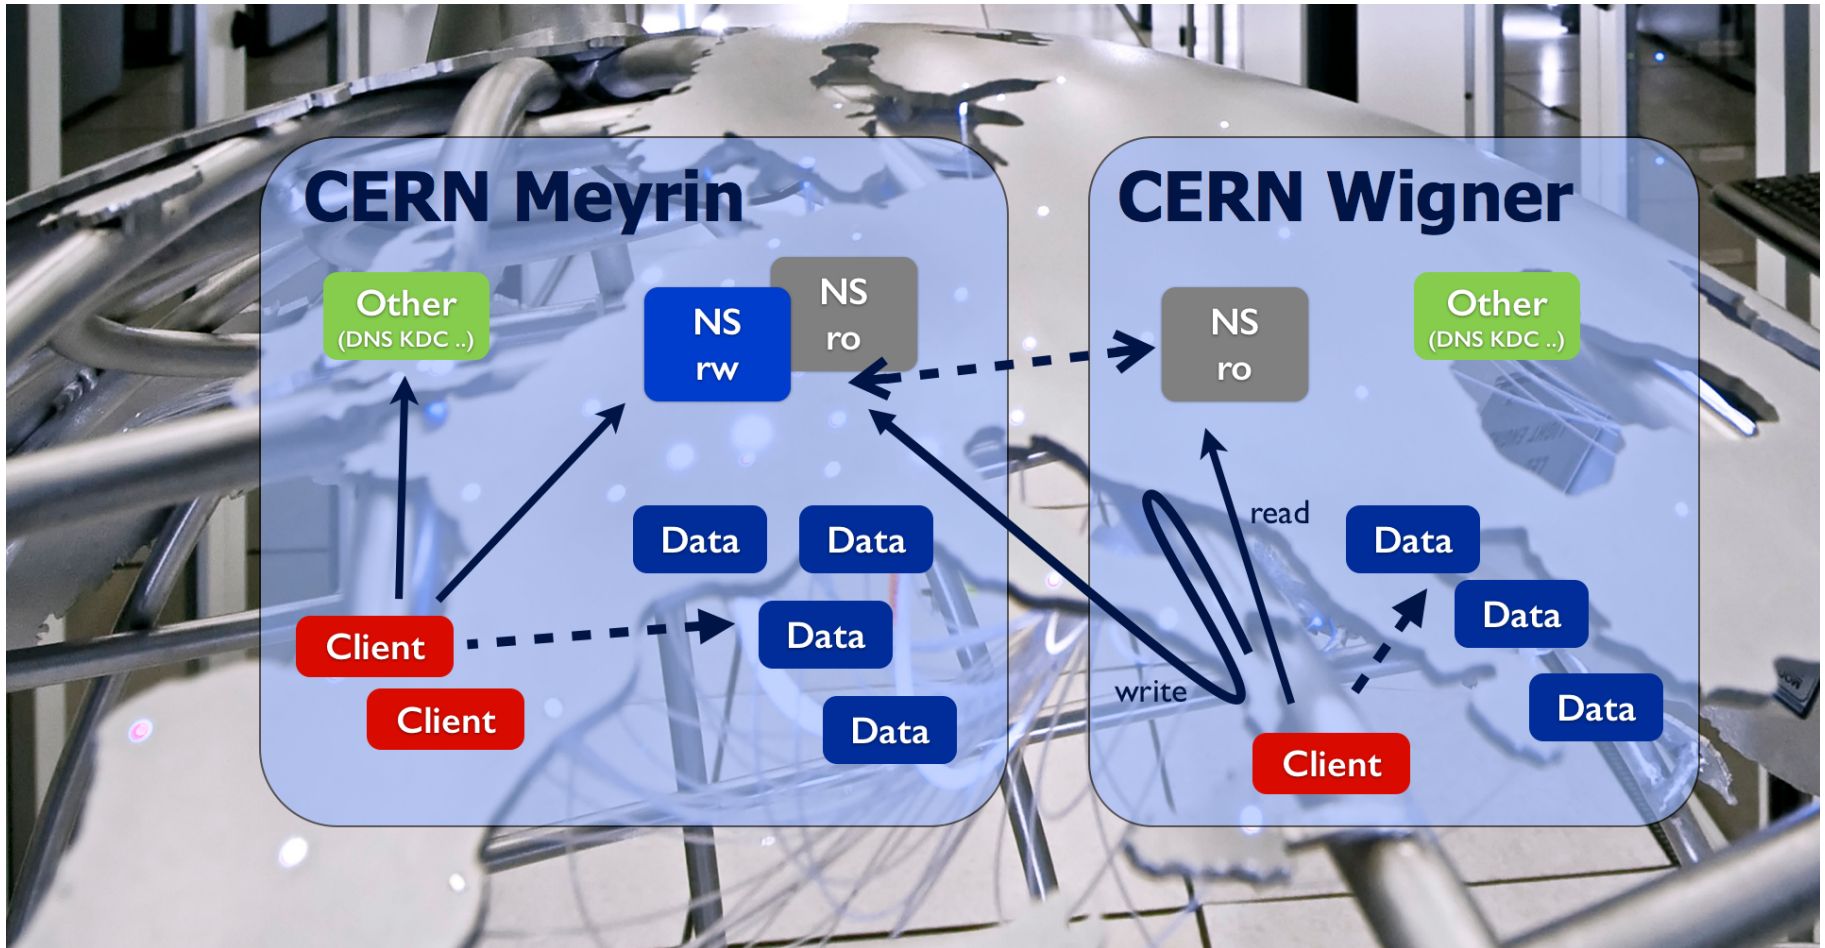


Wigner Computer Centre

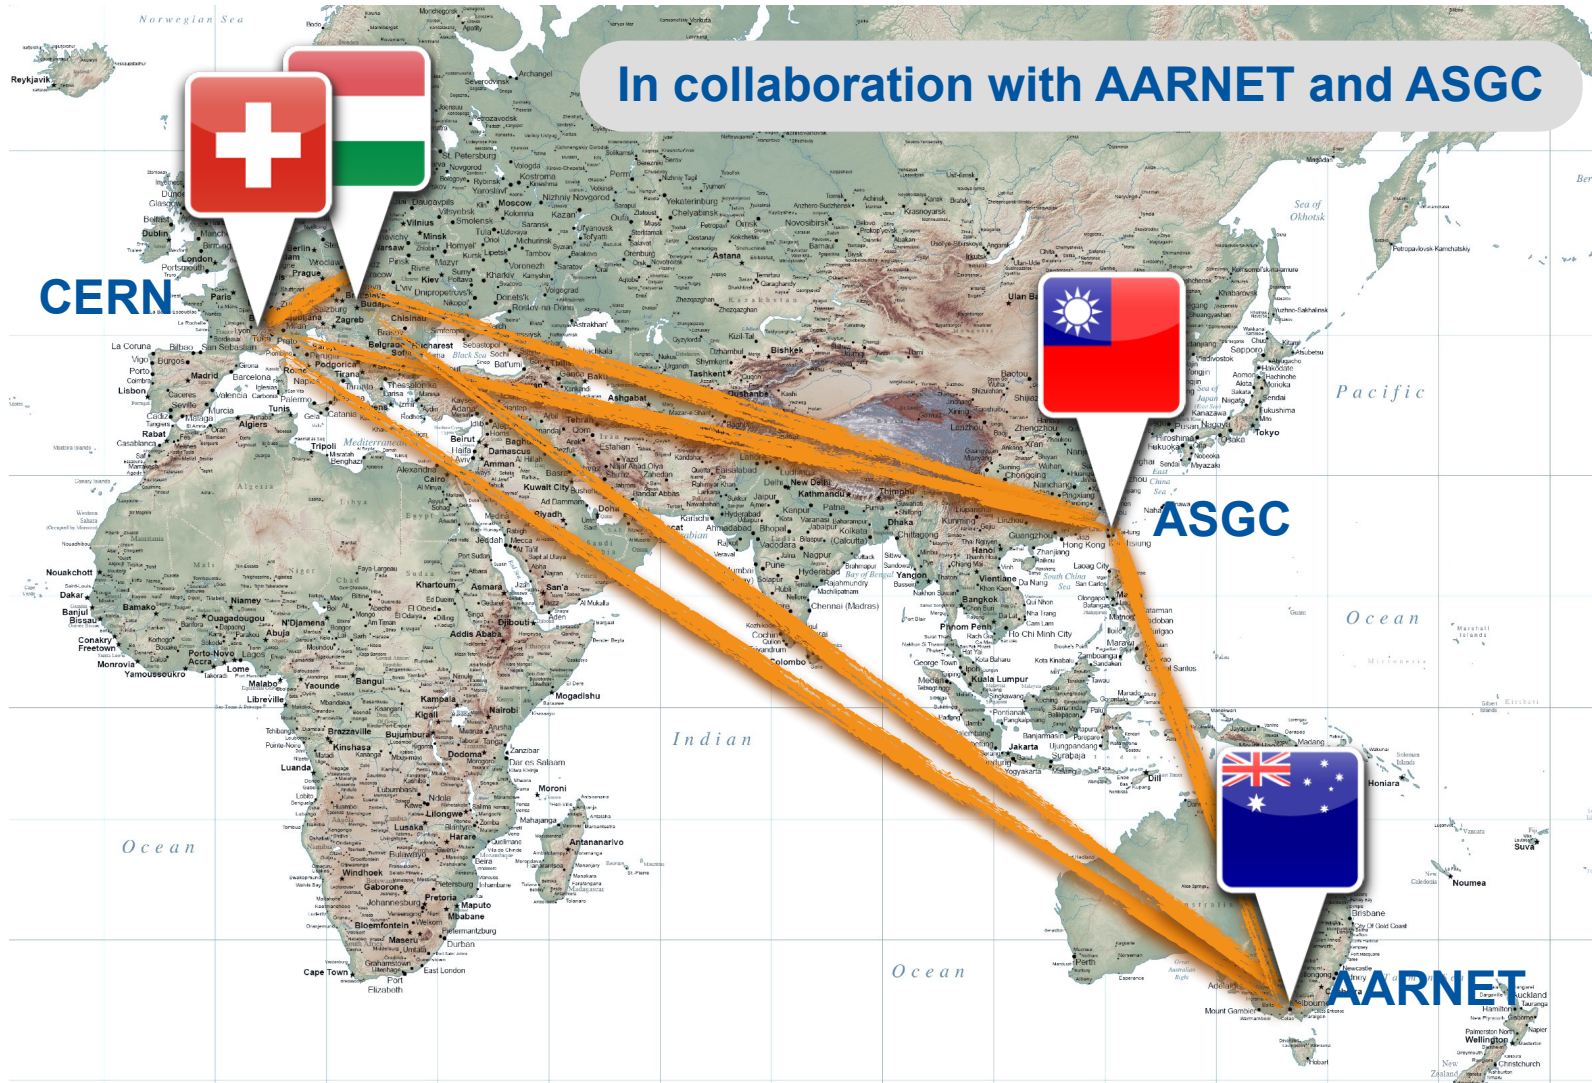


Waiting for the deployment of the third 100Gb/s link

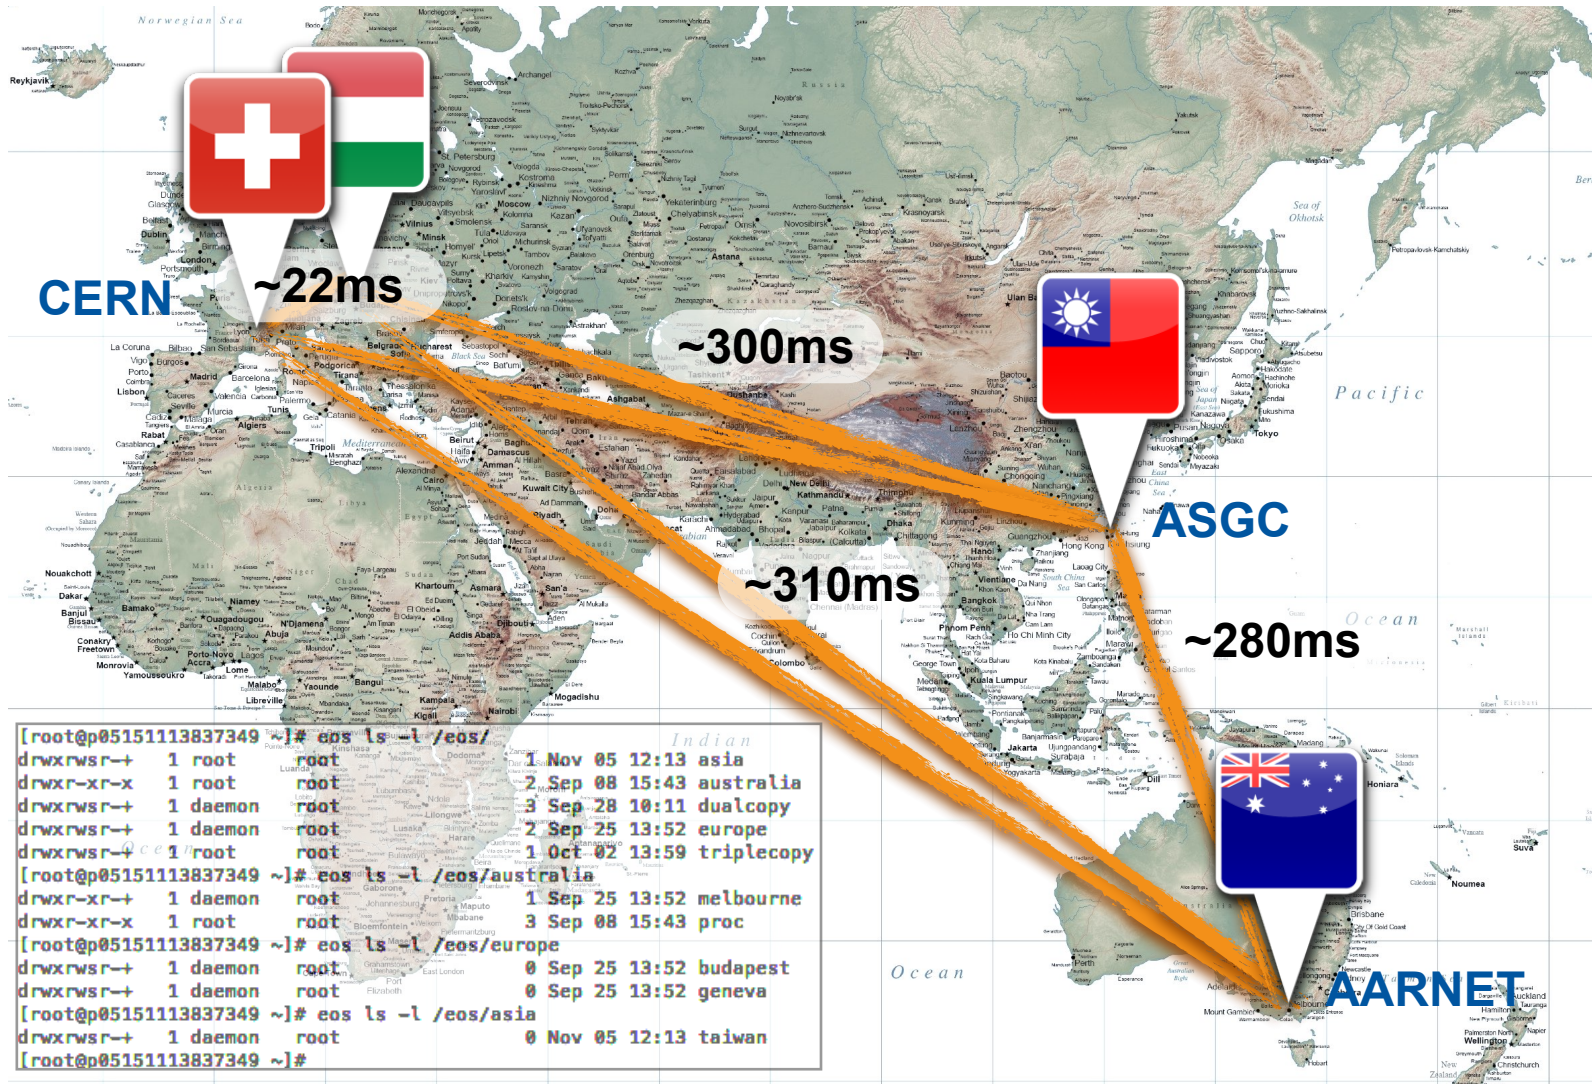
EOS across two sites



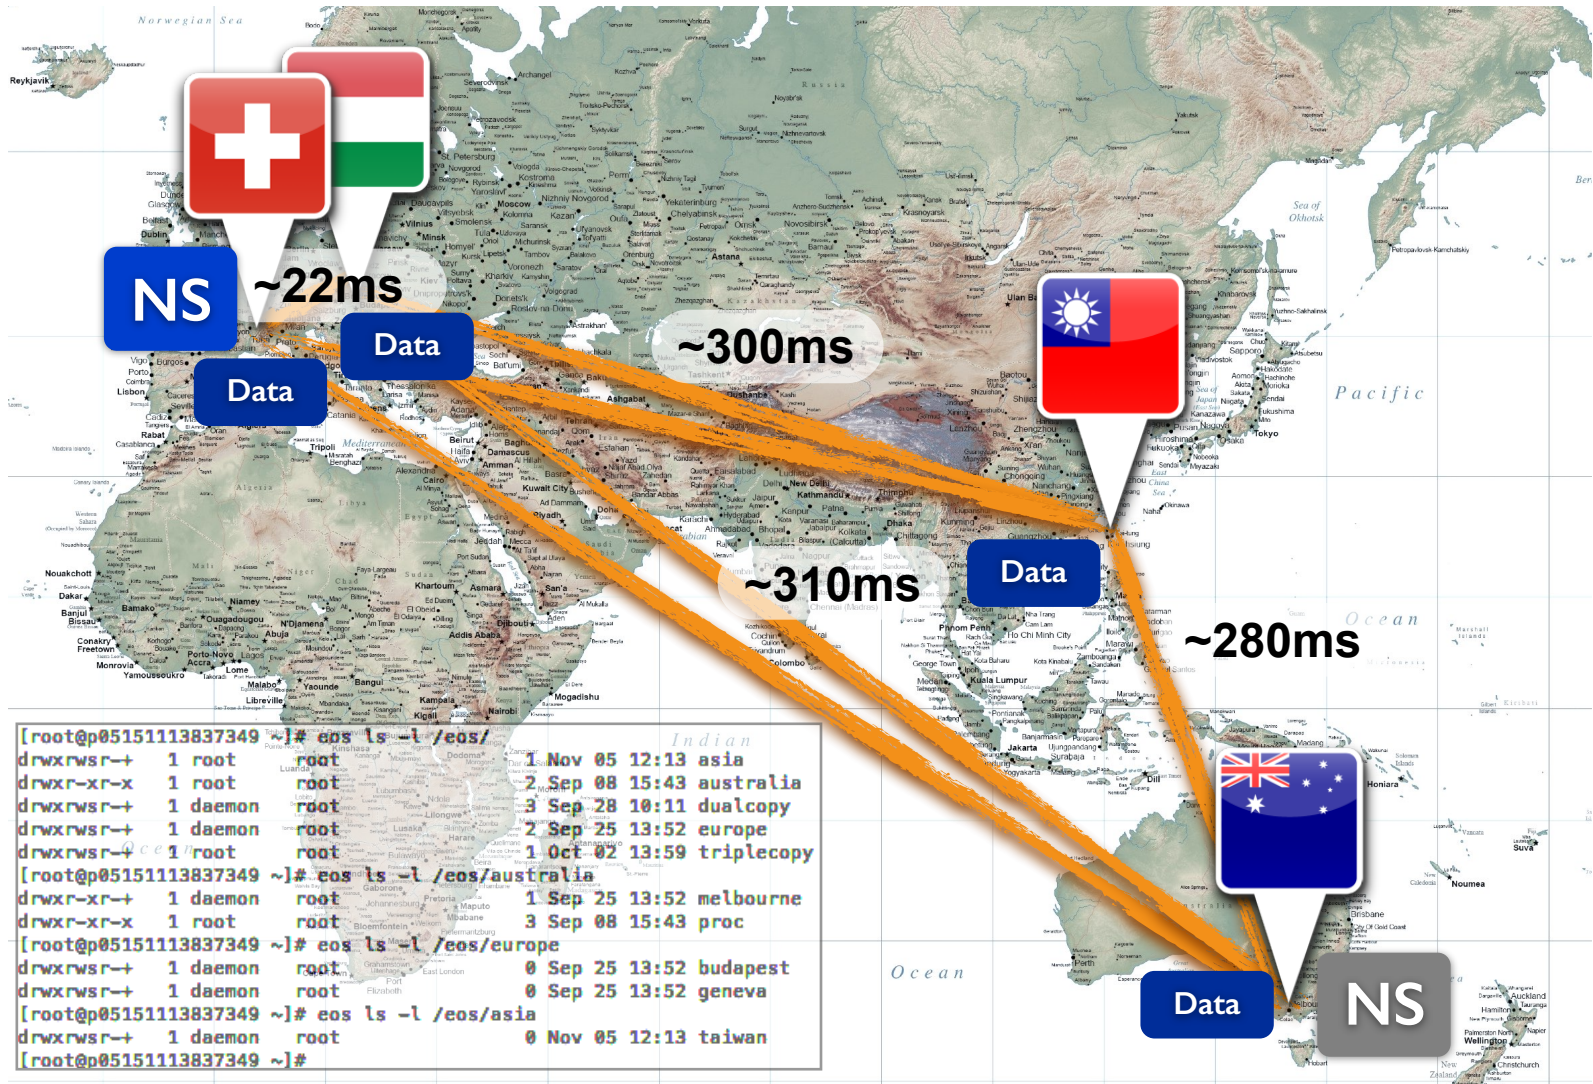
R&D - EOS World-Wide Deployment



R&D - EOS World-Wide Deployment



R&D - EOS World-Wide Deployment



R&D - EOS World-Wide Deployment

/eos

/asia

/taiwan

/australia

/melbourne

/europe

/geneva

/budapest

/dualcopy

/gva-bud



/mel-gva

/mel-bud

/triplecopy

/mel-gva-bud

/mel-gva-tpe



R&D - EOS World-Wide Deployment

/eos

/asia

/taiwan

/australia

/melbourne

/europe

/geneva

/budapest

/dualcopy

/gva-bud




/mel-gva

/mel-bud

/triplecopy

/mel-gva-bud

/mel-gva-tpe



R&D - EOS World-Wide Deployment

/eos

/asia

/taiwan

/australia

/melbourne

/europe

/geneva

/budapest

/dualcopy

/gva-bud



/mel-gva

/mel-bud

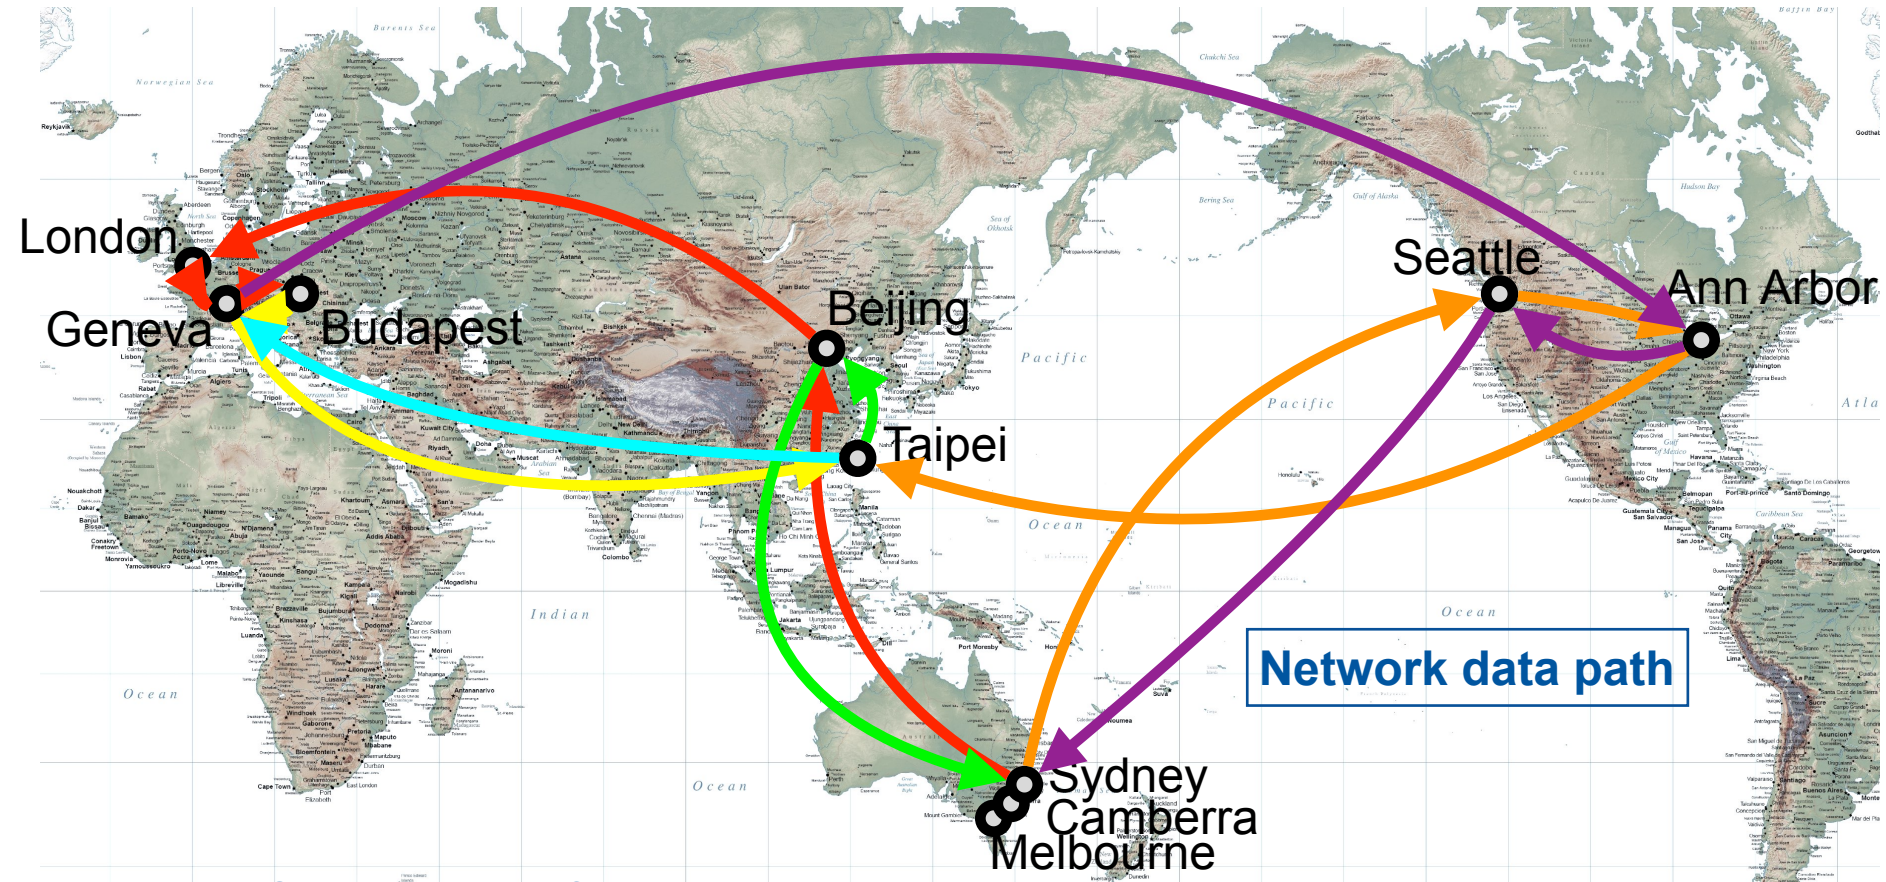
/triplecopy

/mel-gva-bud

/mel-gva-tpe



R&D - EOS World-Wide Deployment




- Streaming performance good
 - possible problems in case of packet drops (tcp window)
 - tcp settings could be optimised
- Latency in read hidden by the read-only NS
- Latency in write to contact the read-write NS

What is CERNBox ?

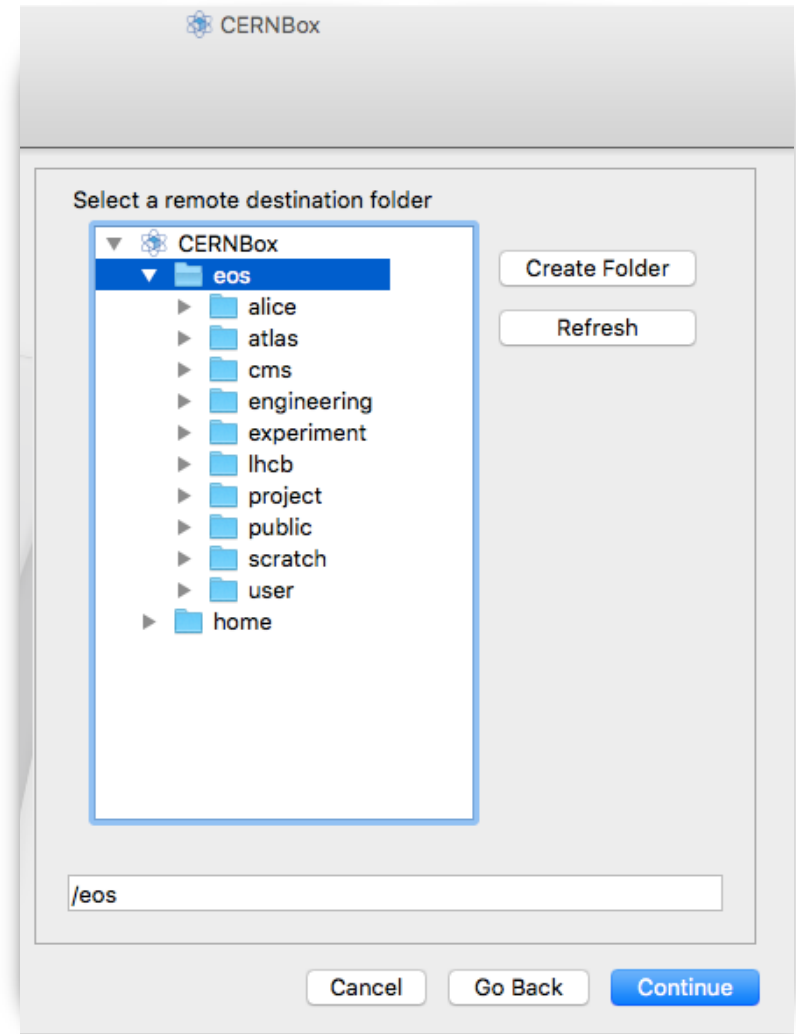



CERNBox provides a cloud synchronisation service

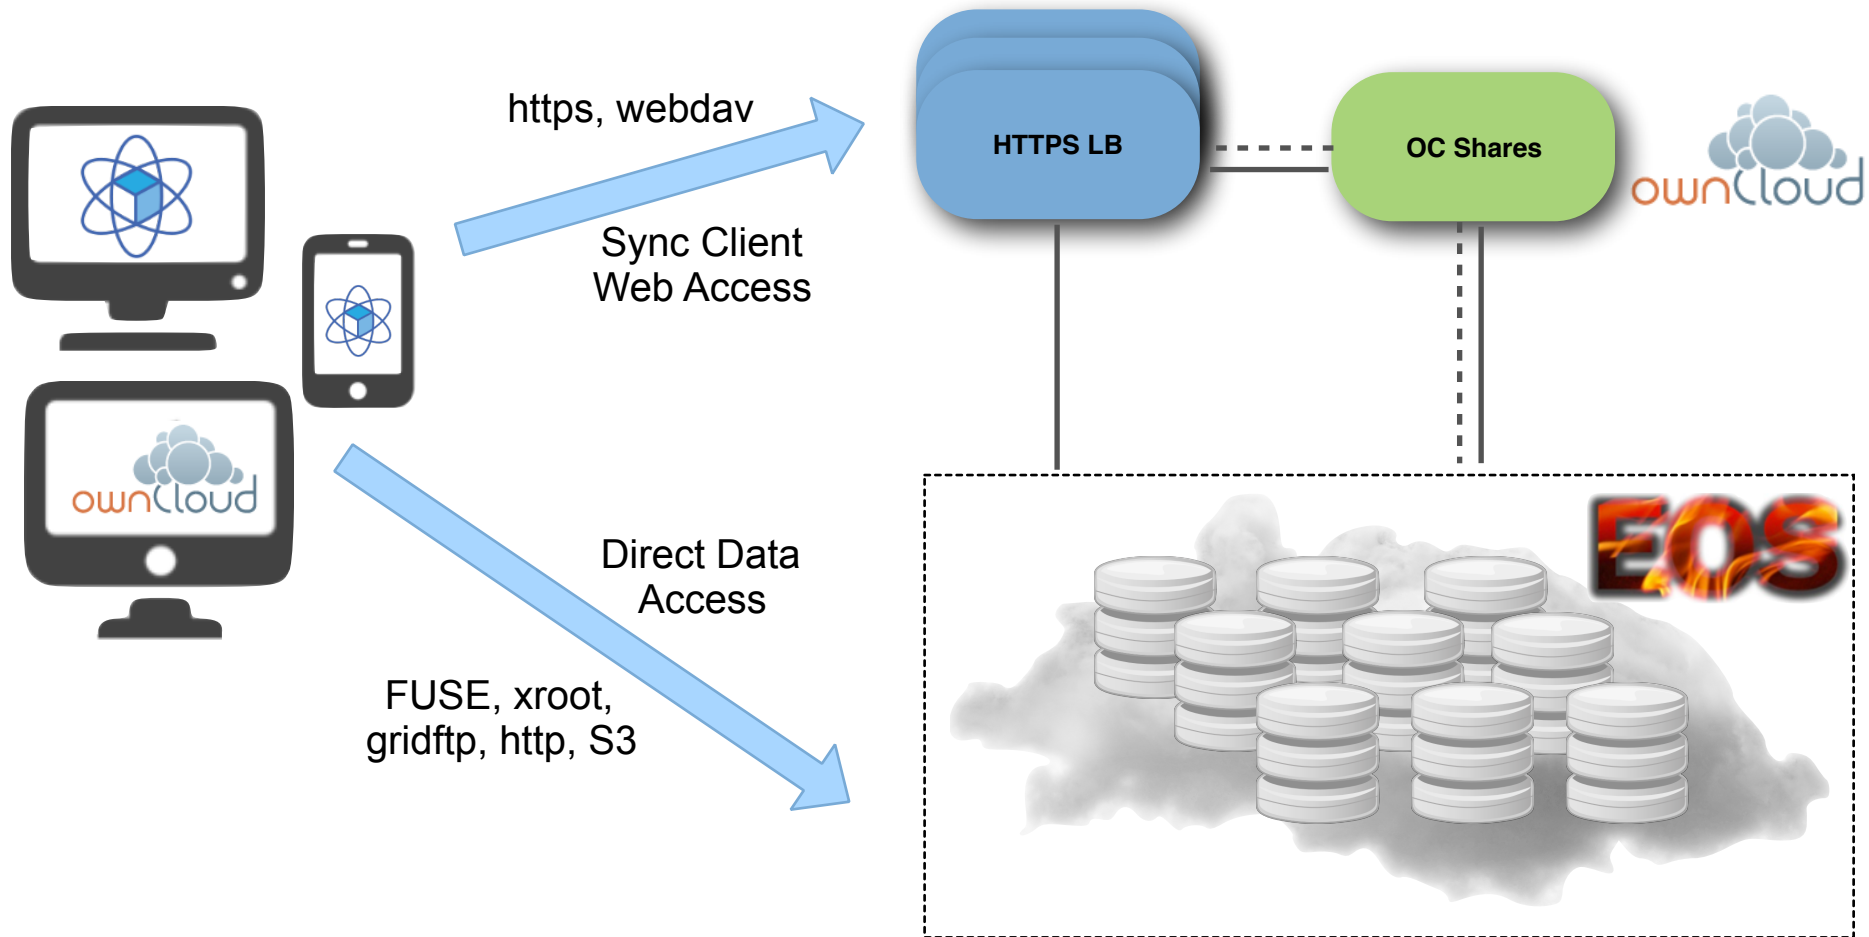
- Available for all CERN users (1TB/user)
- Synchronise files (data at CERN) and offline data access
- Easy way to share with other users
- All major platforms supported
- Based on **ownCloud** integrated with **EOS**



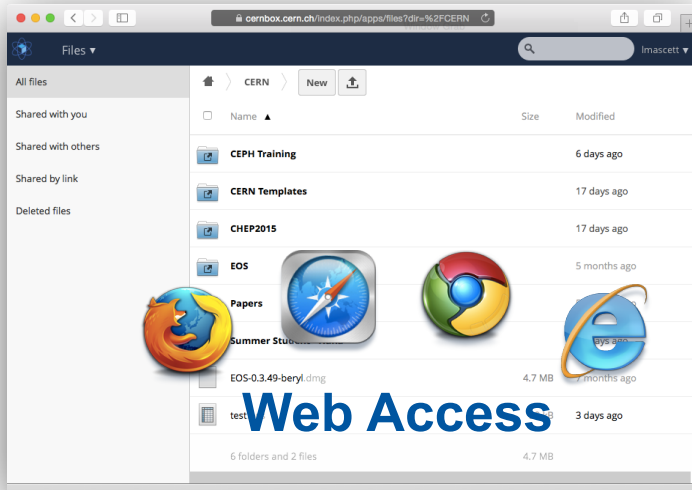
Bring data closer to our users: CERNBox



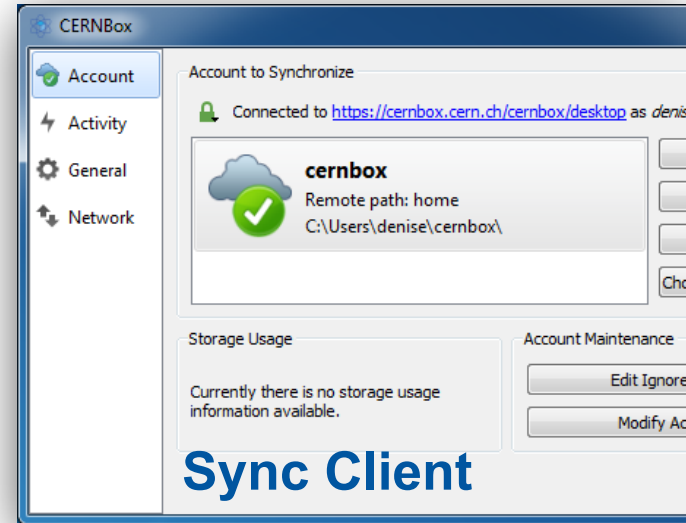
CERNBox Architecture




Available Access Methods



Mobile App



Sync Client



WebDAV



SAMBA



FUSE



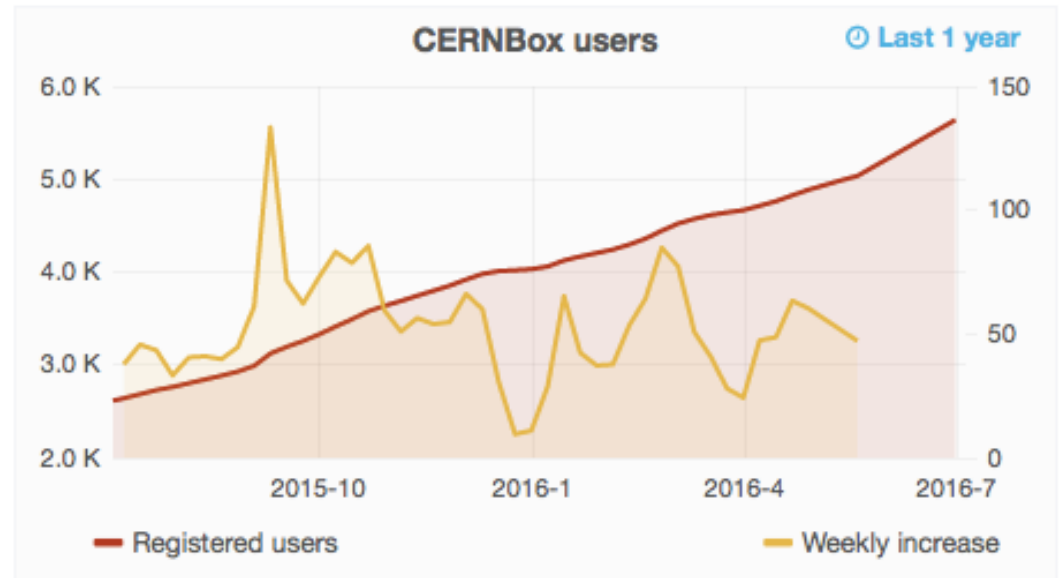
Directly from the storage backend (xroot, http, s3, ...)




EOSUSER/CERNBox Numbers

EOS offers “virtually unlimited” cloud-storage for our end-users


Users	5612
# files	83 Million
# dirs	11 Million
Quota	1TB/user
Used Space	173 TB
Deployed Space	1.3 PB




EOS/CERNBox HTTP Operations



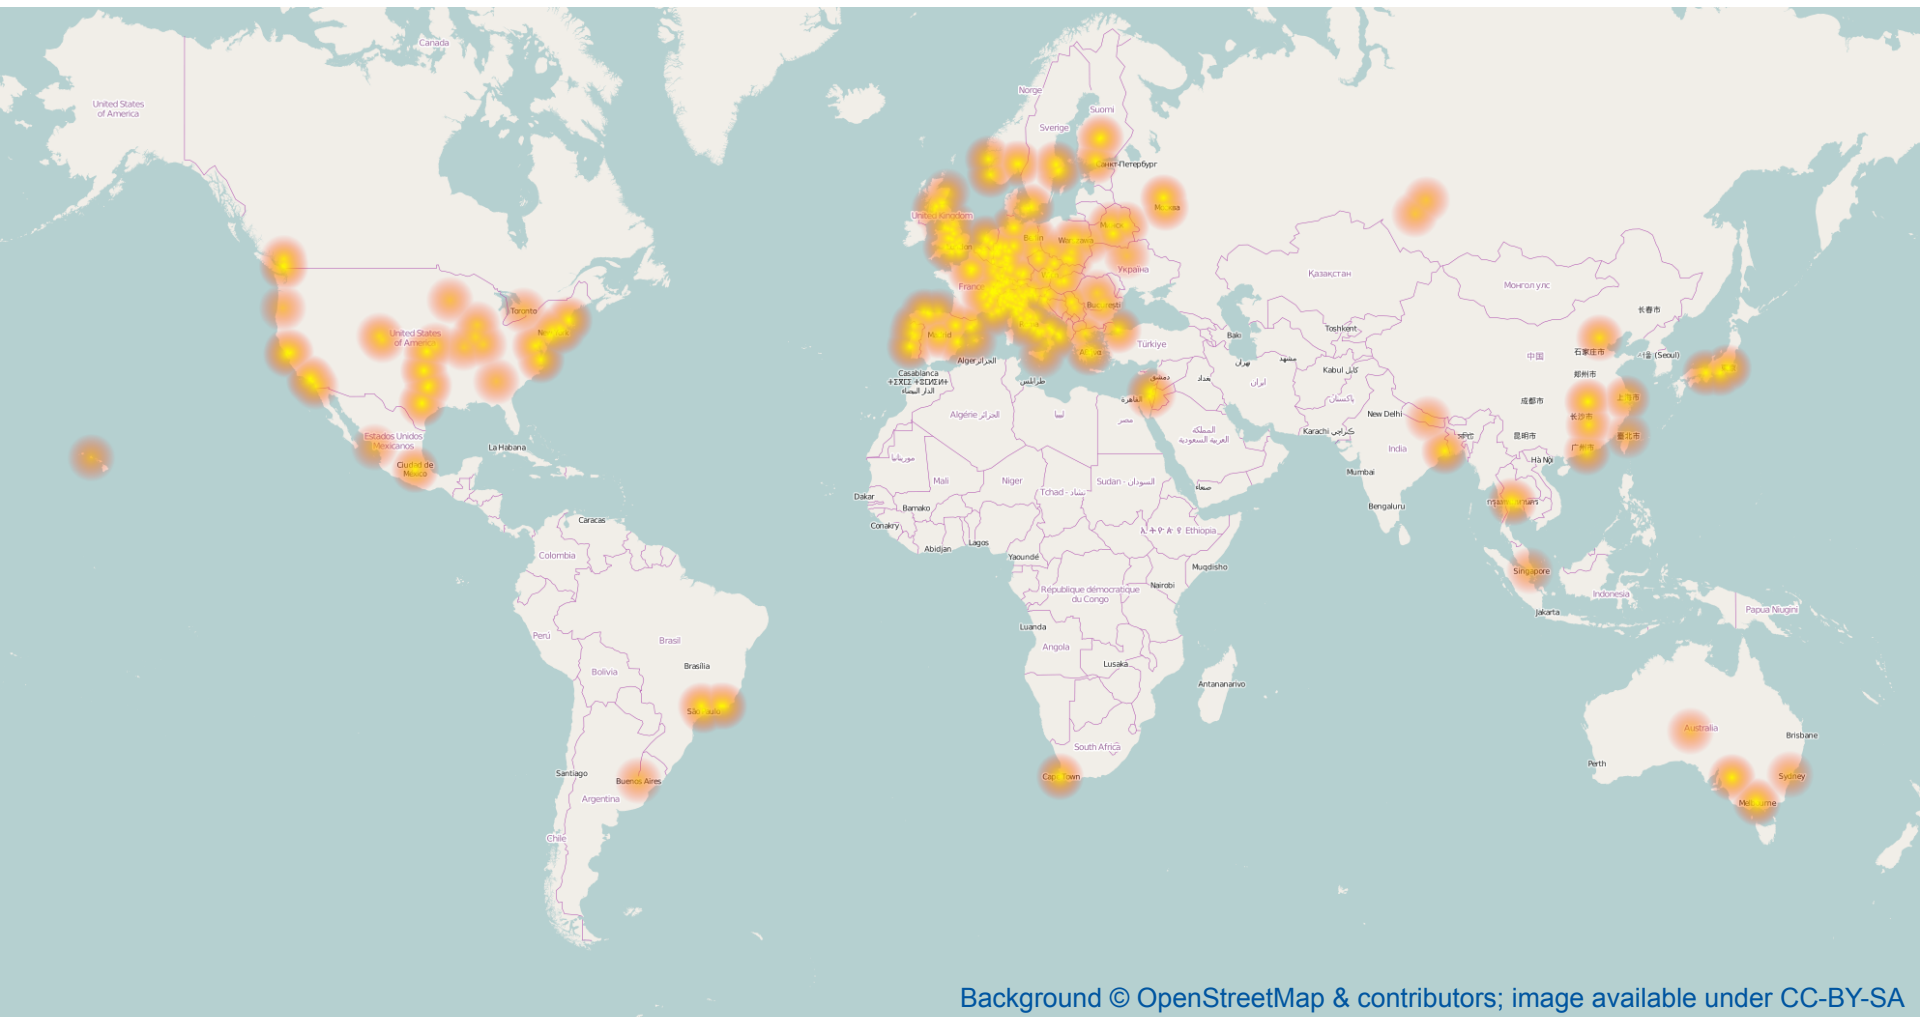
PUT and GET Requests



Distinct IPs

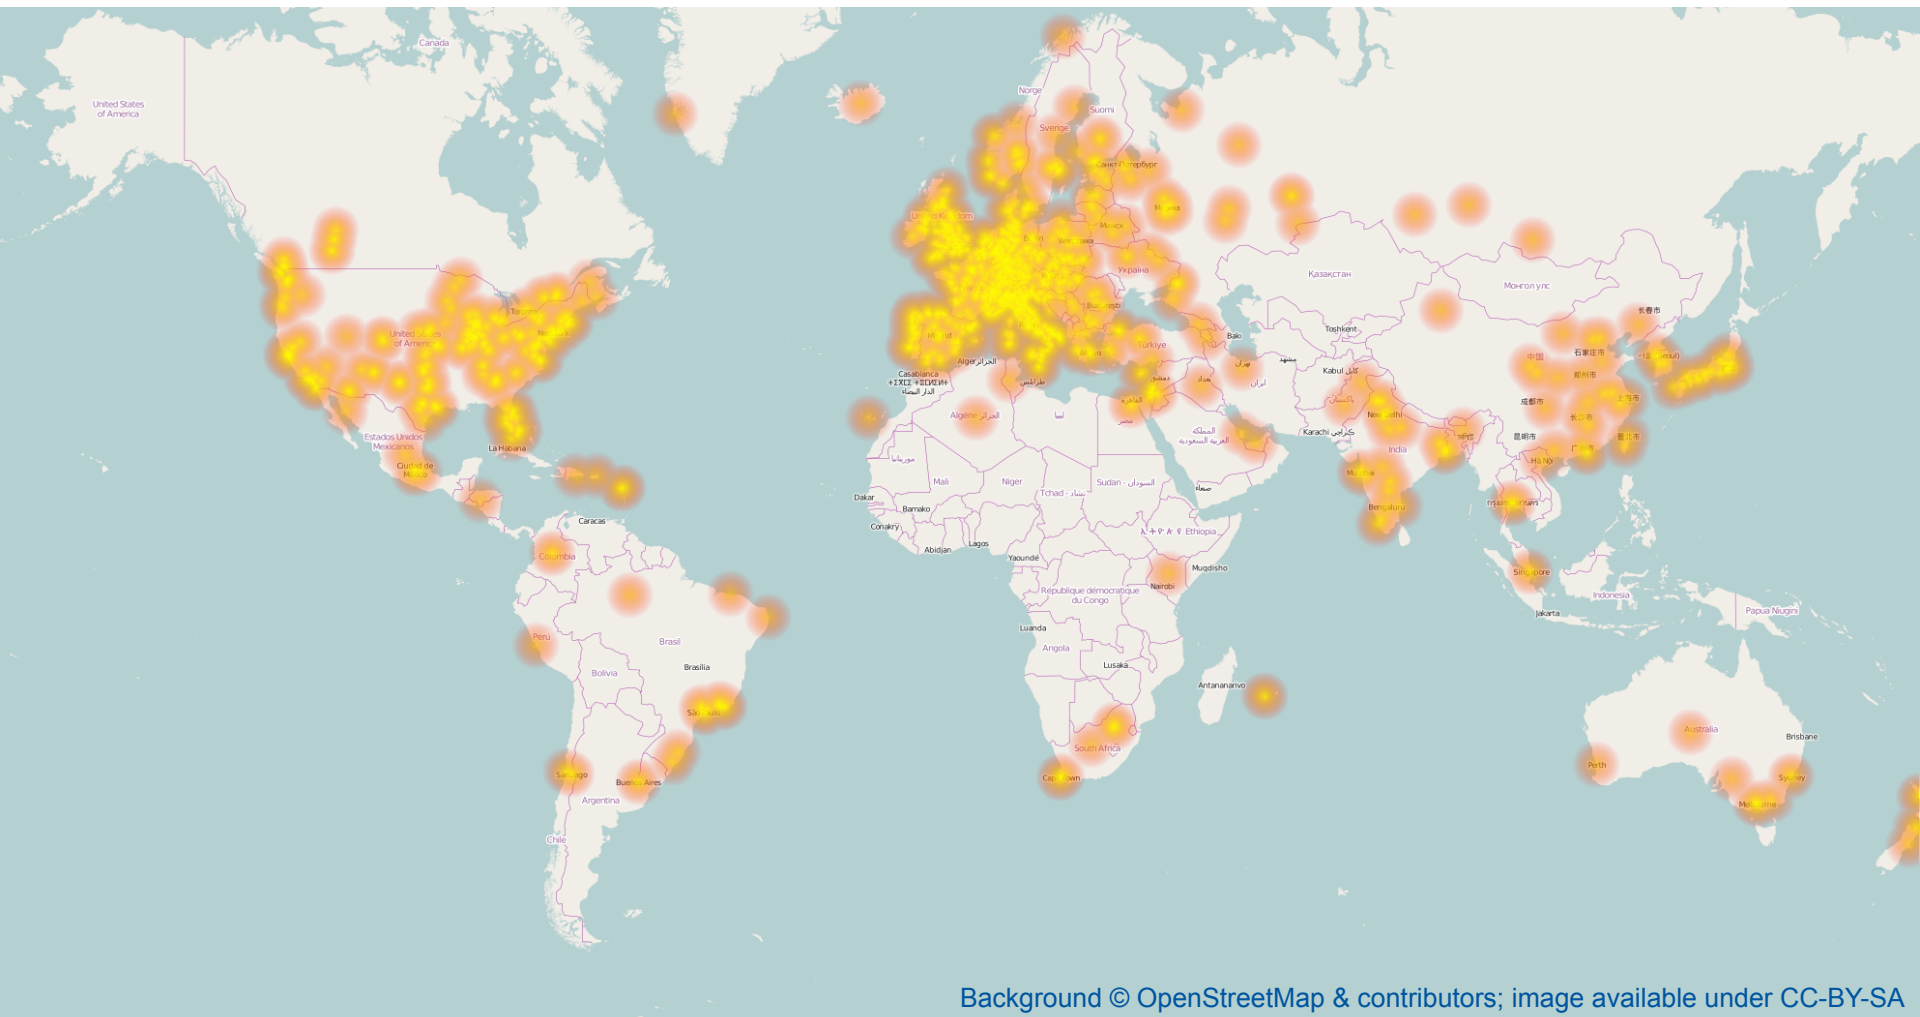


Nov 2015: Geolocation Active Users



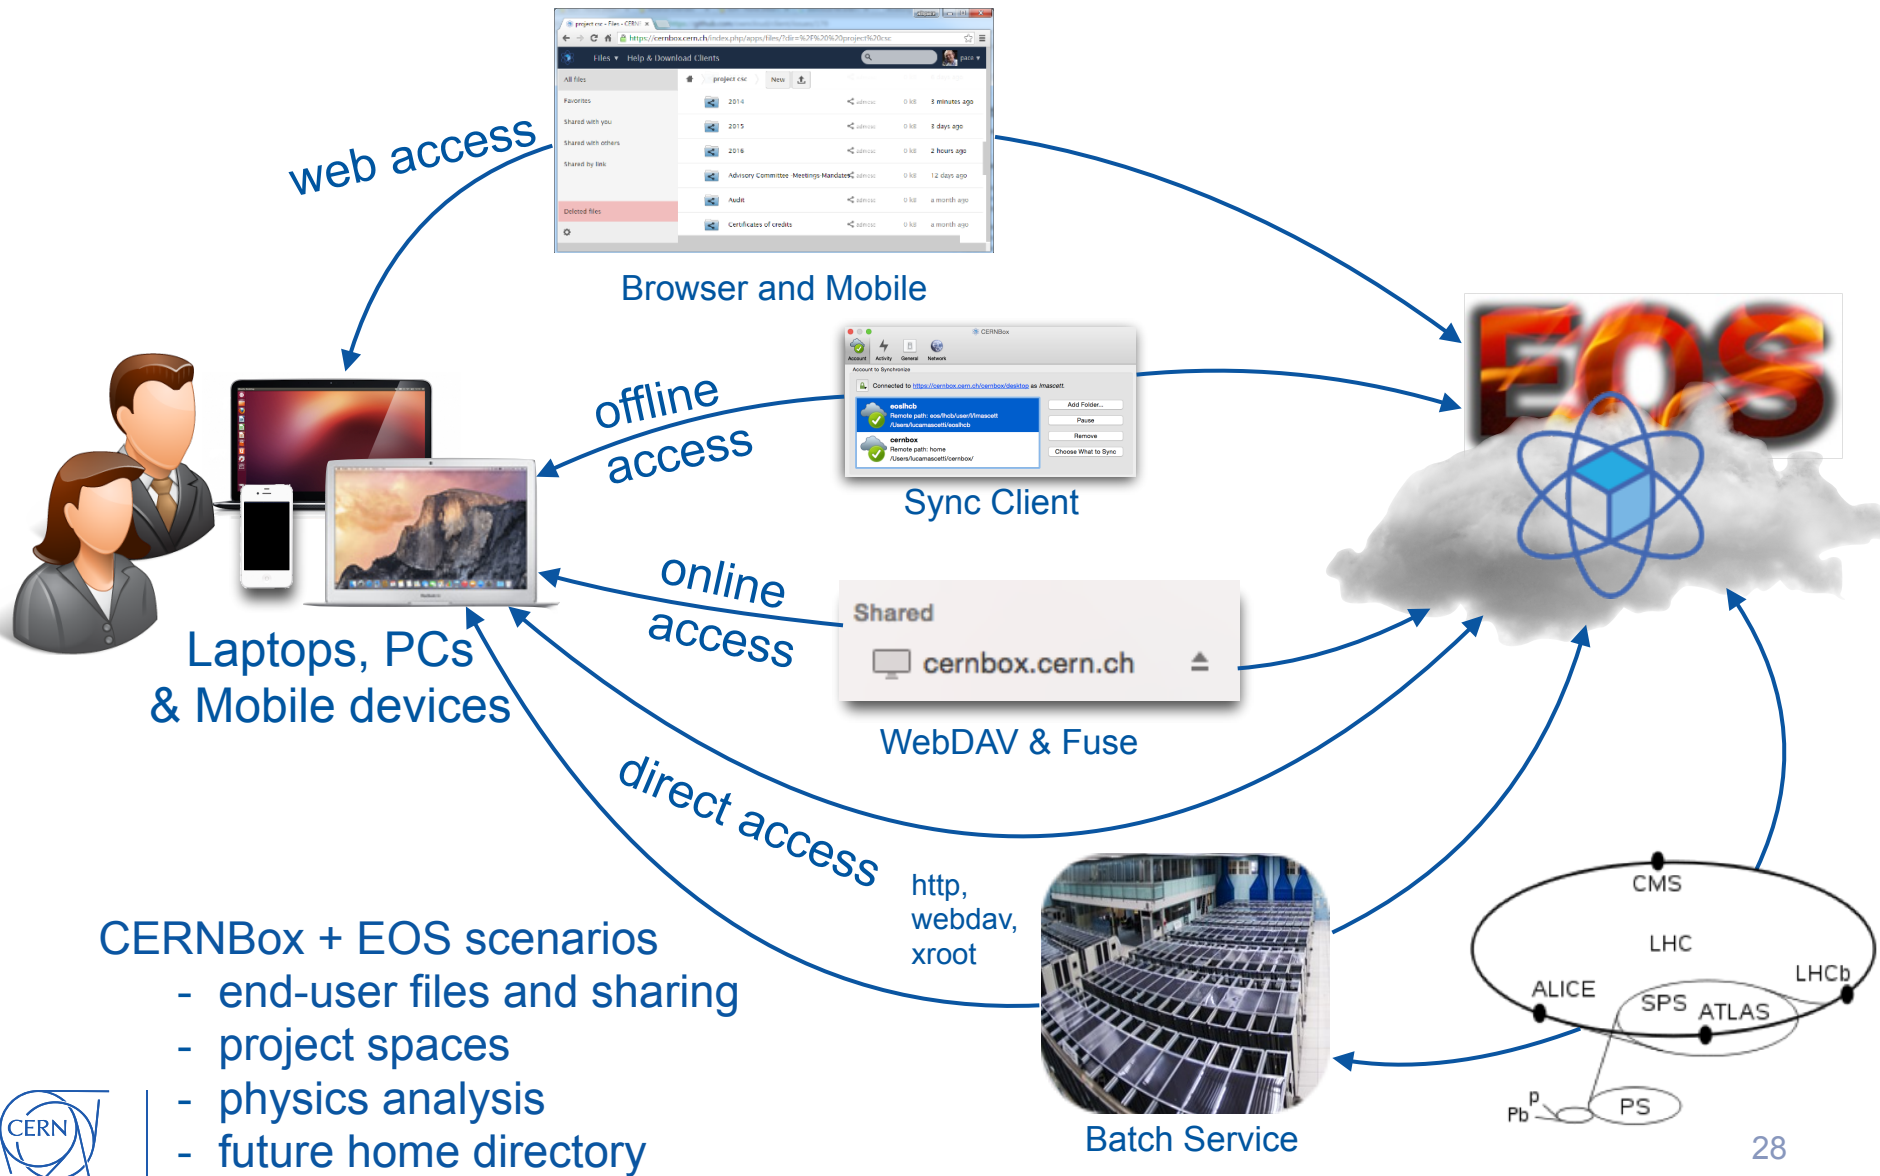
Background © OpenStreetMap & contributors; image available under CC-BY-SA

Dec 2015: Geolocation Active Users




Background © OpenStreetMap & contributors; image available under CC-BY-SA

Future Home Directory (\$HOME)



Summary and Outlook

- **EOS** provides a **very flexible** platform for a large community
 - integrated in Tier-0 workflow by ATLAS & CMS
 - more than 6k users storing data today
- Demonstrated unprecedented scalability
 - largest **low-cost** HEP storage installation site today
 - almost 200 PB and 50k disks
- **Strategic** direction for CERN based disk storage
 - for physics data (user/group/grid)
 - as ‘new-style’ home directory via **CERNBox**



www.cern.ch