
CERN

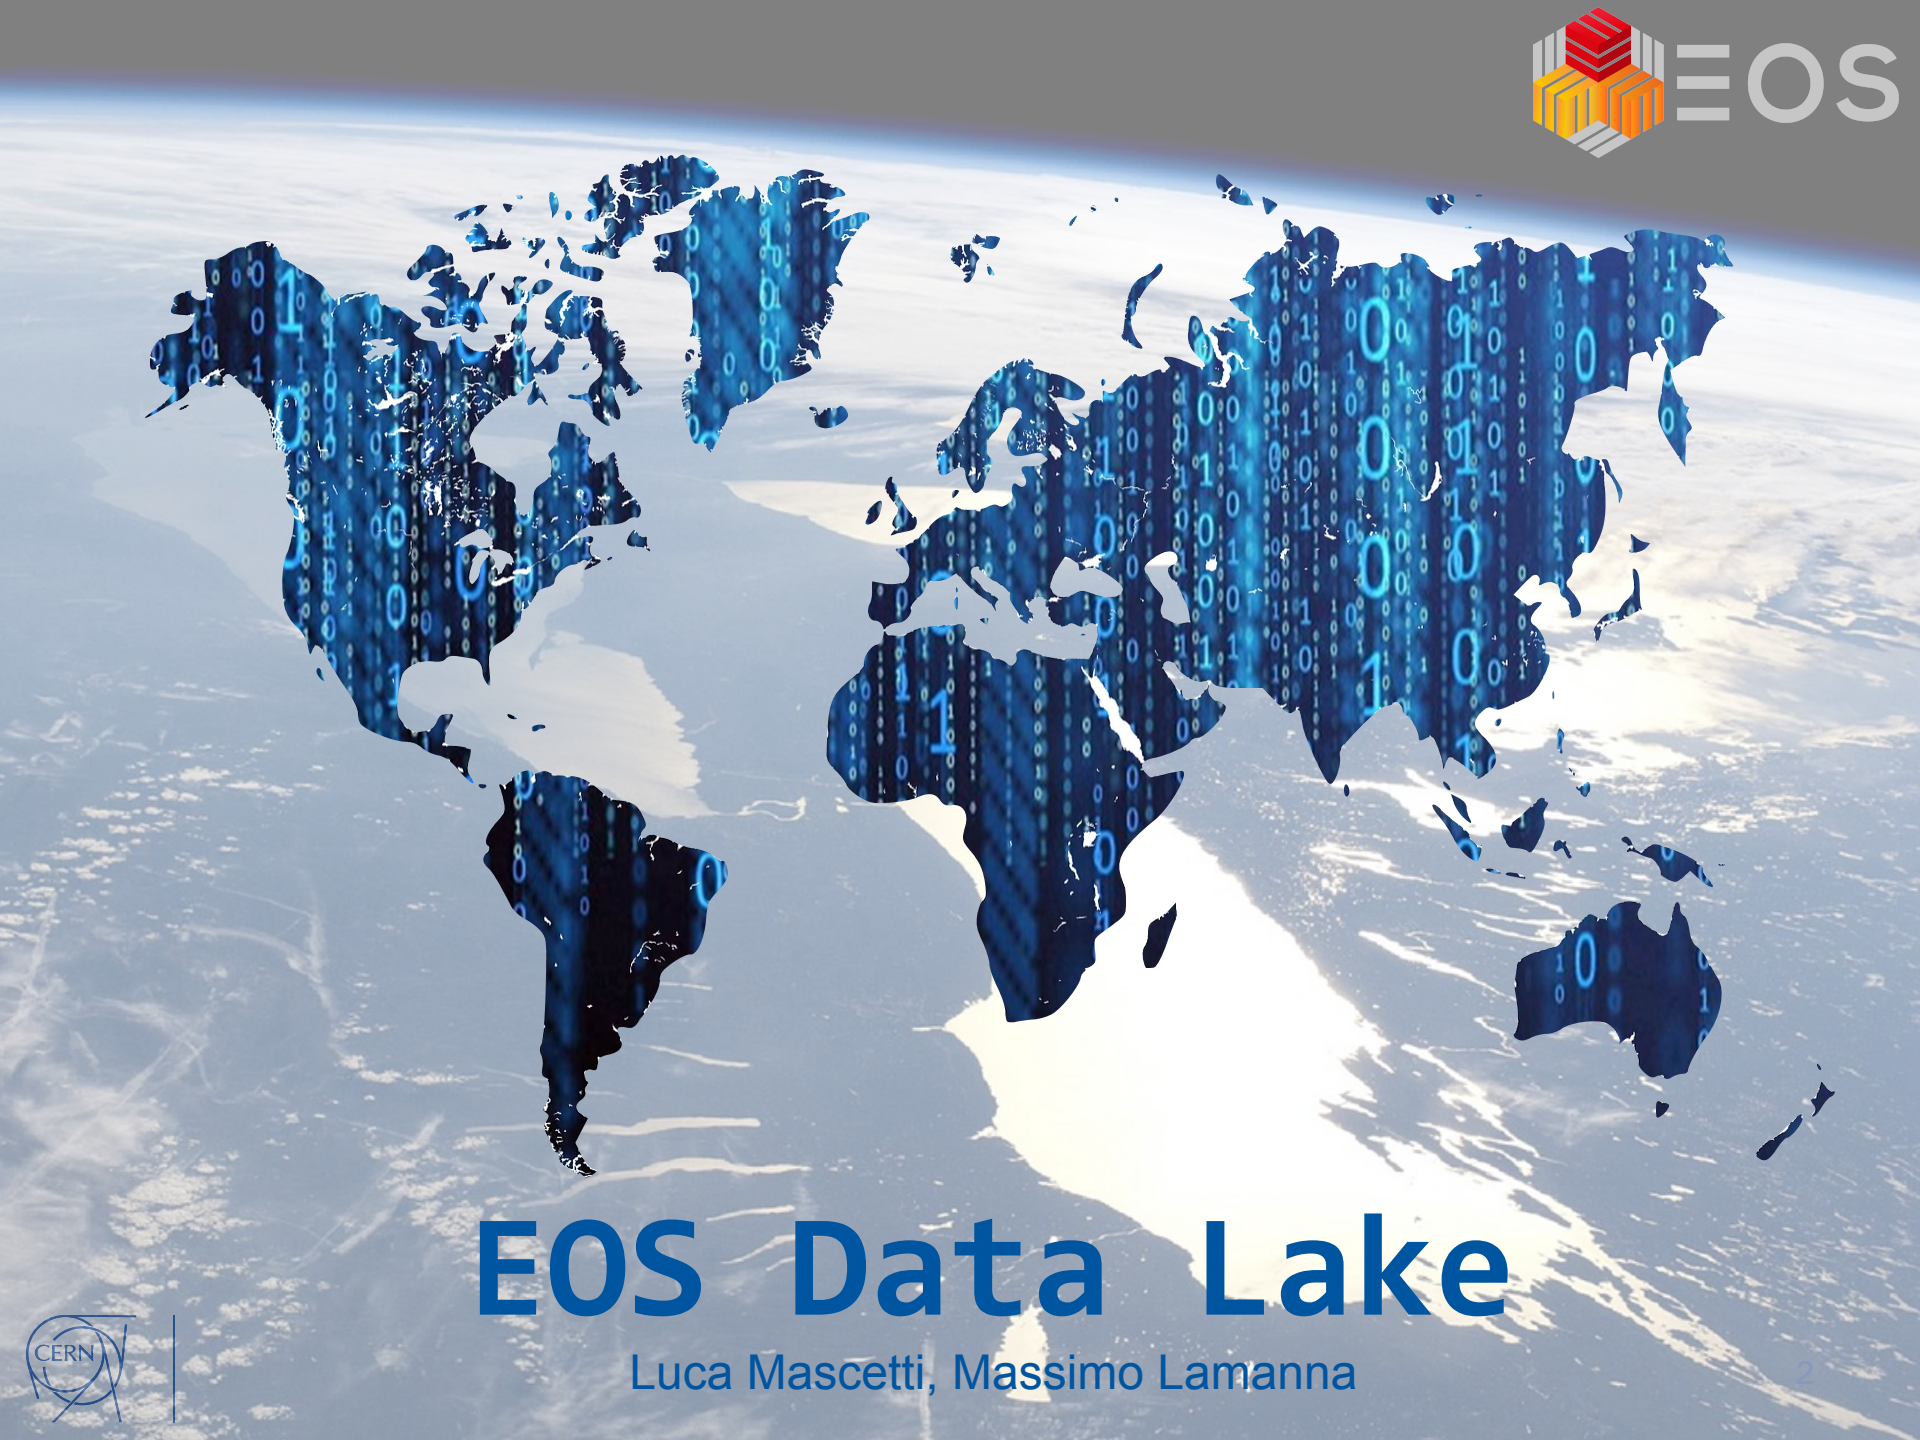

EOS Data Lake

Luca Mascetti, Massimo Lamanna

"WLCG" DATA LAKE

- n sites
- k replicas
- k<<n
- latency > 1ms

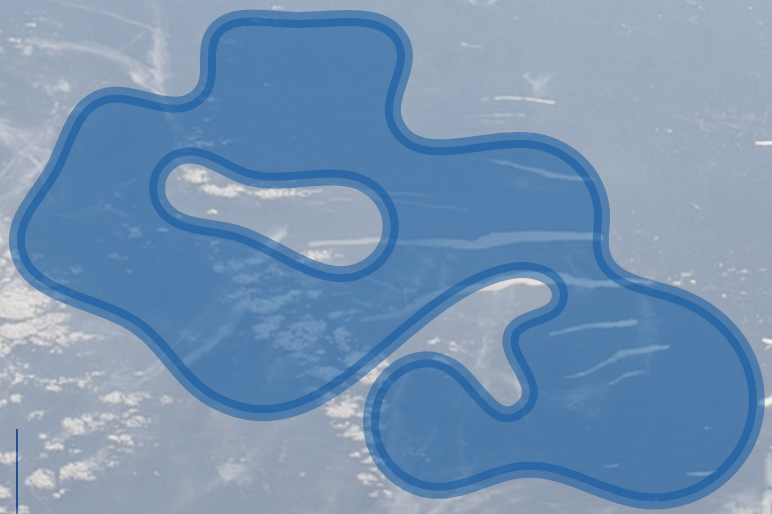

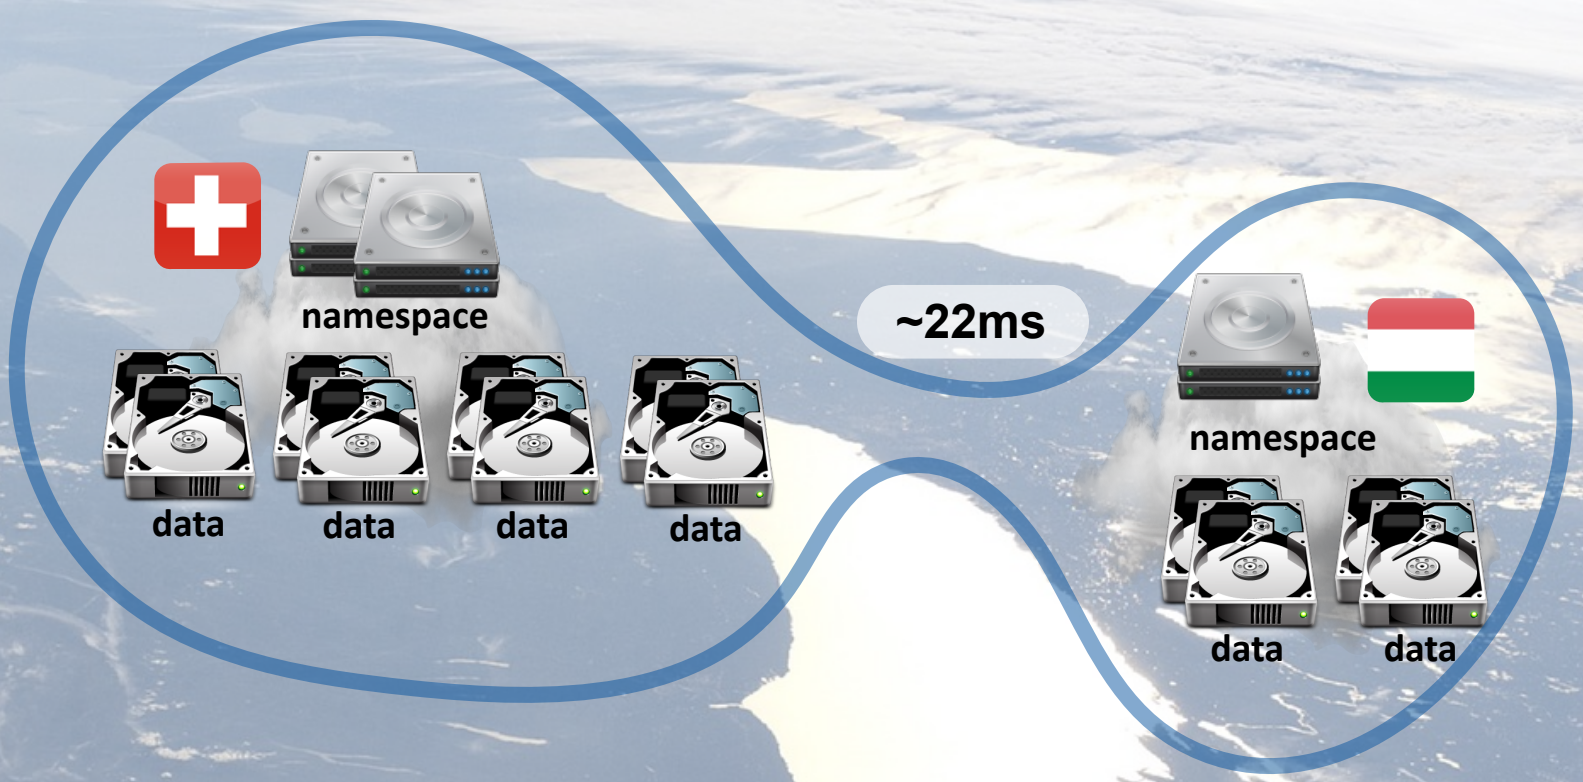
namespace
data
data
data
data
~22ms
namespace
data
data

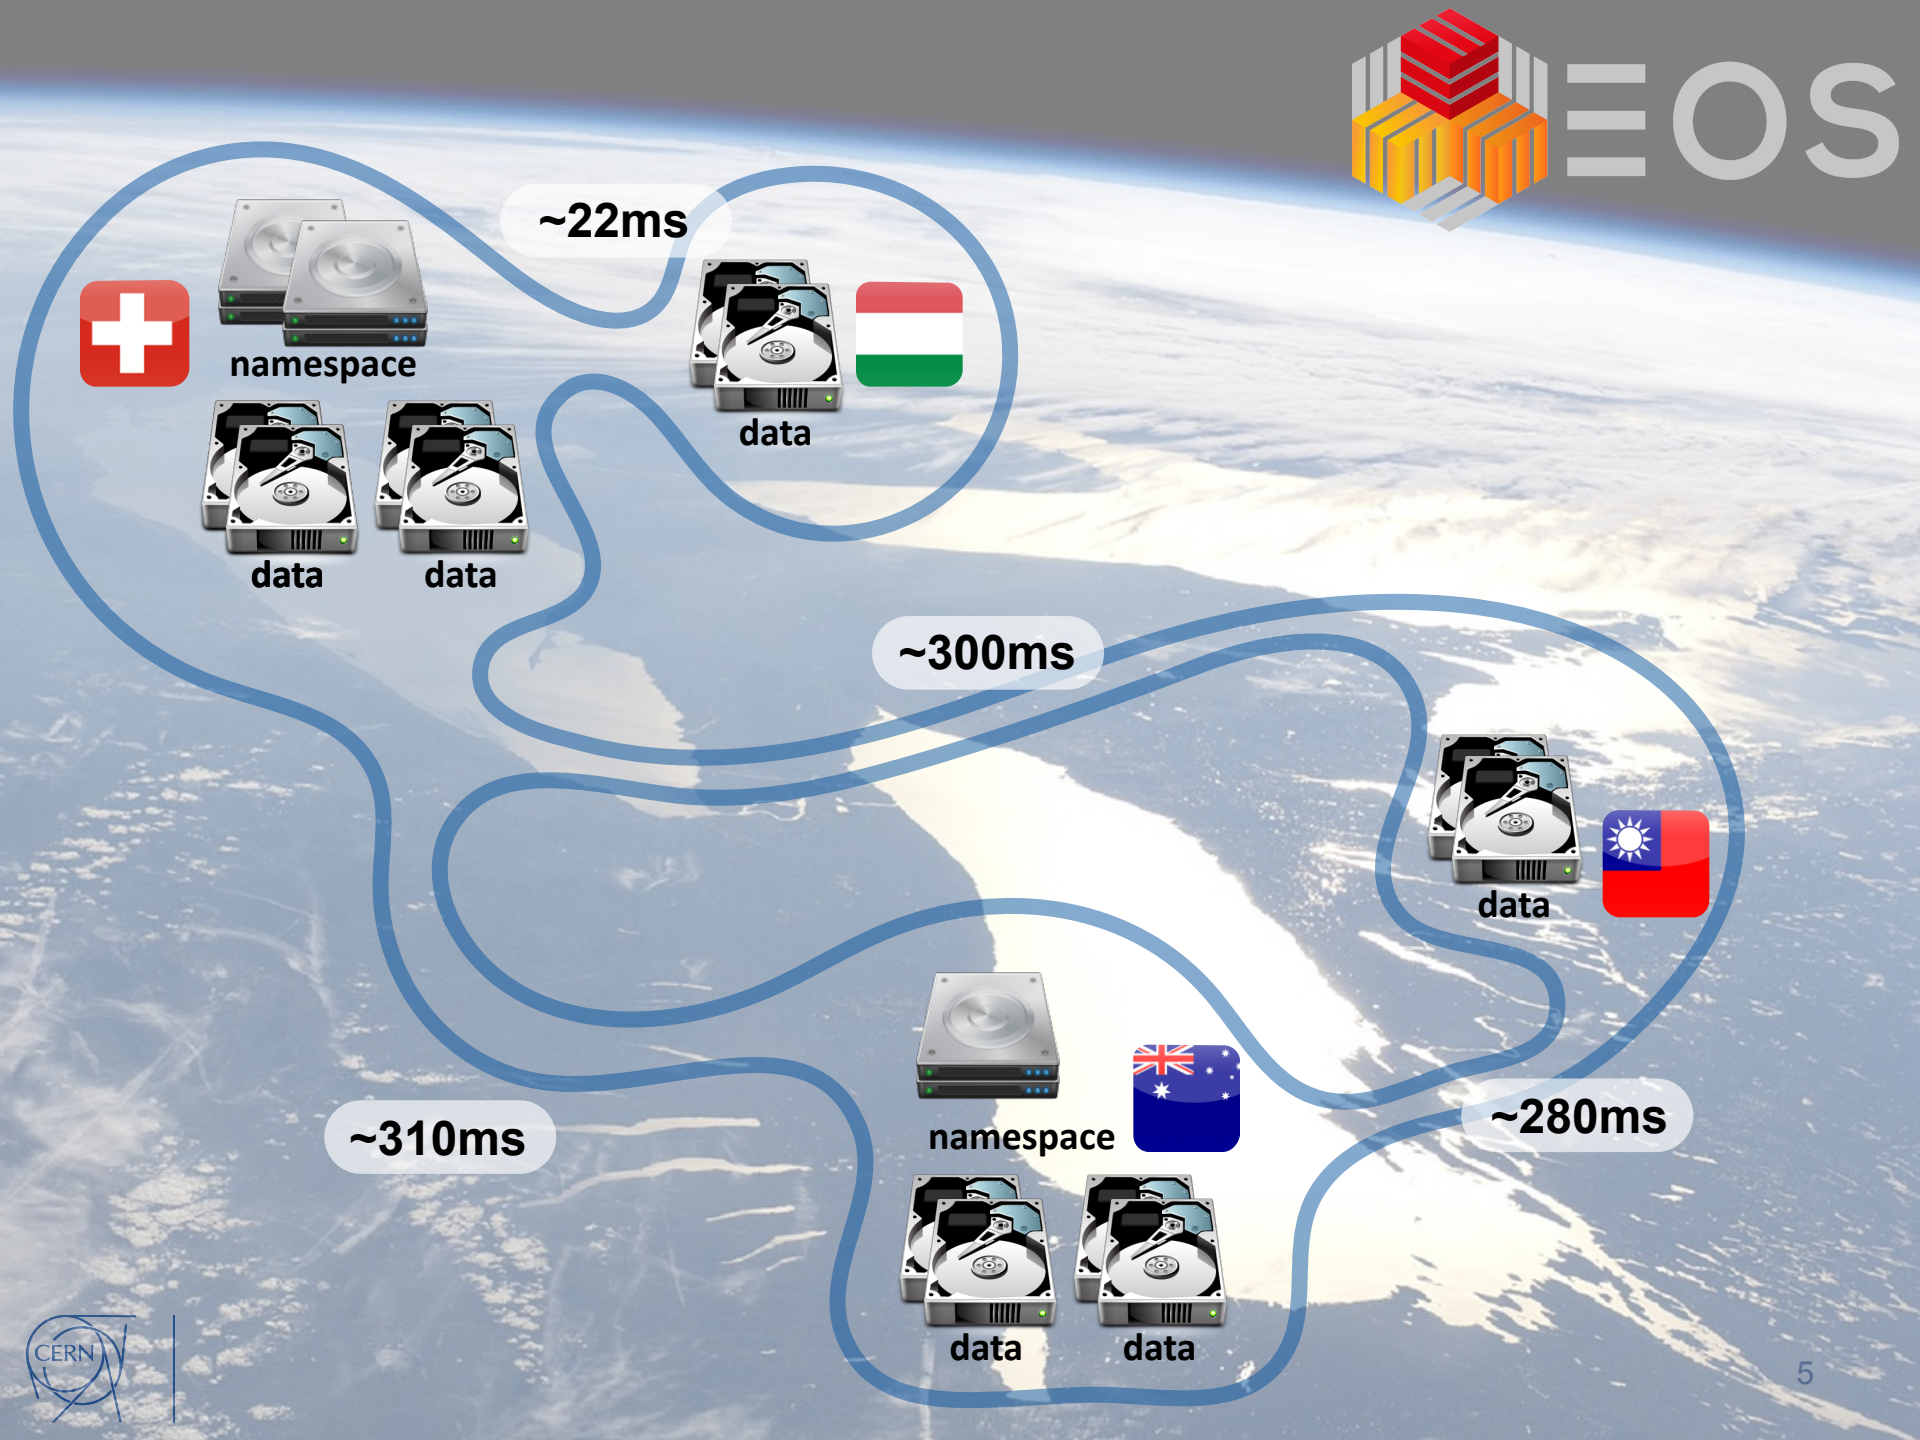
~22ms
namespace
data
data
data
~300ms
data
namespace
data
data
~310ms
~280ms

Storage Pools Overview

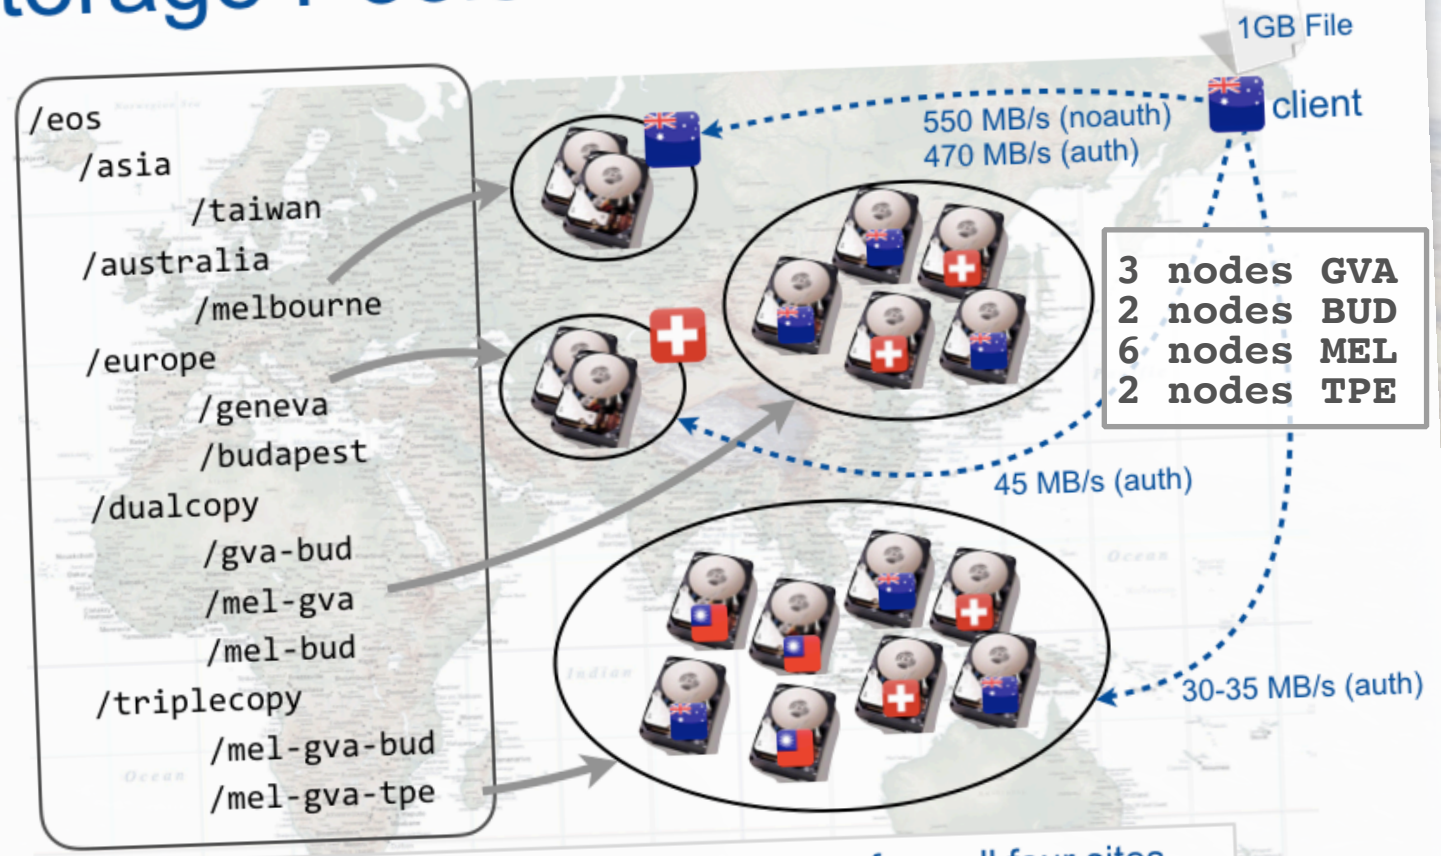

Storage pools were created with filesystems from all four sites.
Files were replicated according to the different configured policy
(e.g. 3 replicas: MEL-GVA-TPE).

Make data access easy
Make Analysis simple
Facilitate Science

aarnet
Australia's Academic and Research Network

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

aspera RSYNC

- 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

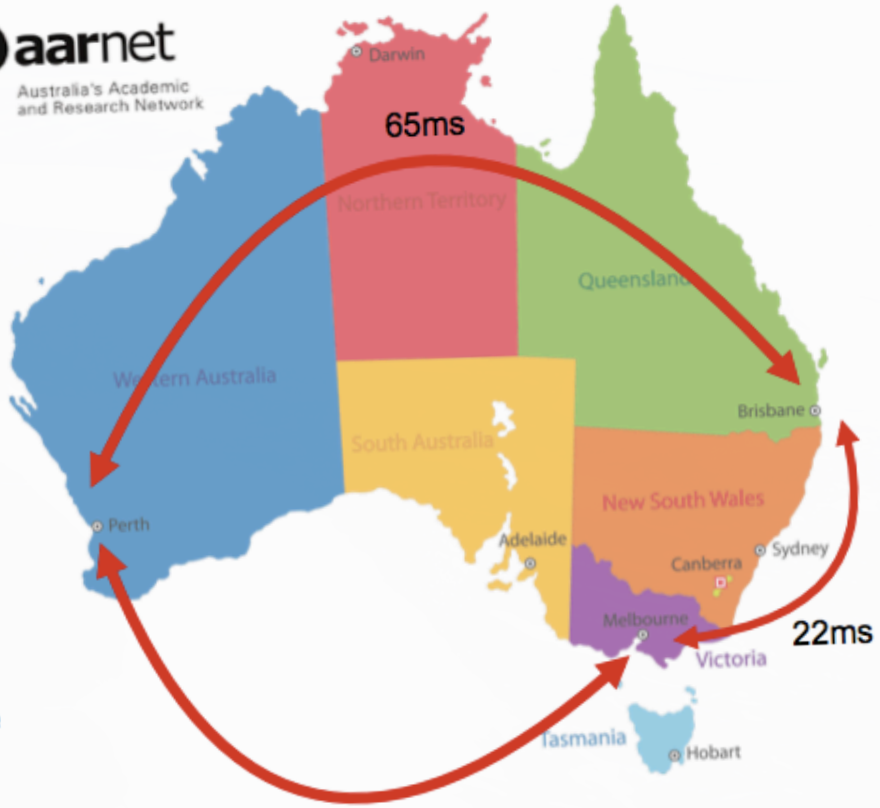

"This system is presently running 0.3.187, and has been so trouble free that I keep forgetting it's there." David Jericho -AARNetSolutions Architect

Background image: Shutterstock

CERN

EOS

National Research Centre (NRC)
"Kurchatov Institute"

Big Data Technologies Laboratory
http://bigdatalab.nrcki.ru/

Federation topology

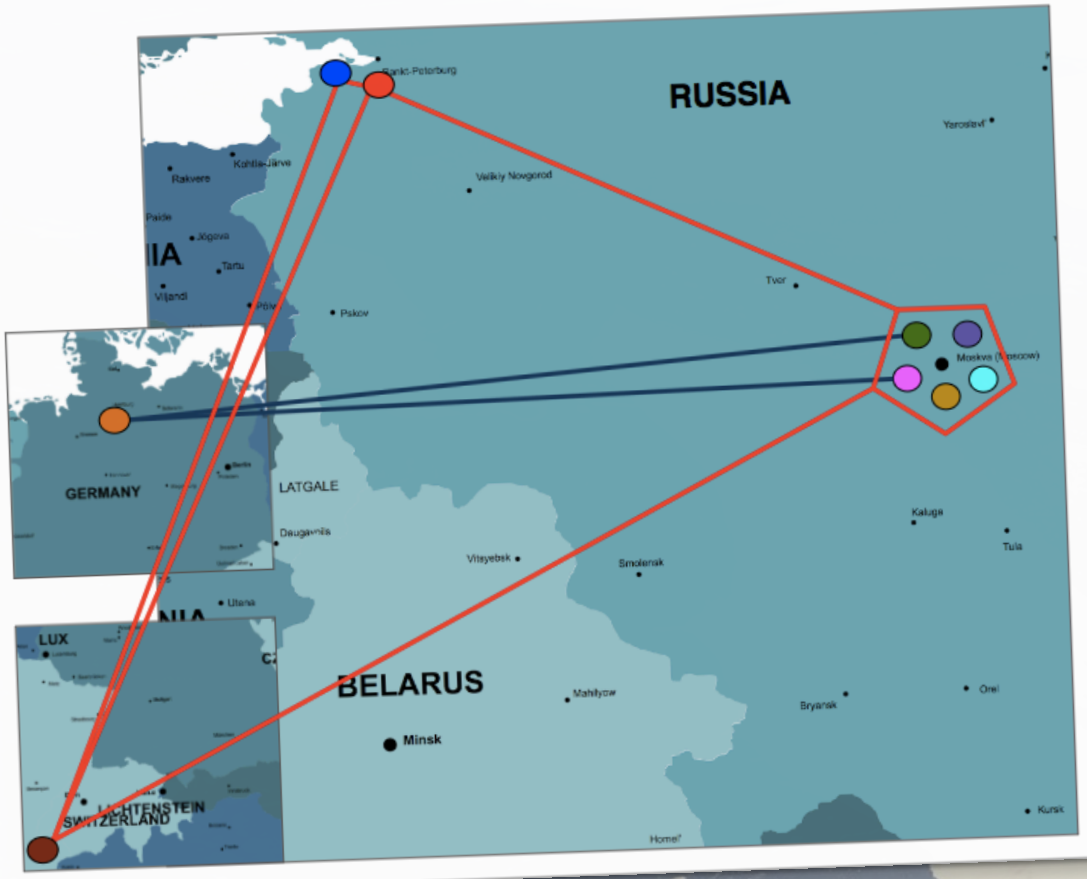

Storage Management

Are we able to reduce the complexity
and the expertise globally required?

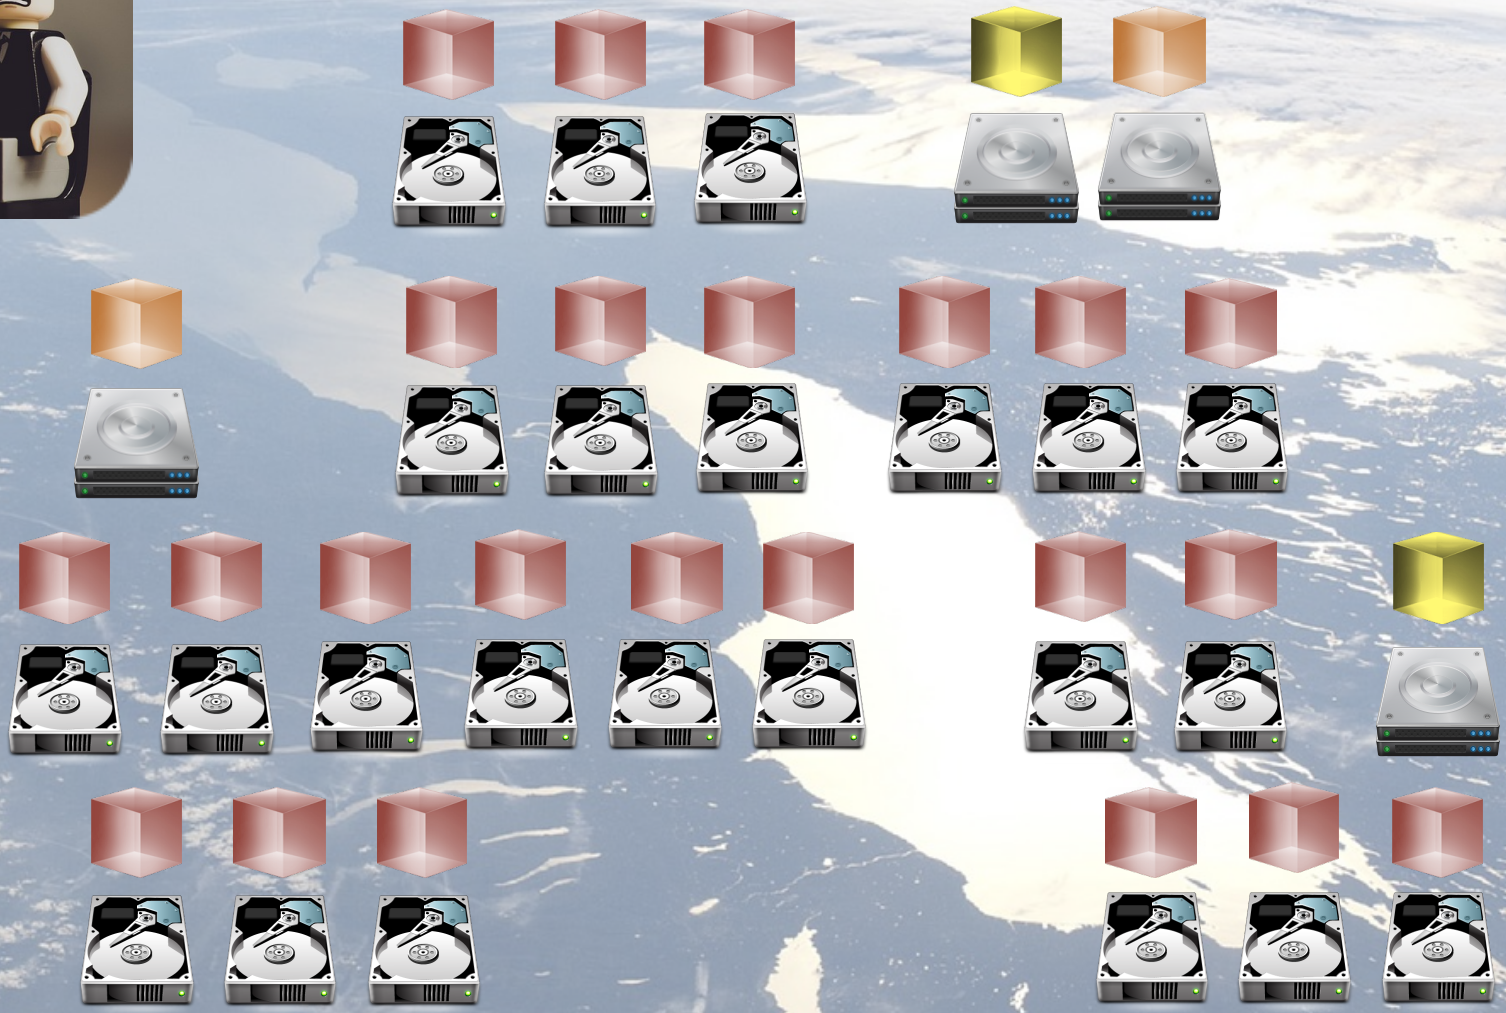

EOS

Decoupling storage software and hardware maintenance

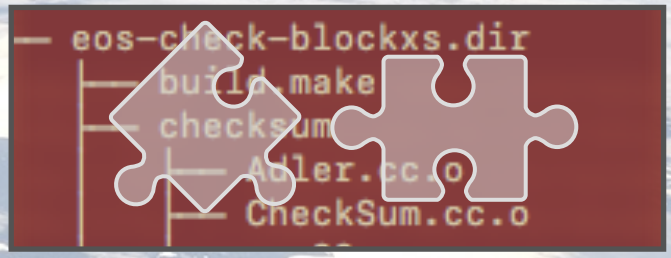

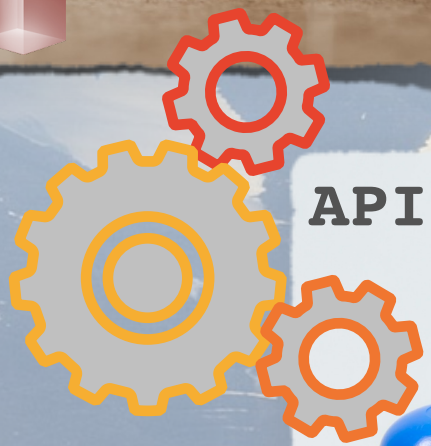

API*

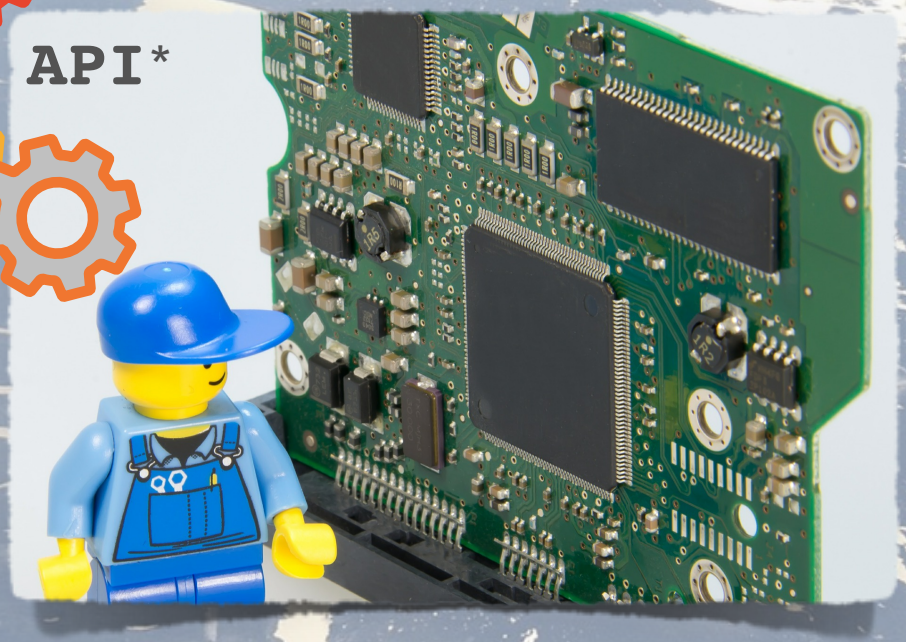

* Already in use by CERN repair team and sysadmins team

Network Evolution

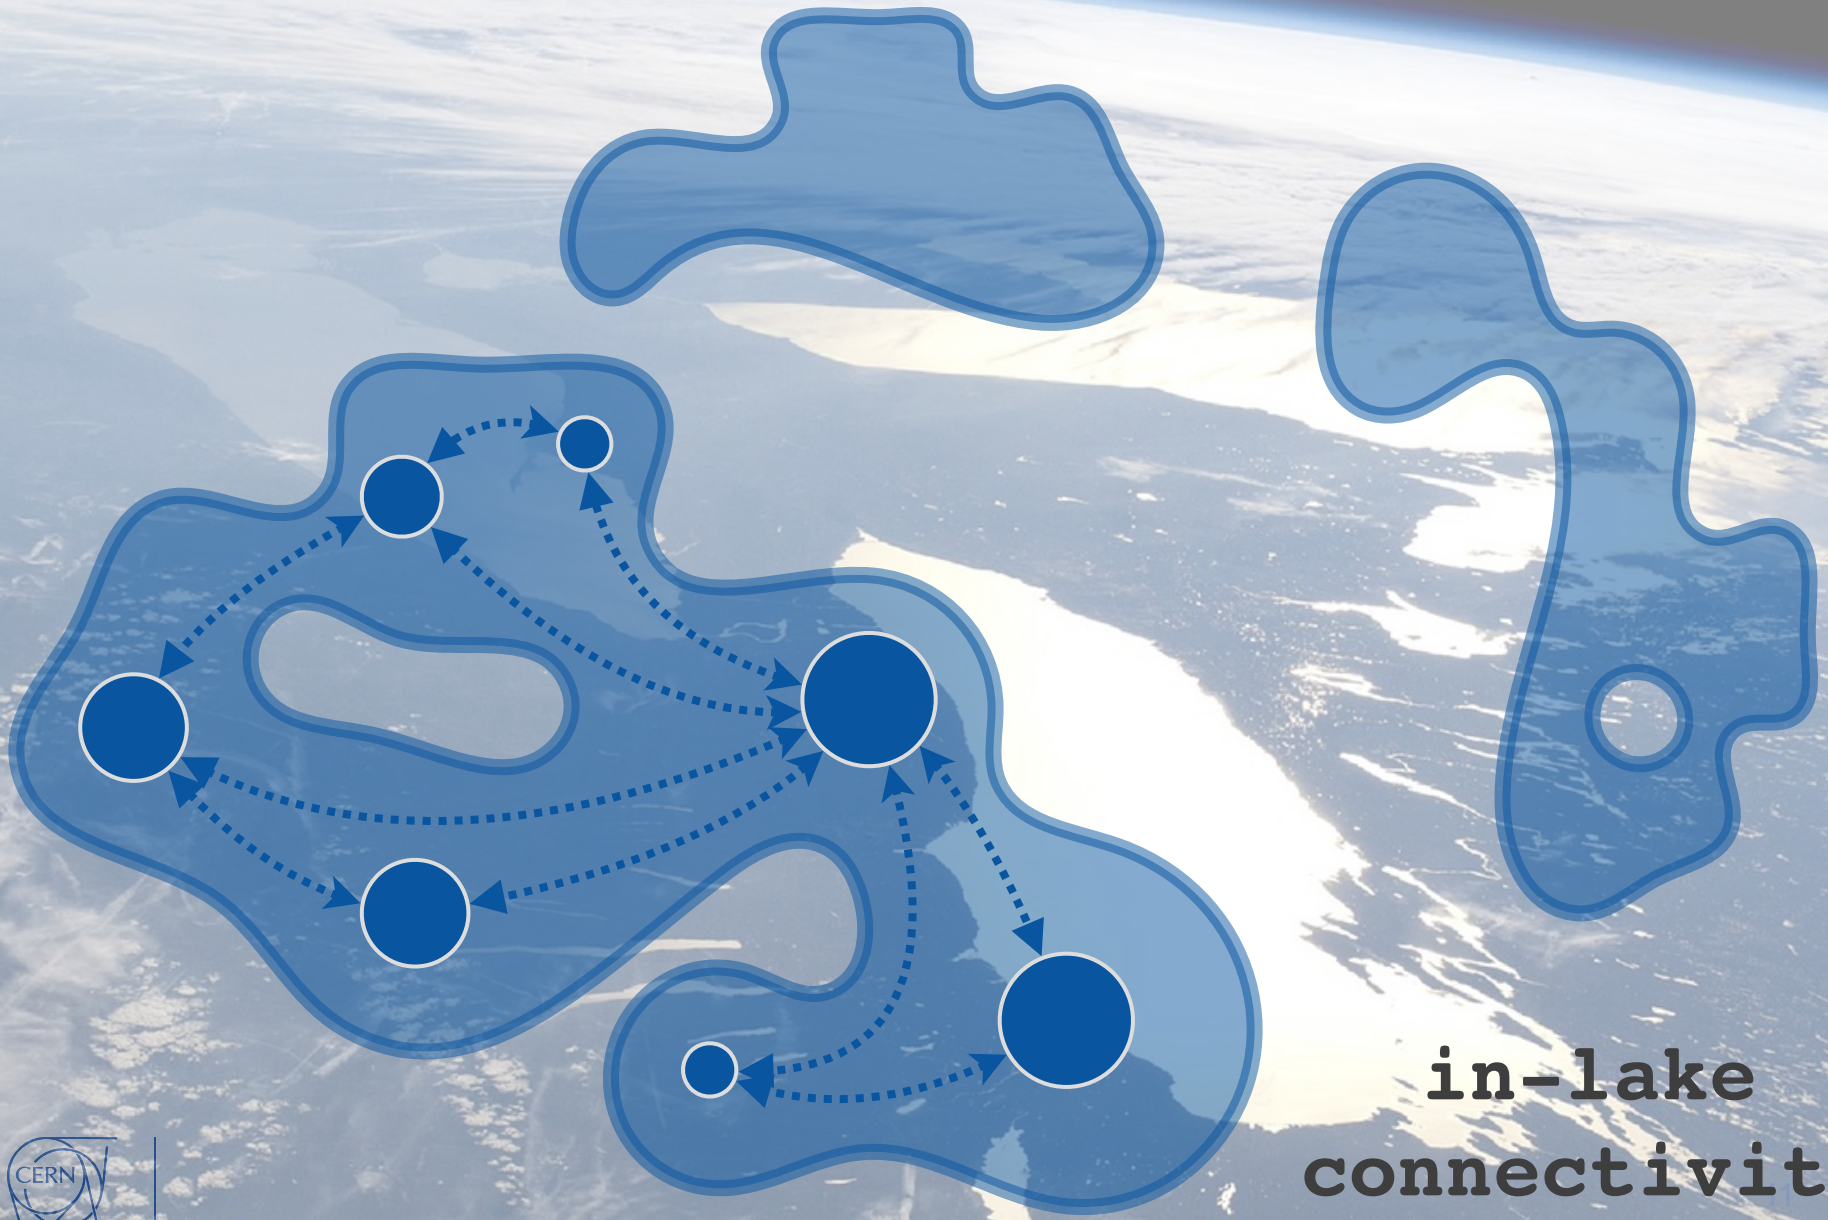

in-lake
connectivity

Network Evolution

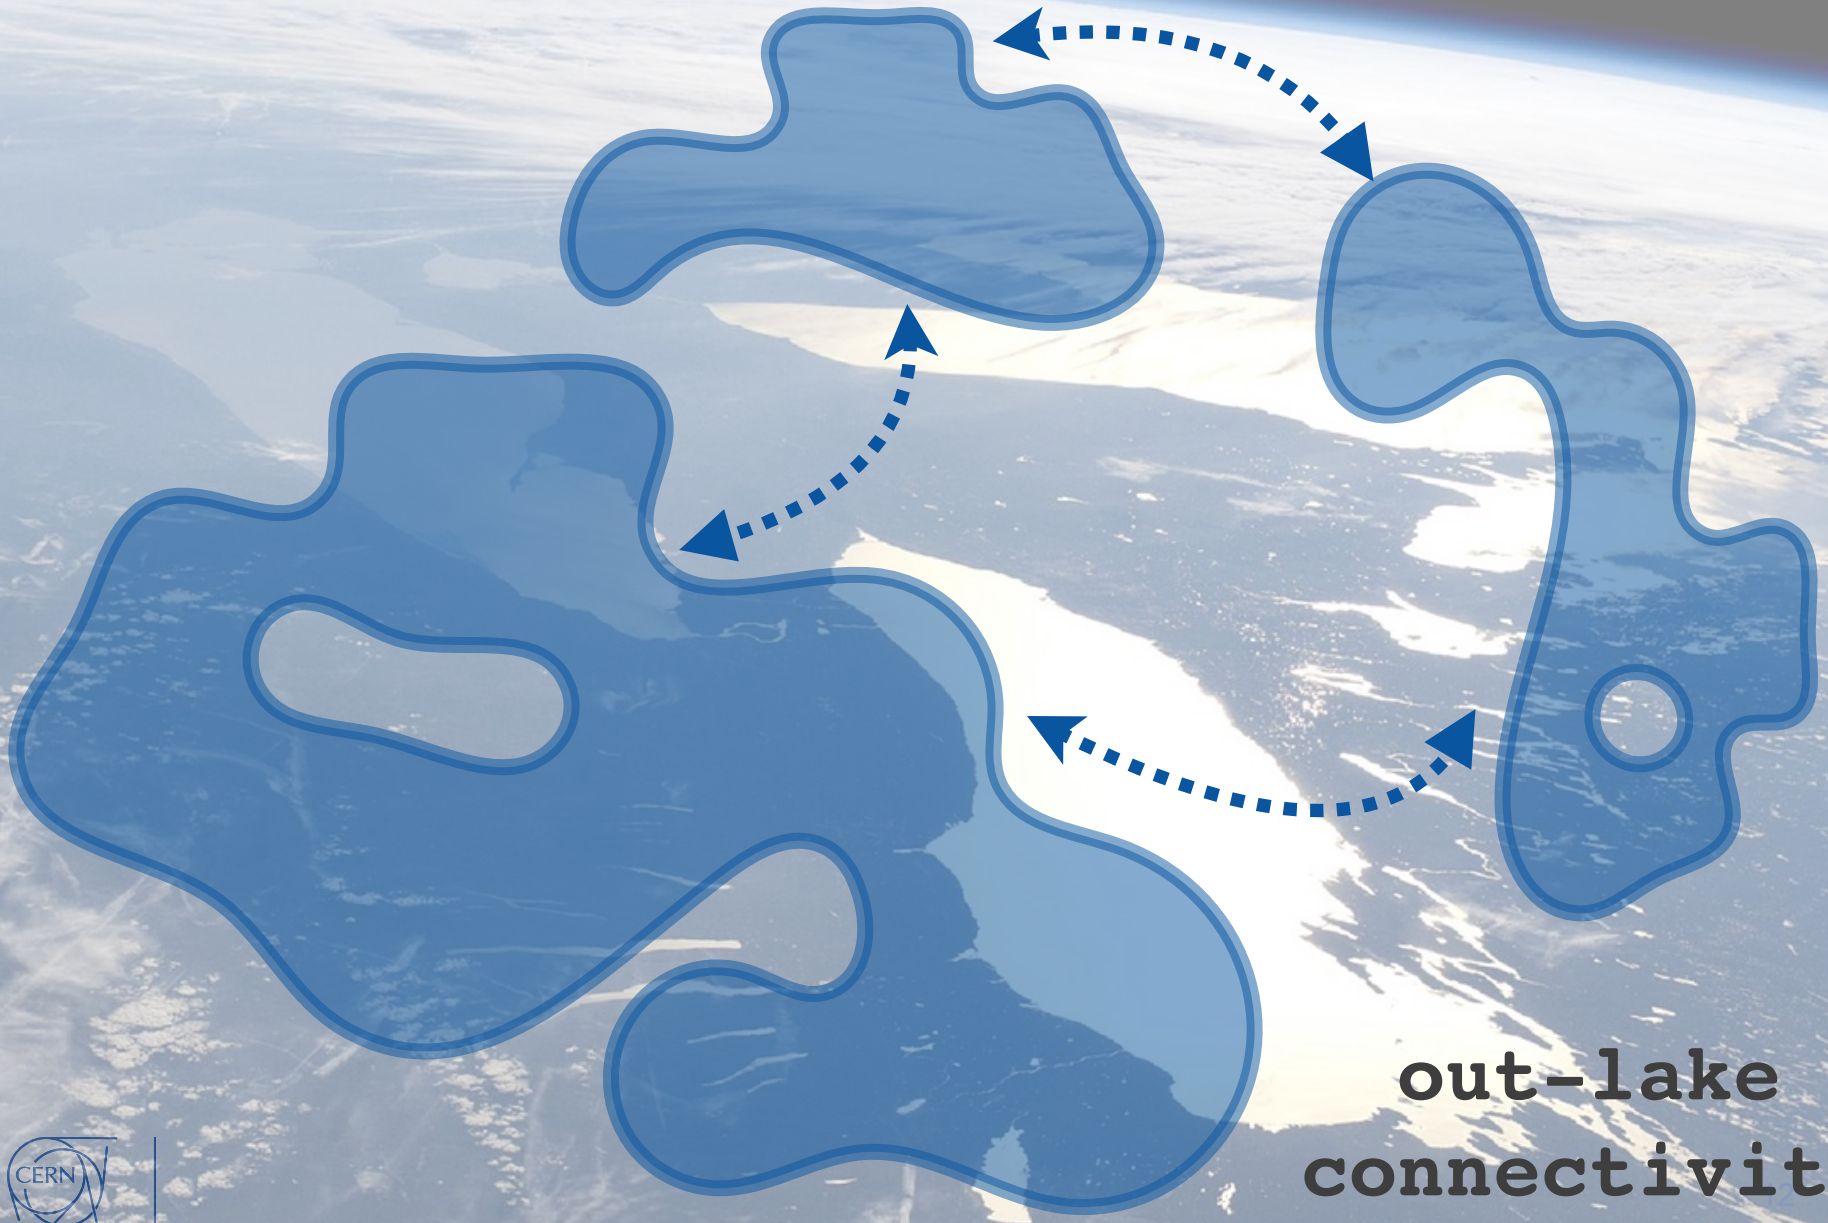

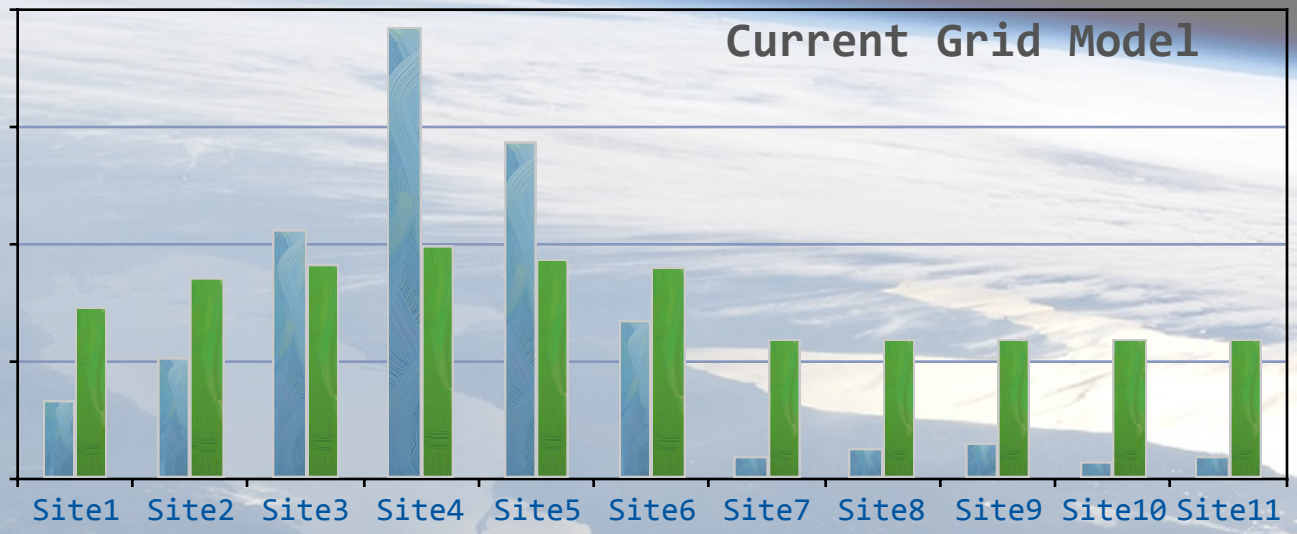

Resource and Manpower Consolidation

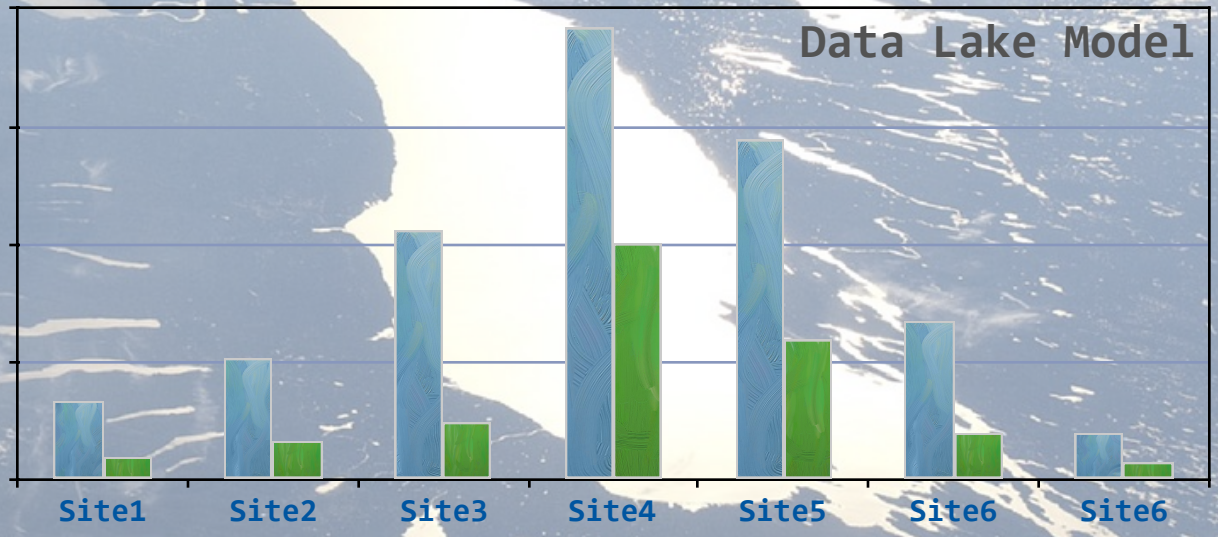

Thanks for the attention!

www.cern.ch