


Storage and VOSpace @ OATS

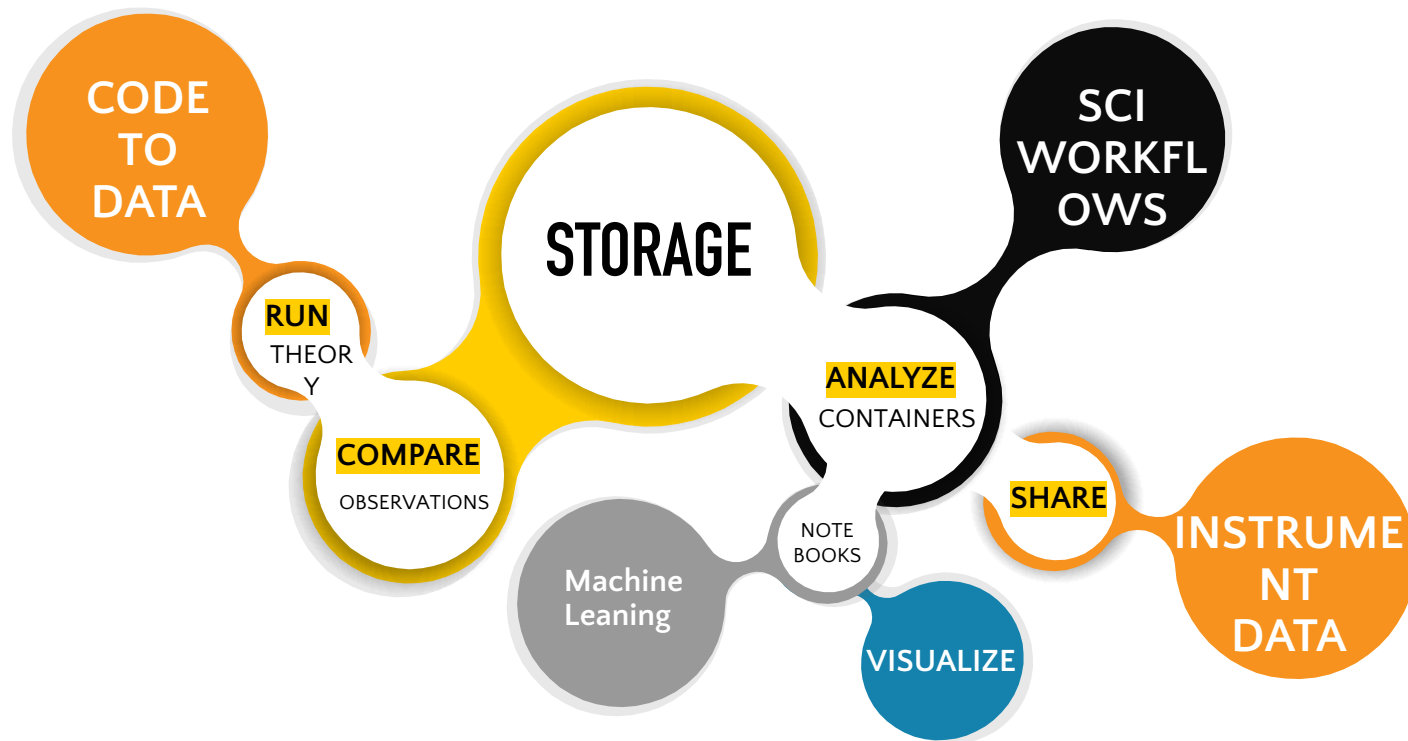
*GIULIANO TAFFONI, MARCO MOLINARO, SARA
BERTOCCO*

IVOA 2018 – 29 MAY 2018 – VICTORIA - CA



Horizon 2020

Why are we talking about storage

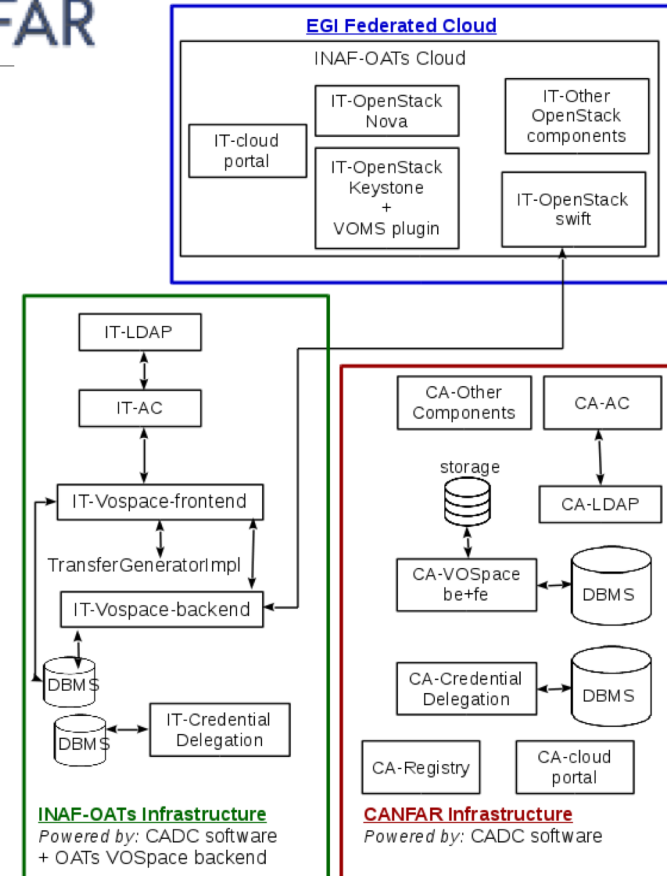


EGI-Engage: EGI CANFAR collaboration




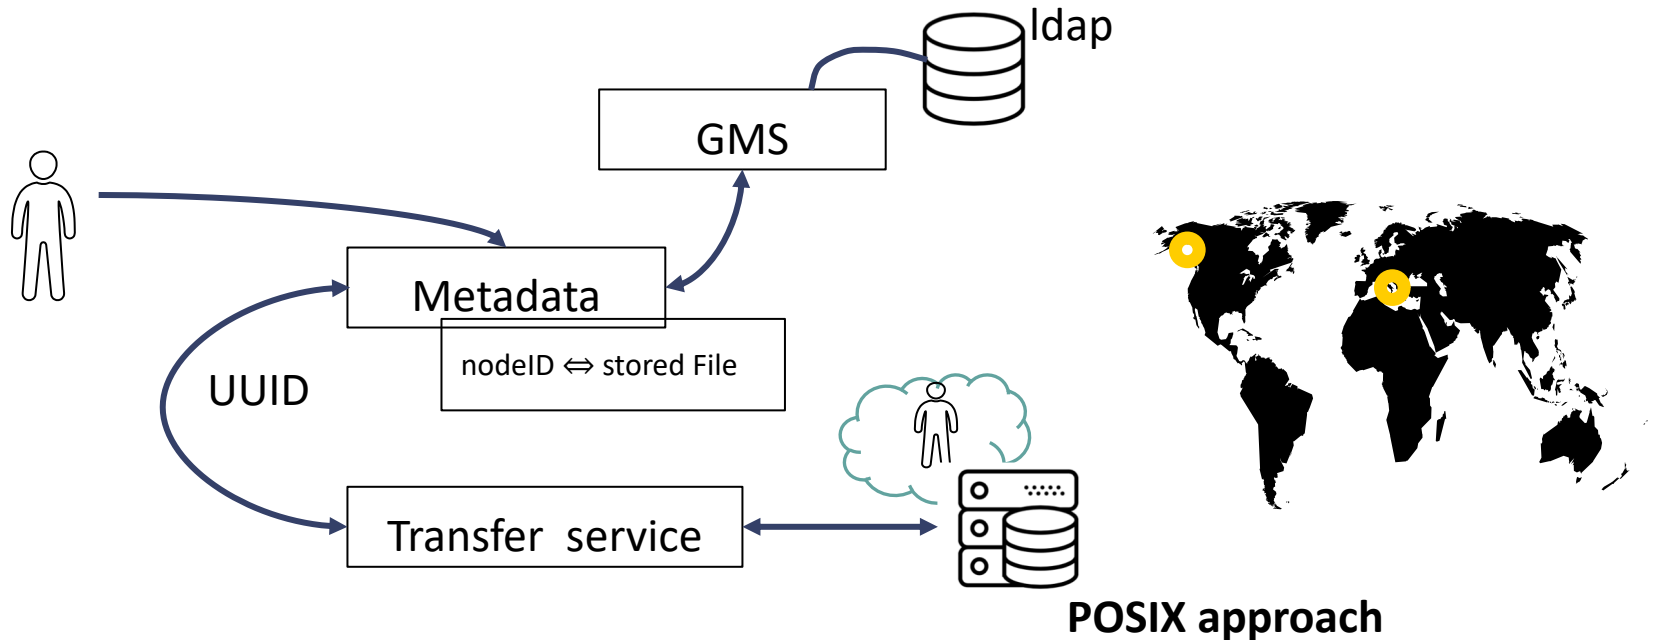
“building an Hybrid Cloud that **interconnects** the CANFAR cloud infrastructure and the EGI Federated Cloud, both remaining **unique entities**, [...] bounded together by means of a **common authentication model** and by the IVOA standards implementation.”

- the Access Control Service
- the VOSpace service
- the Credential Delegation Service



«Cloud access to interoperable IVOA-compliant VOSpace storage», A&C paper



EGI-CANFAR VOSpace implementation




Share data

FUSE mount to make computation but it is slow

EGI-CANFAR VOSpace Cloud implementation



Share data



ExaScale Projects

Co-Design

of a ground- breaking platform capable of scaling peak performance to 400 PFlops per **30MW**

CPU

Tape-out of a novel EuroExa **Arm-based** processing unit and integration **FPGA** for prototyping and **data-flow** acceleration

Interconnect

Novel/unique hybrid, geographically-addressed, low latency high bandwidth switching **interconnect**

Accelerators

A homogenized software platform including **heterogeneous acceleration support;**

Applications

A rich mix of key HPC applications

Storage

Scalable high performance low latency **data management** based on **BeeGFS;**



20MEuro to deployment of a **demonstrator** of a scalable, integrated and operational petaflops level prototype or **400 PFLOPS**.

HPC Storage: BeeGFS

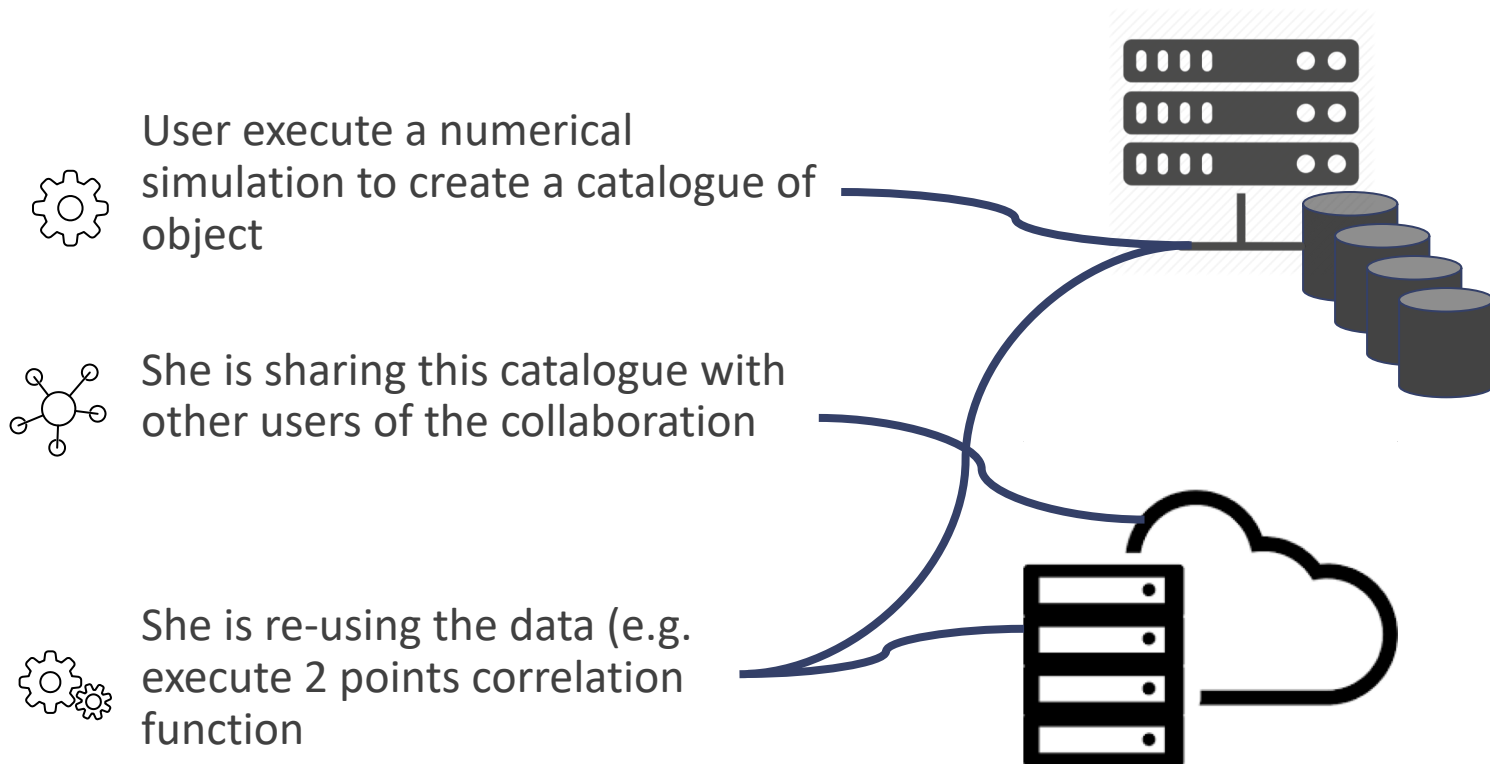
Optimized for performance
(low latency and high
throughput)

Scalable

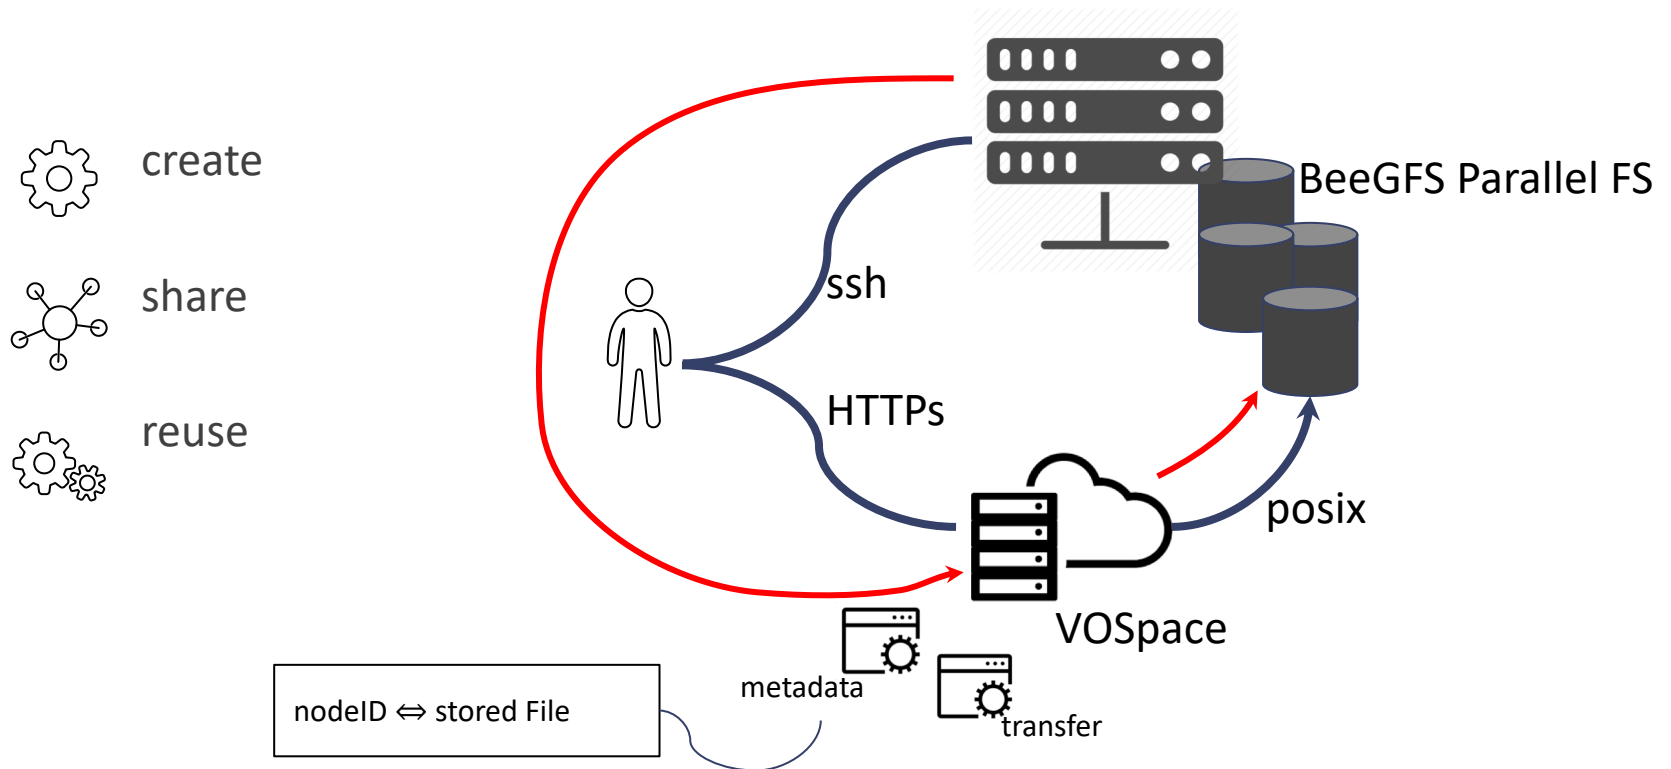
Redundancy on demand




Use case #2




The standard implementation



User space != Data space




Userspace == dataspace



How can we share authorizations?

WARNING:


User can manipulate data live on the Storage
User can manipulate Authz → proxy



NOTE:

Users must exist on my cluster
Group must be local (no “groups” interoperability)
Group must exist on cluster

One GMS over 3 services



Conclusions

We are experimenting on VOSpace for HPC (and HPCA)

There is not yet a real solution from our side.

New challenges in storage are emerging from scientific needs: data “close” to computing, intelligent multi-tier storage (WAN, LAN, etc)

Object Storage as new HPC and HPDA infrastructure