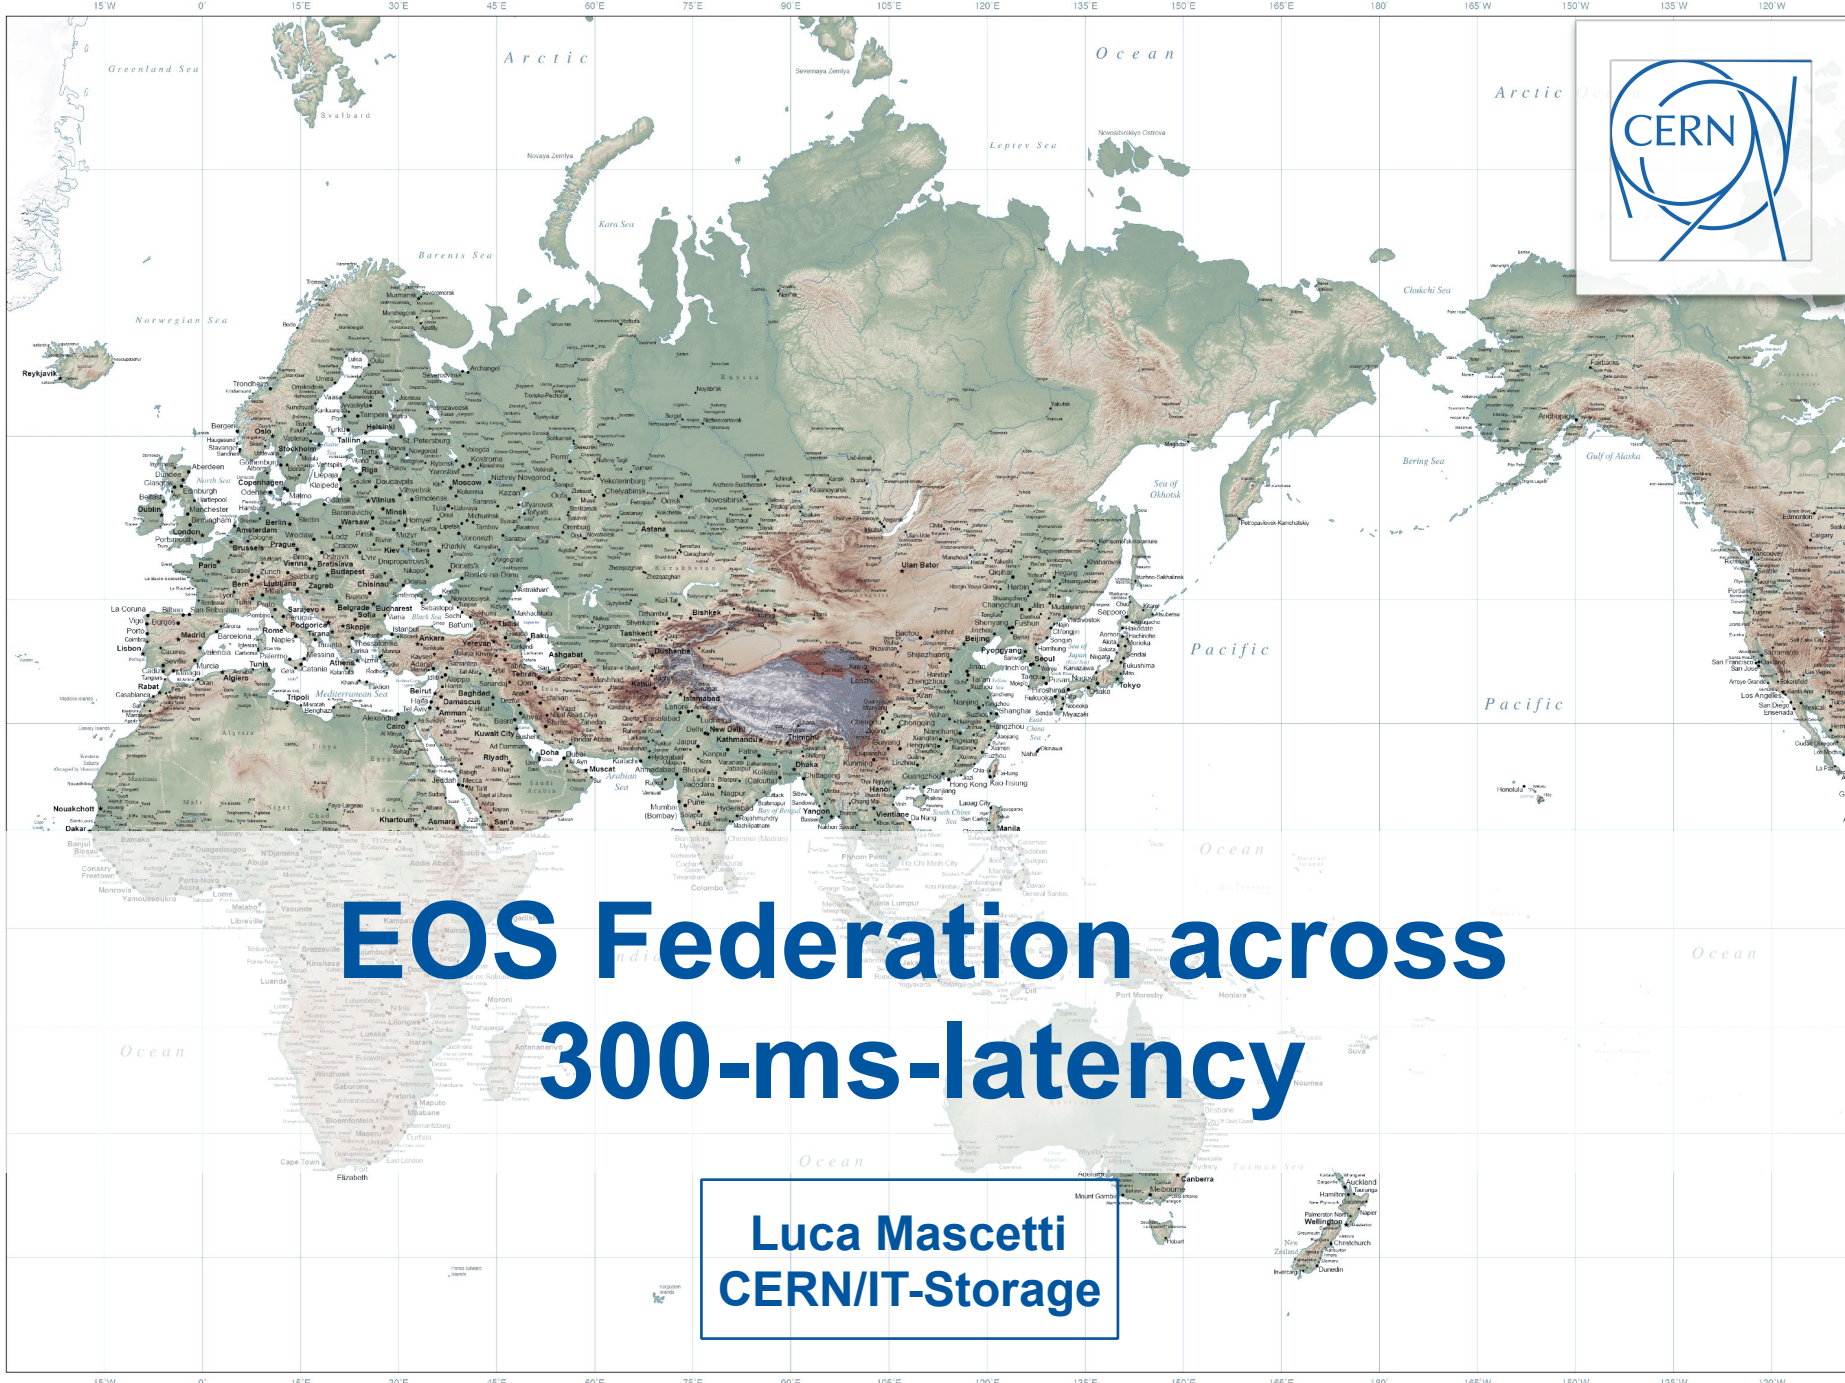






EOS Federation across 300-ms-latency


Luca Mascetti
CERN/IT-Storage




EOS Service @ CERN

Total Space
250 PB

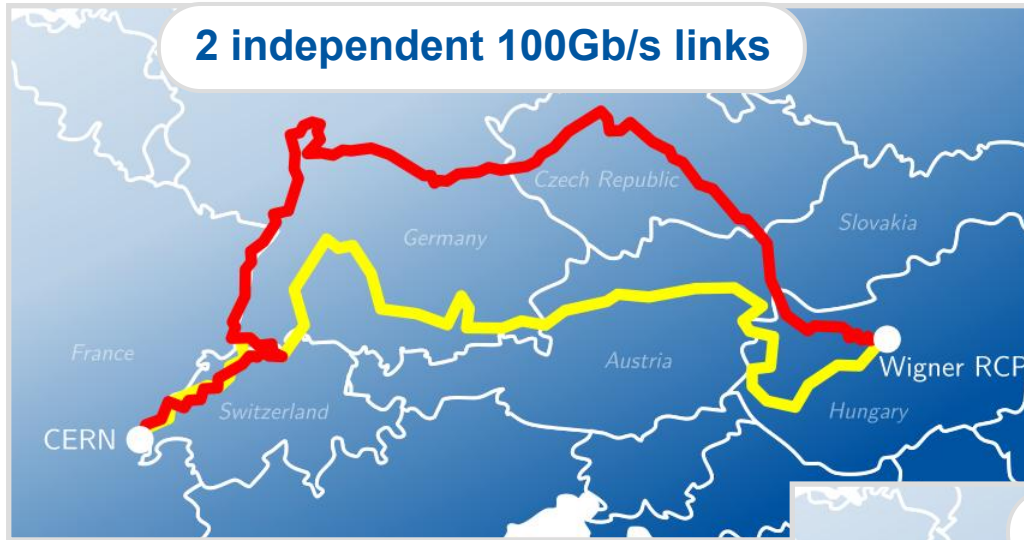
Files Stored
2.112 Bil



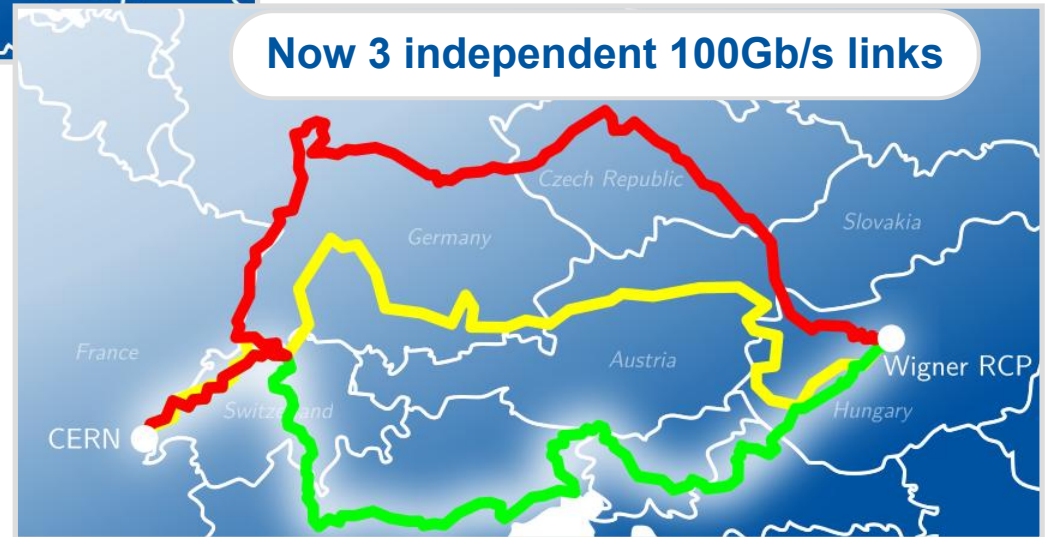
- Disk Storage for all LHC and physics data
 - and for CERNBox
- Deployment across two computer centres
 - CERN and Wigner RCP
 - Three 100Gb links



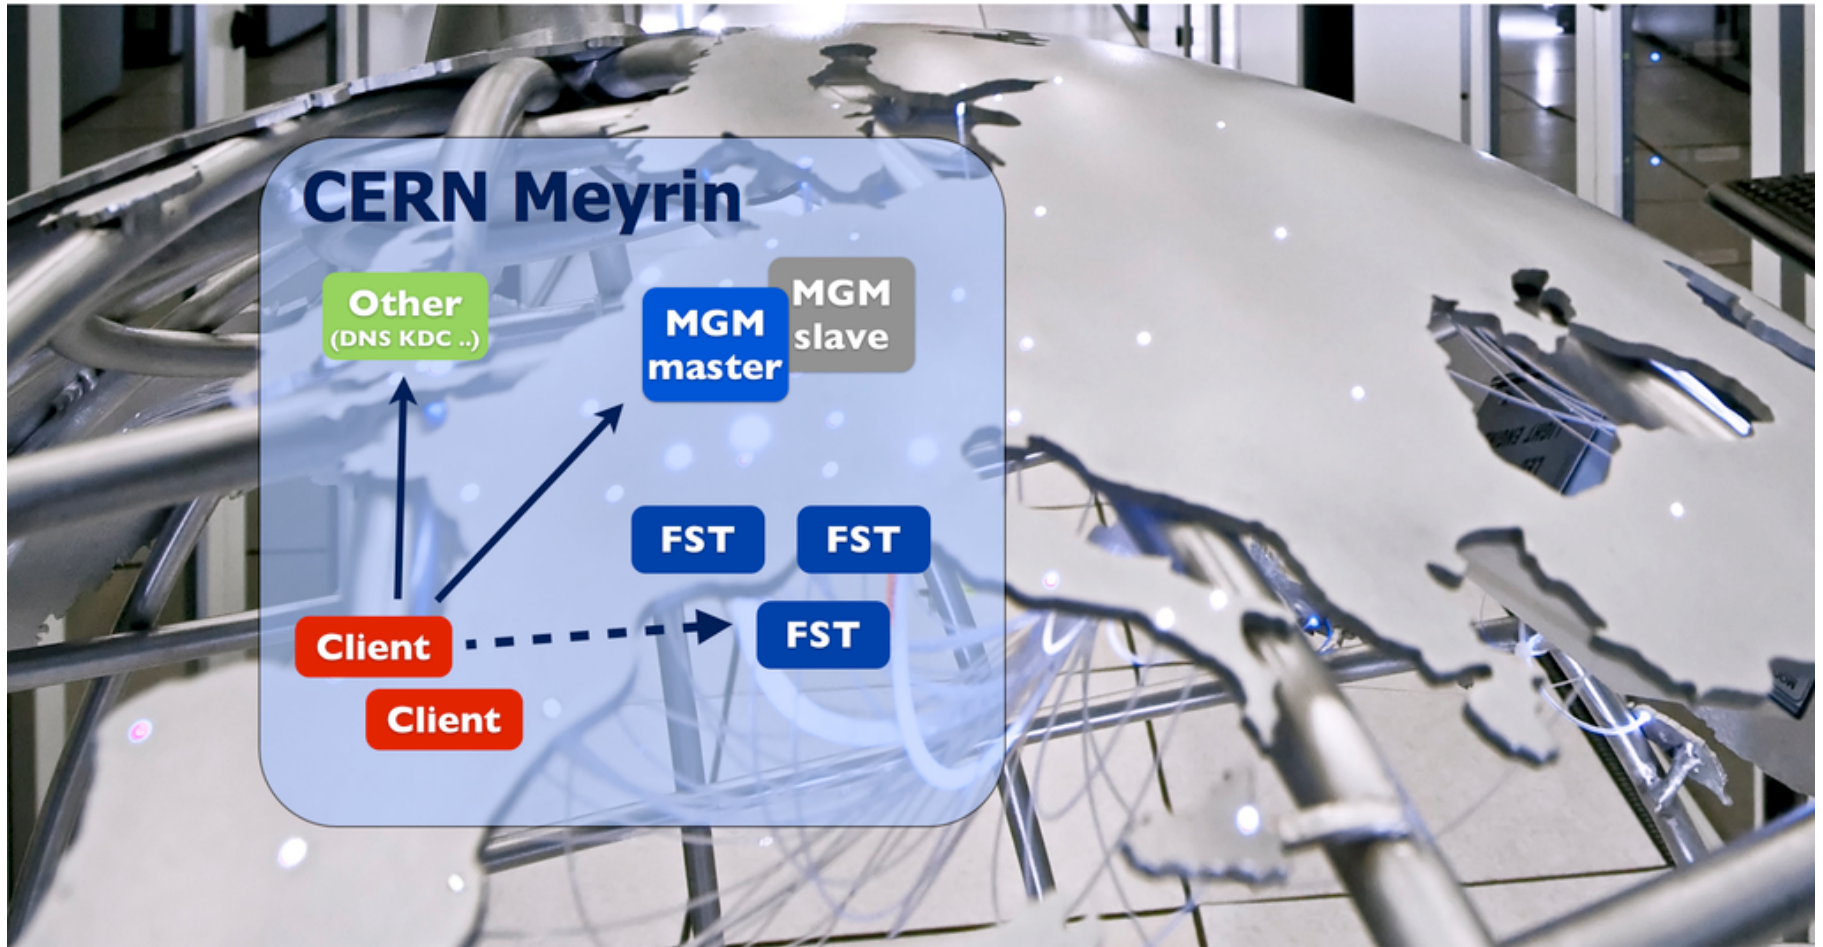
Wigner Computer Centre



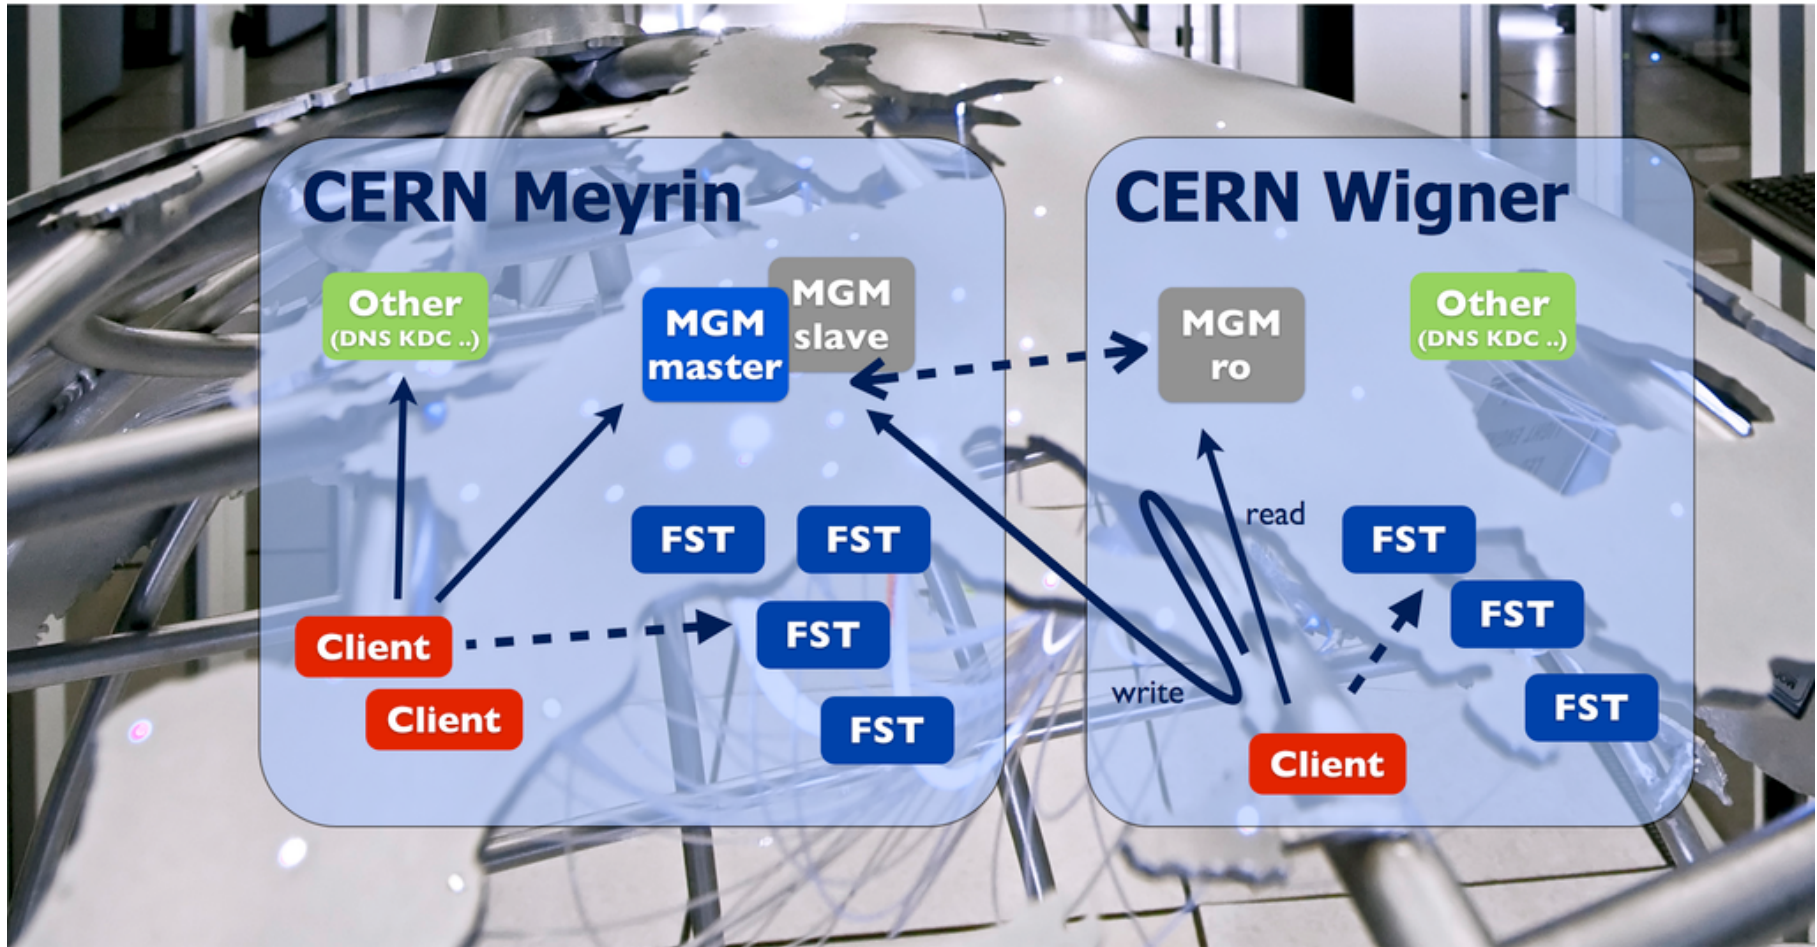
```
tes of data.  
icmp_seq=1 ttl=58 time=22.0 ms  
icmp_seq=2 ttl=58 time=22.1 ms  
icmp_seq=3 ttl=58 time=22.1 ms  
icmp_seq=4 ttl=58 time=22.1 ms  
icmp_seq=5 ttl=58 time=22.0 ms  
icmp_seq=6 ttl=58 time=22.0 ms  
icmp_seq=7 ttl=58 time=22.2 ms
```



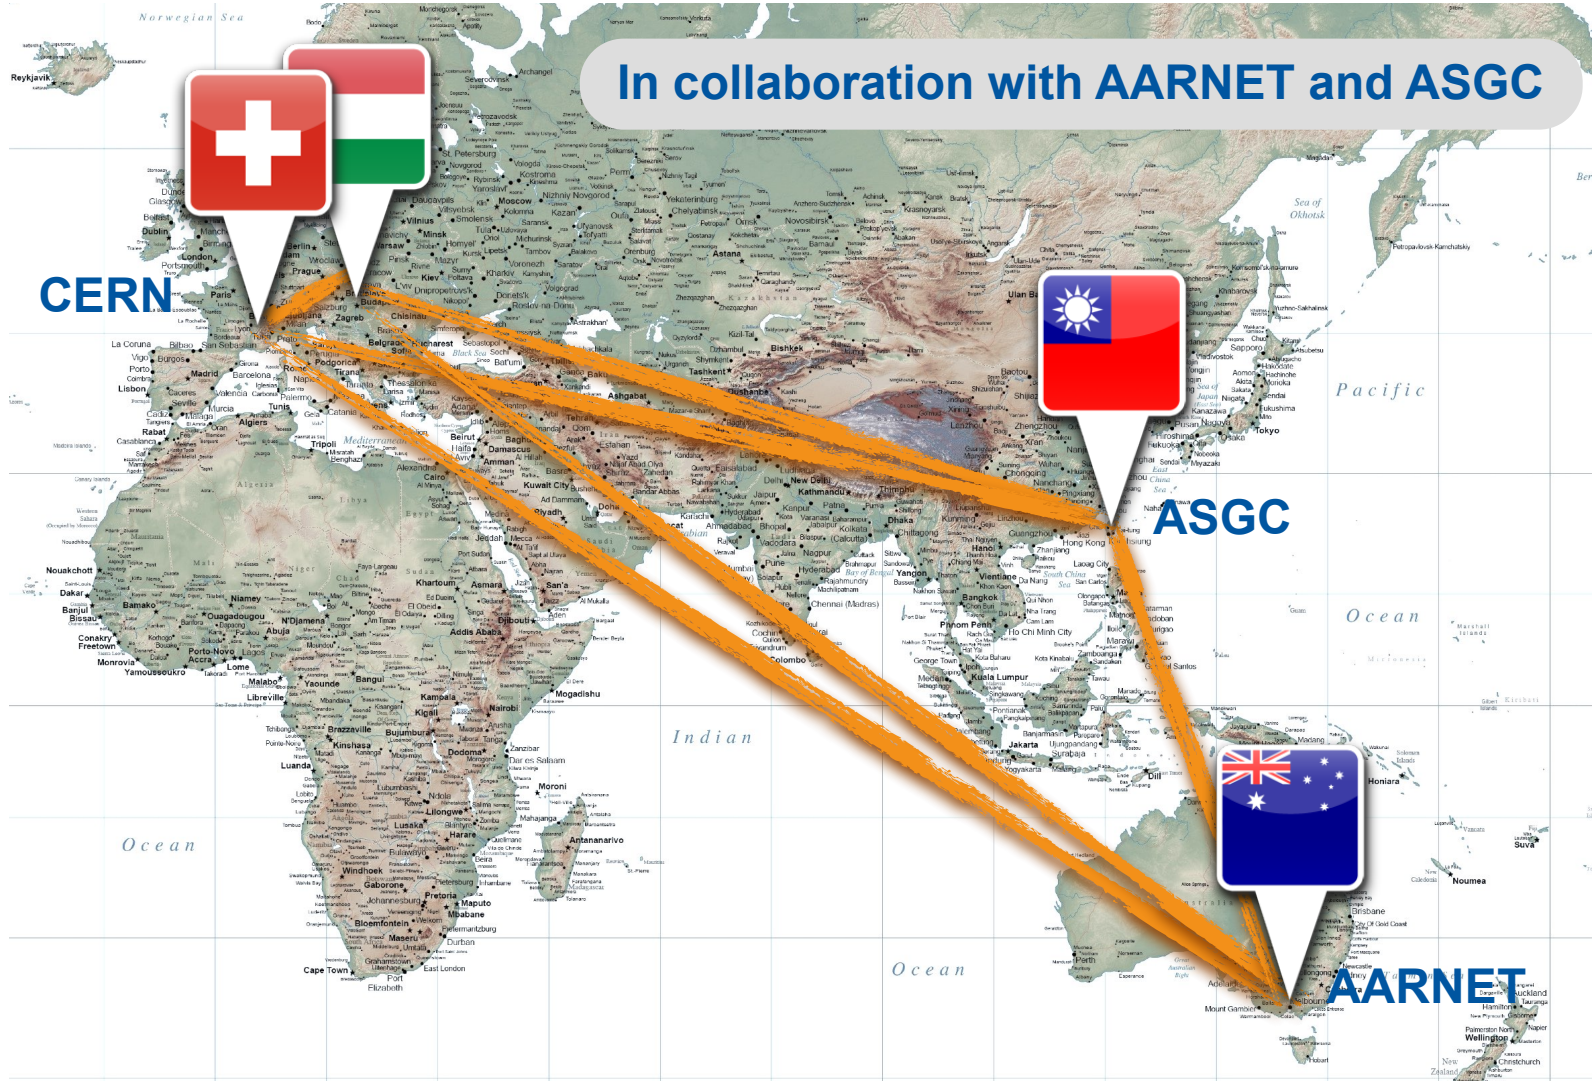
EOS 2011-2013 Deployment



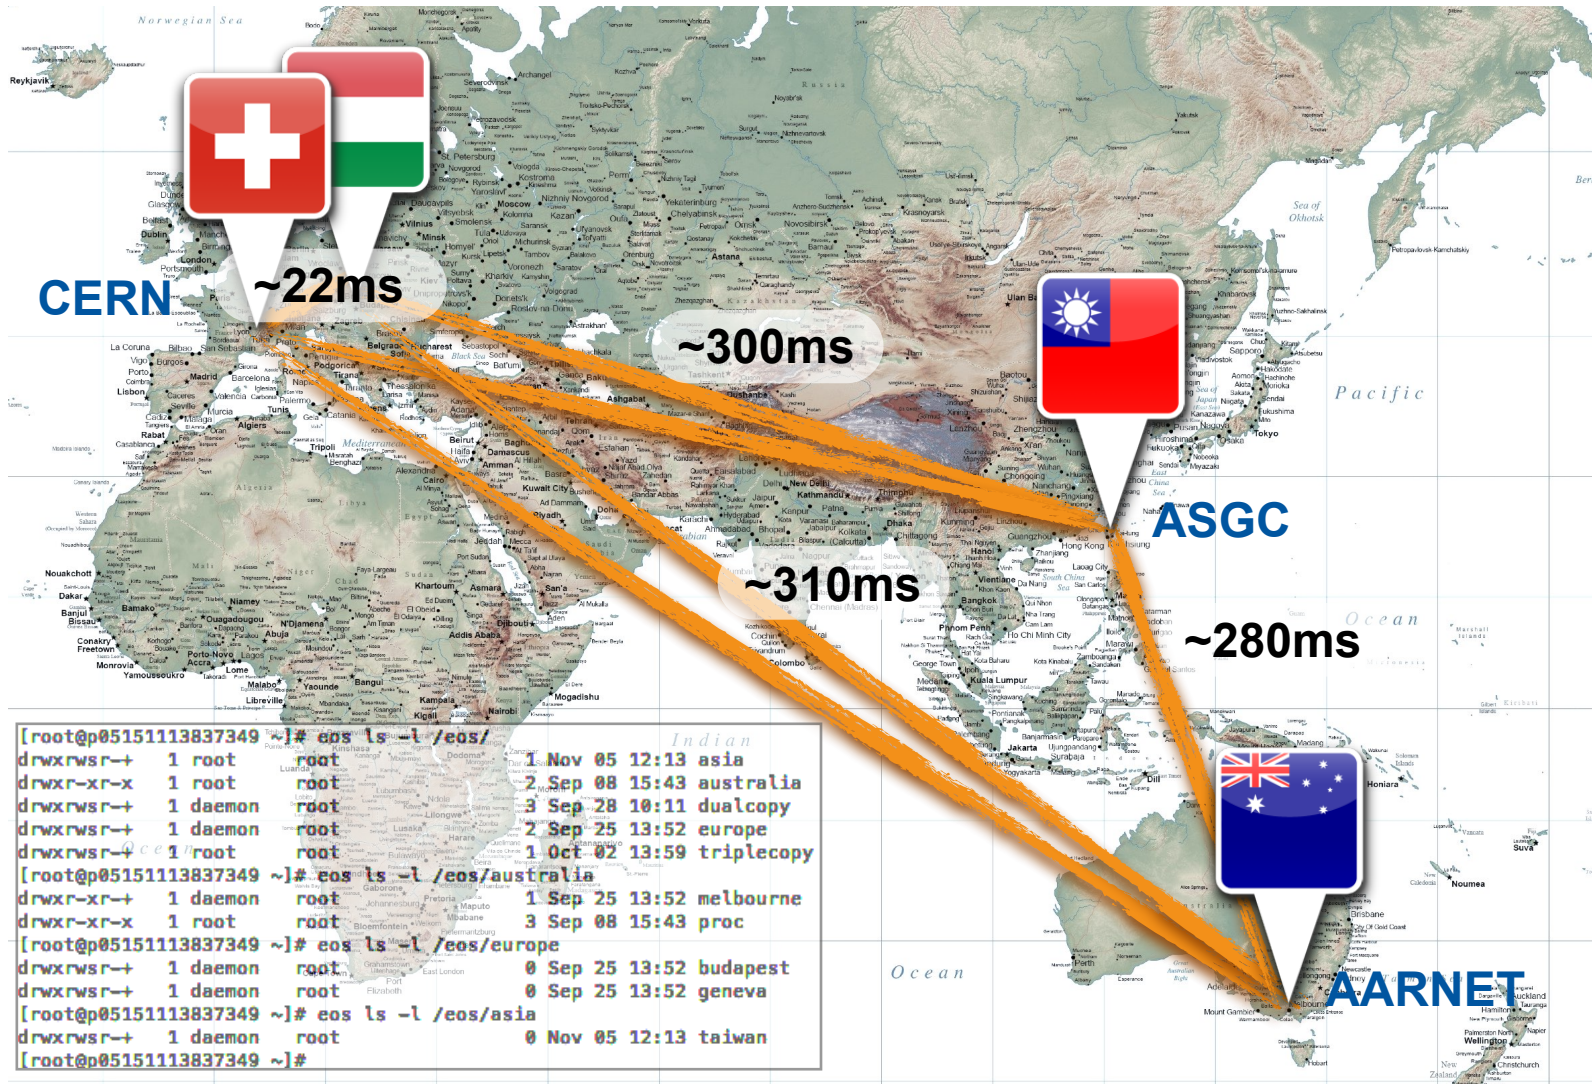
EOS 2014-2017 Deployment



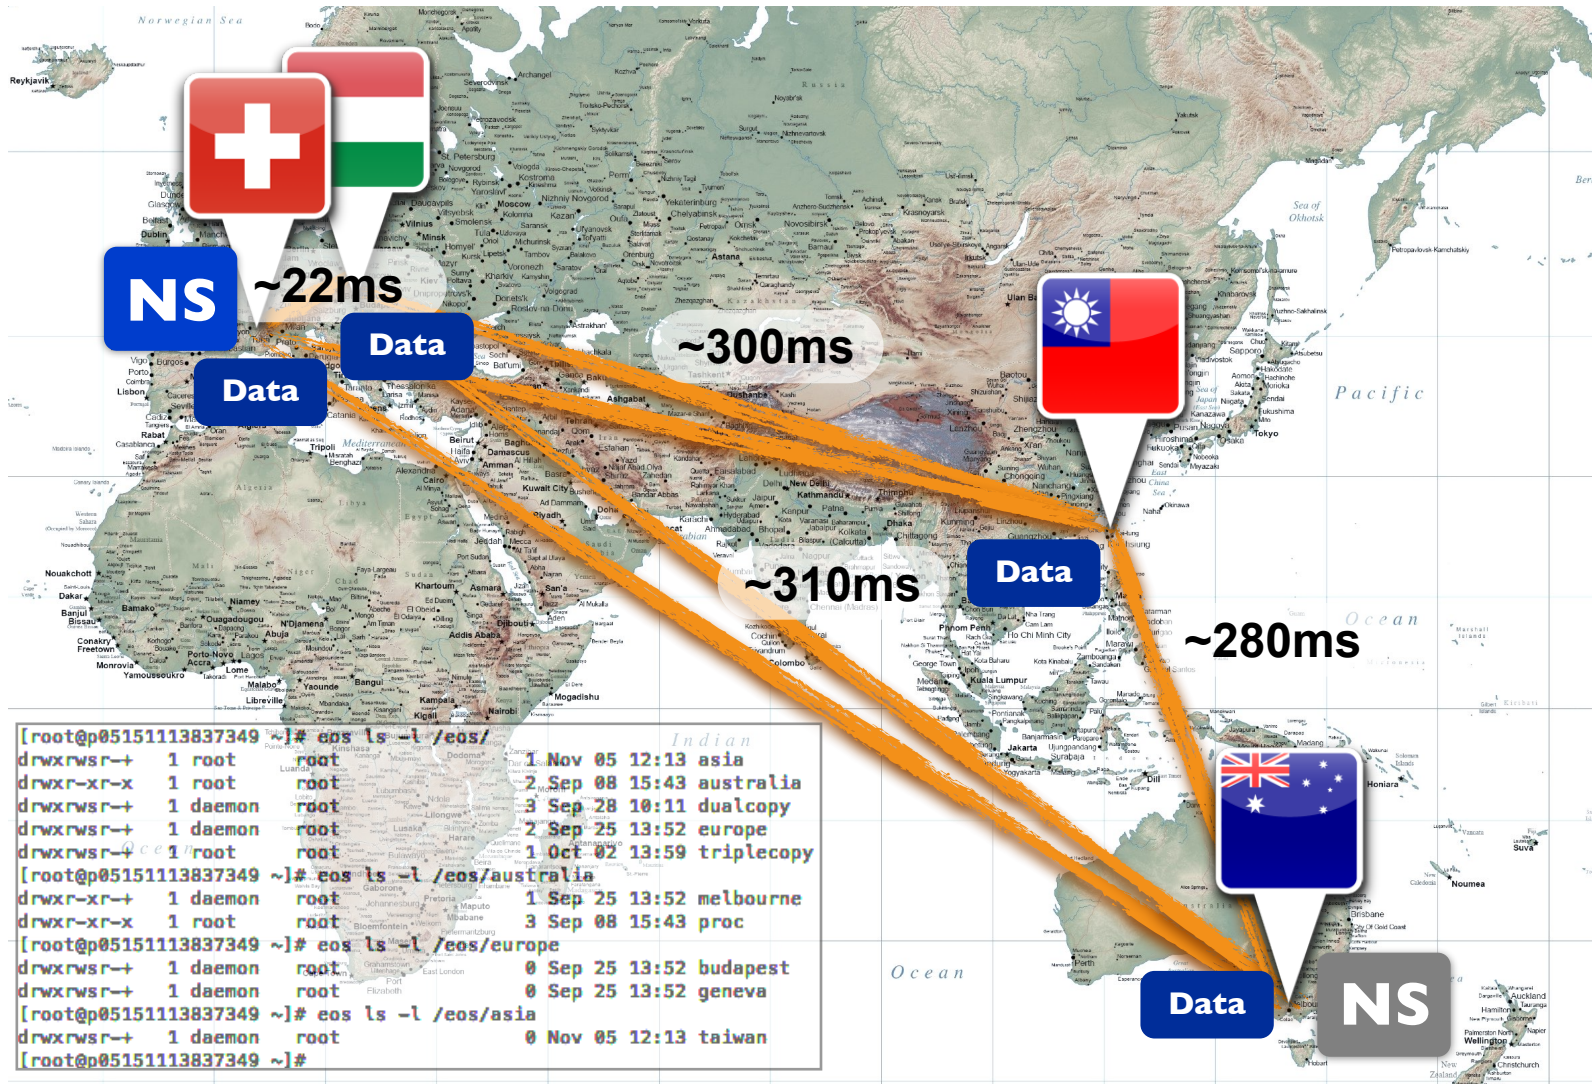
R&D - EOS World-Wide Deployment



R&D - EOS World-Wide Deployment



R&D - EOS World-Wide Deployment



```
[root@05151113837349 ~]# eos ls -l /eos/
drwxrwsr+ 1 root root 1 Nov 05 12:13 asia
drwxr-xr-x 1 root root 2 Sep 08 15:43 australia
drwxrwsr+ 1 daemon root 3 Sep 28 10:11 dualcopy
drwxrwsr+ 1 daemon root 2 Sep 25 13:52 europe
drwxrwsr+ 1 root root 1 Oct 02 13:59 triplecopy
[root@05151113837349 ~]# eos ls -l /eos/australia
drwxr-xr+ 1 daemon root 1 Sep 25 13:52 melbourne
drwxr-xr-x 1 root root 3 Sep 08 15:43 proc
[root@05151113837349 ~]# eos ls -l /eos/europe
drwxrwsr+ 1 daemon root 0 Sep 25 13:52 budapest
drwxrwsr+ 1 daemon root 0 Sep 25 13:52 geneva
[root@05151113837349 ~]# eos ls -l /eos/asia
drwxrwsr+ 1 daemon root 0 Nov 05 12:13 taiwan
[root@05151113837349 ~]#
```



1-GEOTAG Pools

/eos

/asia

/taiwan

/australia

/melbourne

/europe

/geneva

/budapest

/dualcopy

/gva-bud



/mel-gva

/mel-bud

/triplecopy

/mel-gva-bud

/mel-gva-tpe



2-GEOTAG Pools

/eos

/asia

/taiwan

/australia

/melbourne

/europe

/geneva

/budapest

/dualcopy

/gva-bud




/mel-gva

/mel-bud

/triplecopy

/mel-gva-bud

/mel-gva-tpe



3-GEOTAG Pools

/eos

/asia

/taiwan

/australia

/melbourne

/europe

/geneva

/budapest

/dualcopy

/gva-bud



/mel-gva

/mel-bud

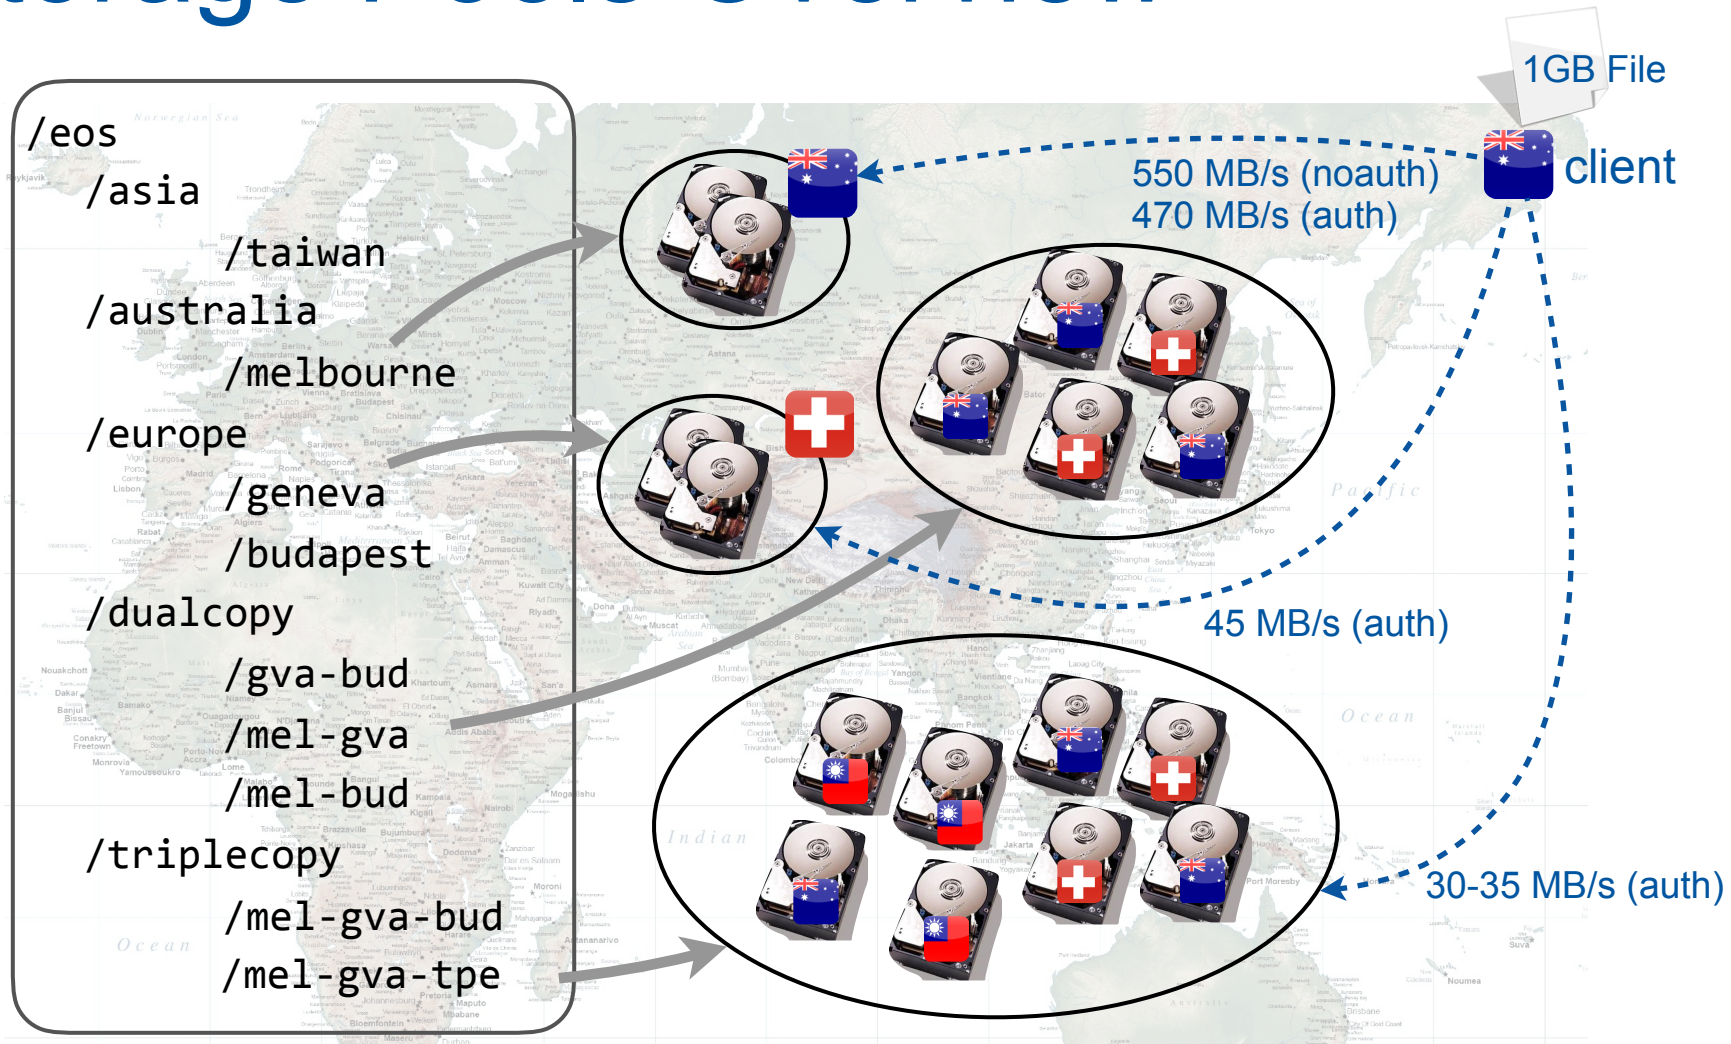
/triplecopy

/mel-gva-bud

/mel-gva-tpe

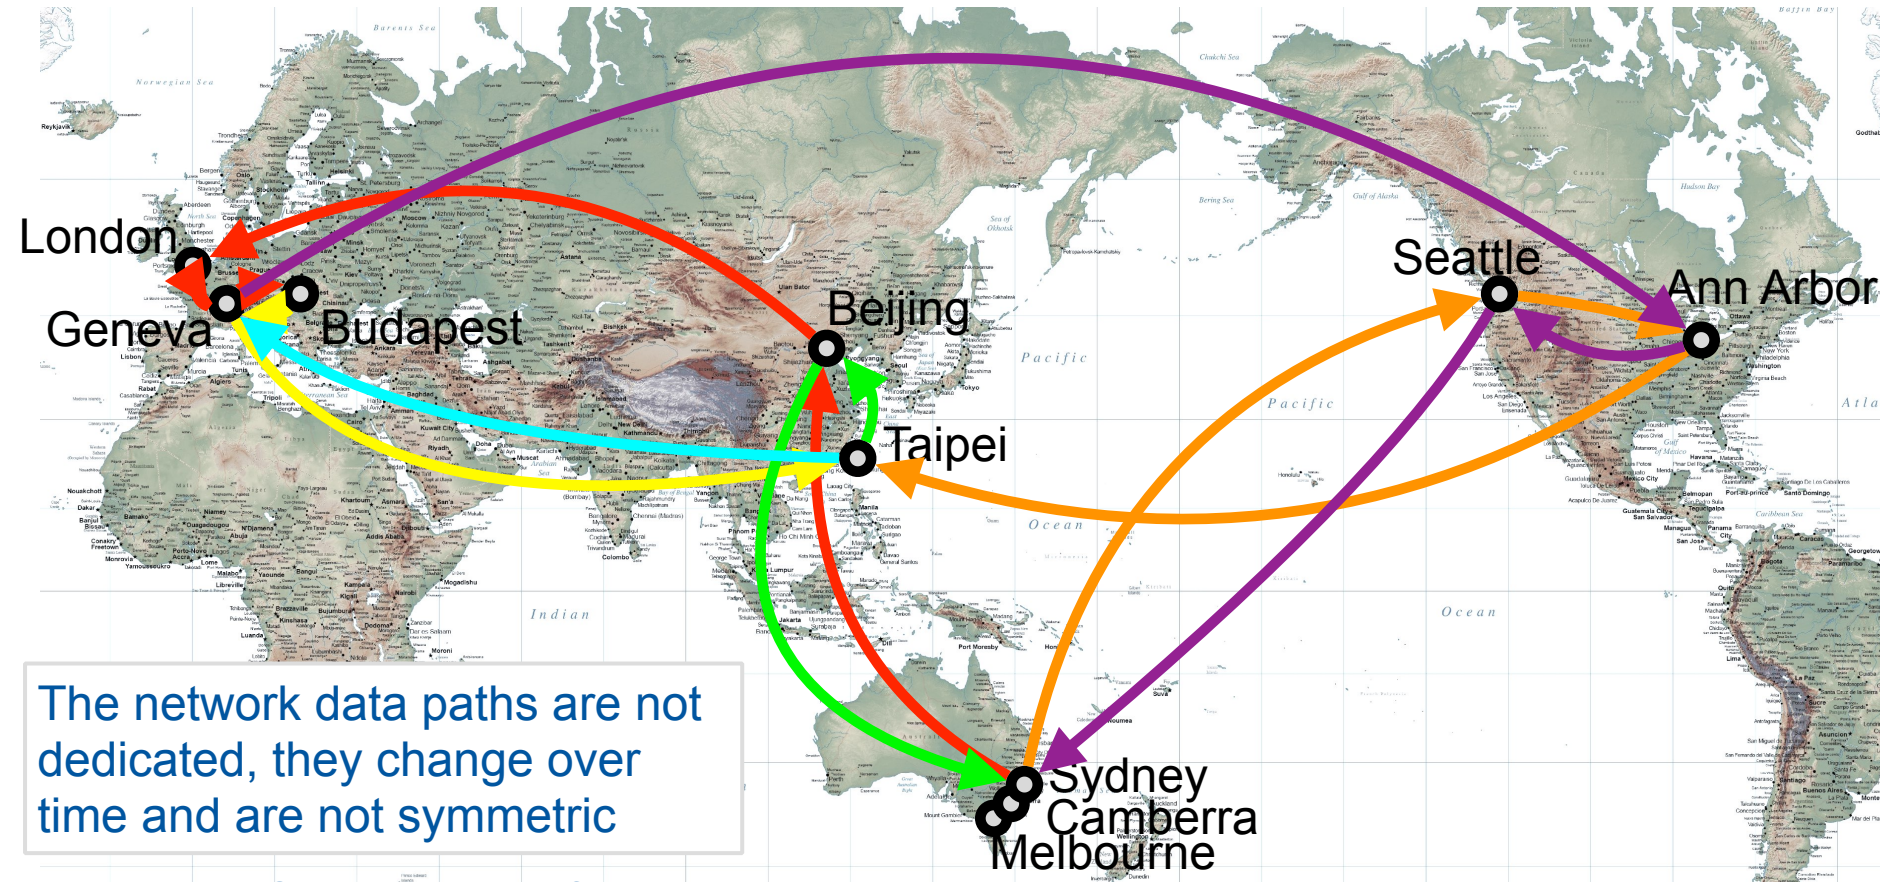


Storage Pools Overview



Storage pools were created with filesystems from all four sites. Files were replicated according to the different configured policy (e.g. 3 replicas: MEL-GVA-TPE).

Network Topology



- 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

AARNET

Make data access easy
Make Analysis simple
Facilitate Science

- Scale-out filesystem underneath the ownCloud app, using the eosd fuse interface for file IO

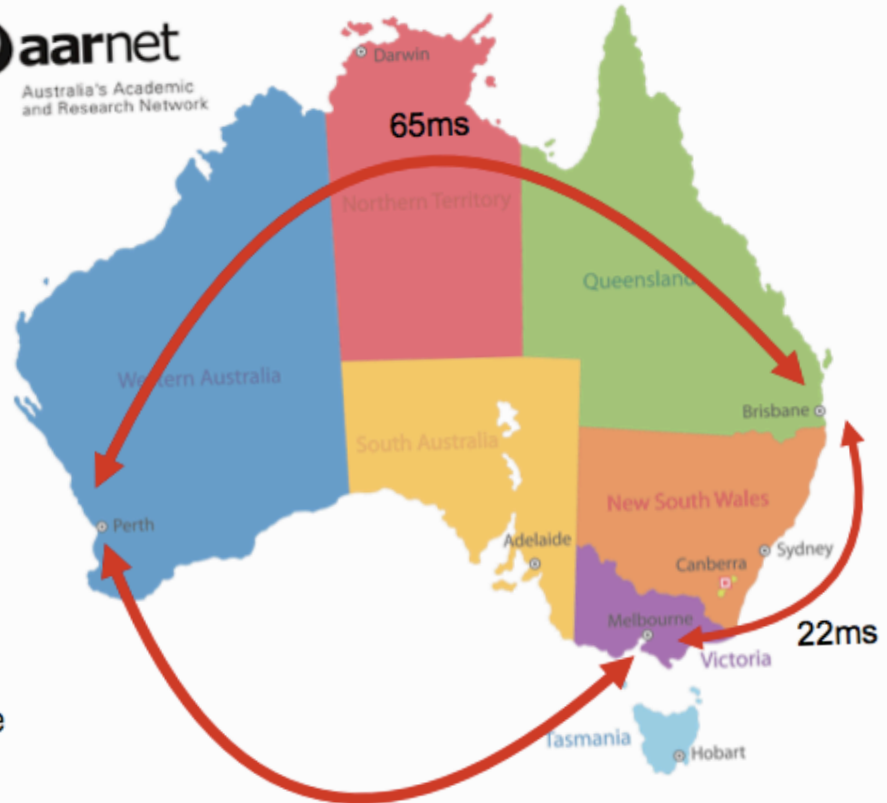


- Geo-distributed setup: Brisbane, Melbourne, Perth
 - ~1PB (scale to ~20PB next year)



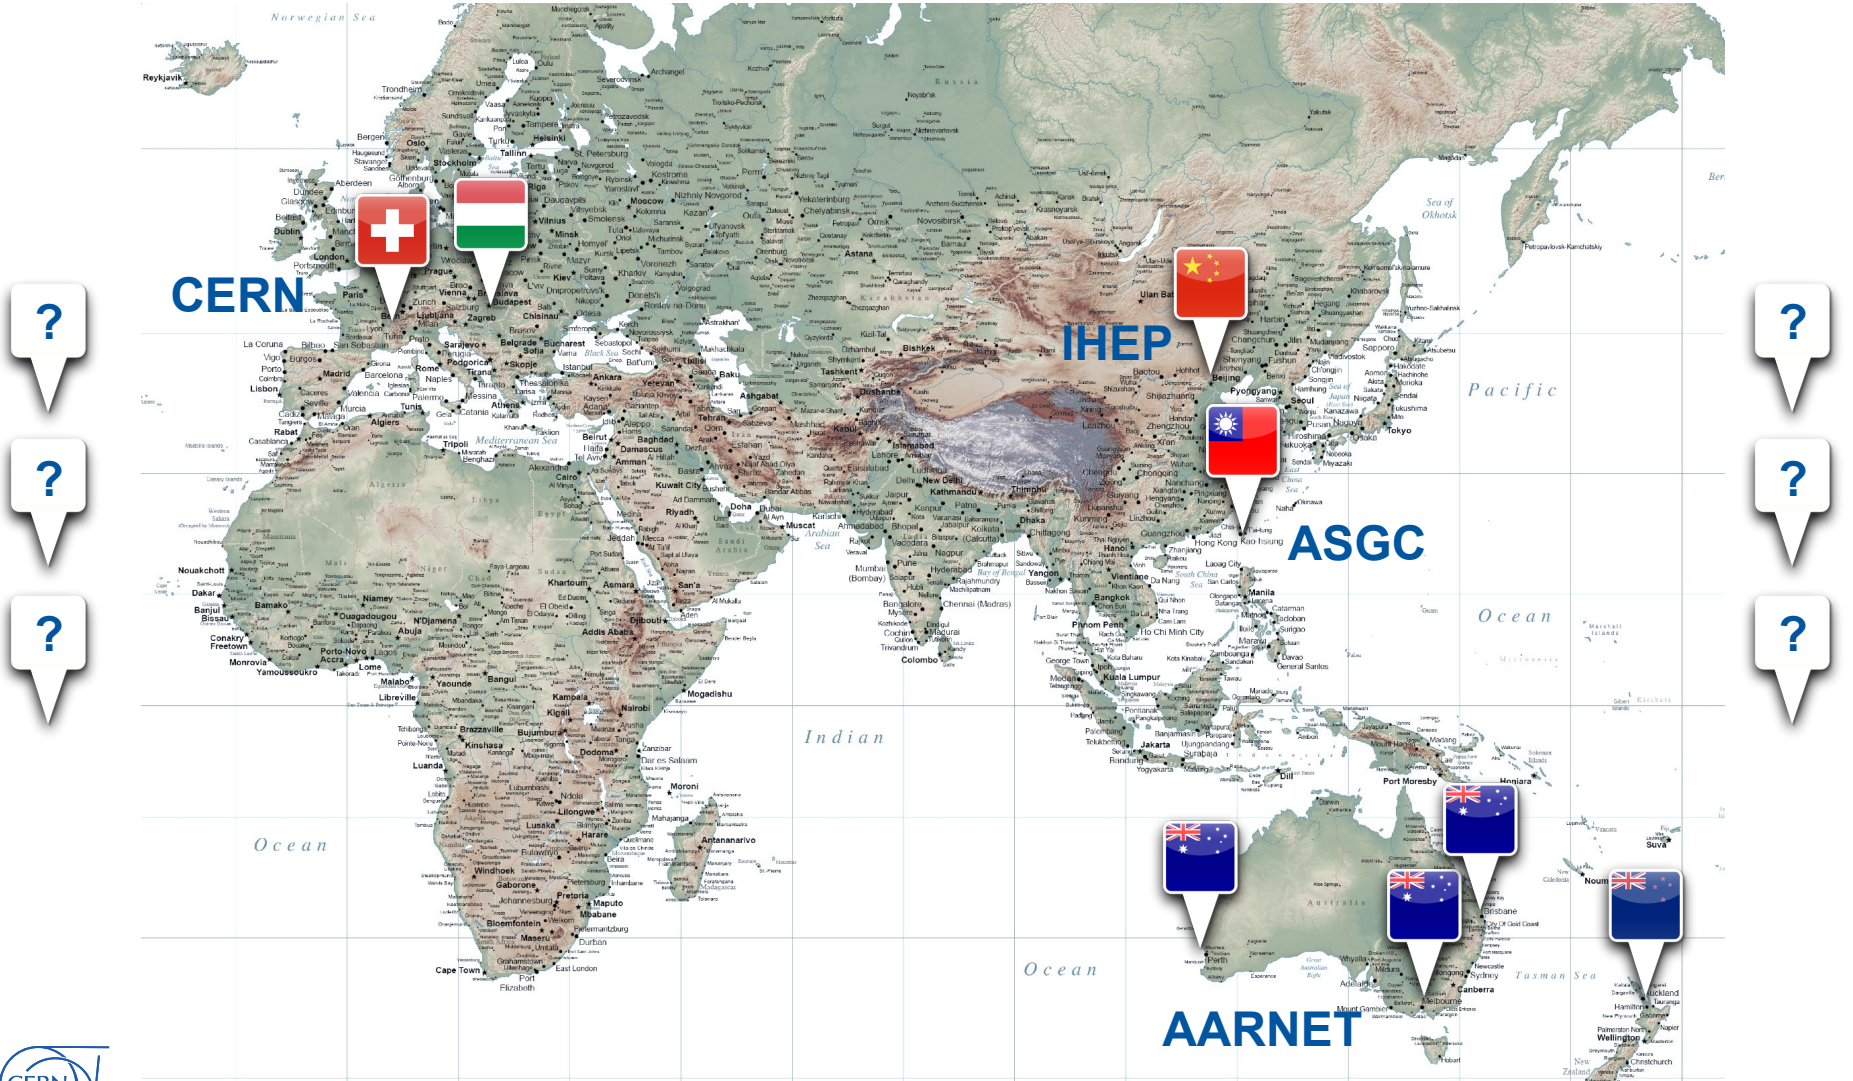
Australian National University

- Australian National University, in Acton Canberra: mirror archives of both genome sequences and open or freely available software distributed among three sites




"This system is presently running 0.3.187, and has been so trouble free that I keep forgetting it's there." David Jericho -AARNetSolutions Architect

R&D - EOS World Wide Deployment - Future



Summary

- EOS provides a very flexible storage platform for a large community
- Demonstrated unprecedented scalability
 - largest low-cost HEP storage installation site today
- EOS confirm its capability in handling multiple sites
 - even with very high latencies.
- Performance using native clients were mainly dictated by the connectivity status of the sites.
- Authentication (for read or for write) affected by the latency to contact the closer namespace.



www.cern.ch