ESCAPE
European Science Cluster of Astronomy & Particle physics ESFRI research Infrastructures

CCIN2P3

Exploring XCache

Paul Musset

CC-IN2P3/CNRS

ESCAPE - The European Science Cluster of Astronomy & Particle Physics ESFRI Research Infrastructures has received funding from the European Union's Horizon 2020 research and innovation programme under the Grant Agreement n° 824064.

Introduction

- CC-IN2P3
- ESCAPE
- Caching and XCache

Funded by the European Union's
Horizon 2020 - Grant N° 824064

CC-IN2P3

- Academic data centre in Lyon, France
- 80 groups/experiments served including:
  - HEP (LHC Tier 1)
  - Astrophysics (LSST, EUCLID, CTA)
  - Groups in other disciplines :
    - Bioinfo : BIOASTER
- Capacities
  - Compute (40kslots)
  - Storage (~100 PB)

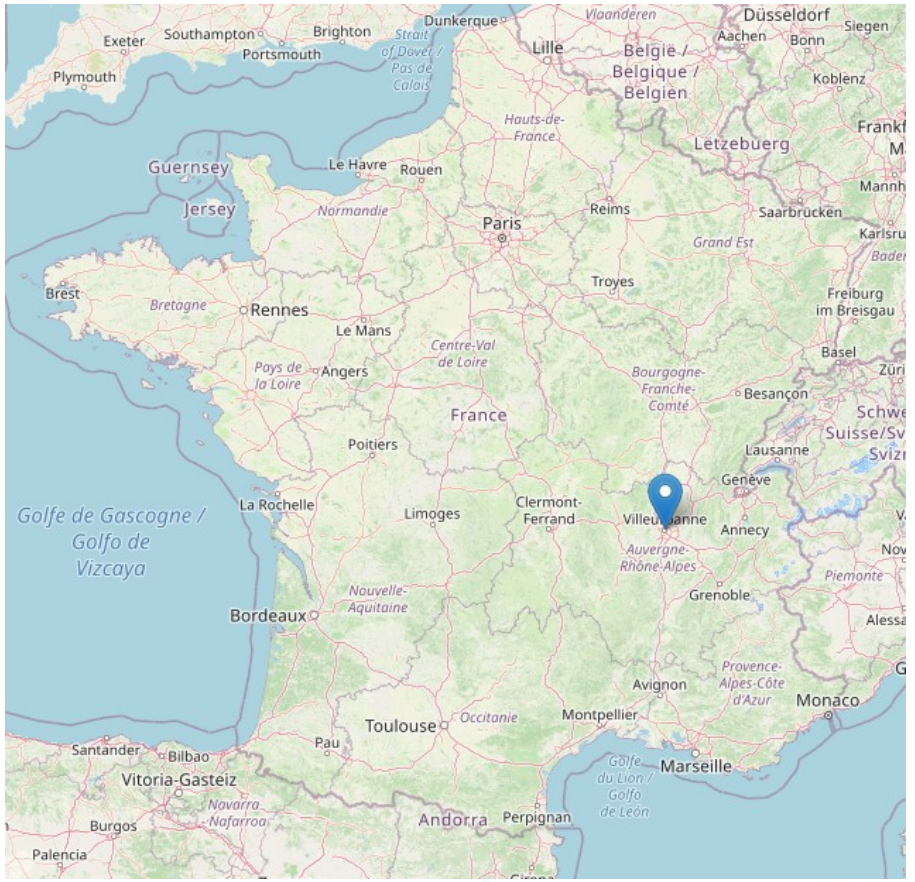

Funded by the European Union's Horizon 2020 - Grant N° 824064

Escape

- Prototype an infrastructure adapted to Exabyte-scale needs of large science projects
- Ensure the sciences drive the development of the EOSC
- Address FAIR data management principles

Funded by the European Union's Horizon 2020 - Grant N° 824064

WP2: Data Infrastructure for Open Science

- Implementation of a Data Lake
- Services operated by the partner institutes to do data managements, storage
- Many services used/tested
  - Rucio, FTS, AAI, XCache

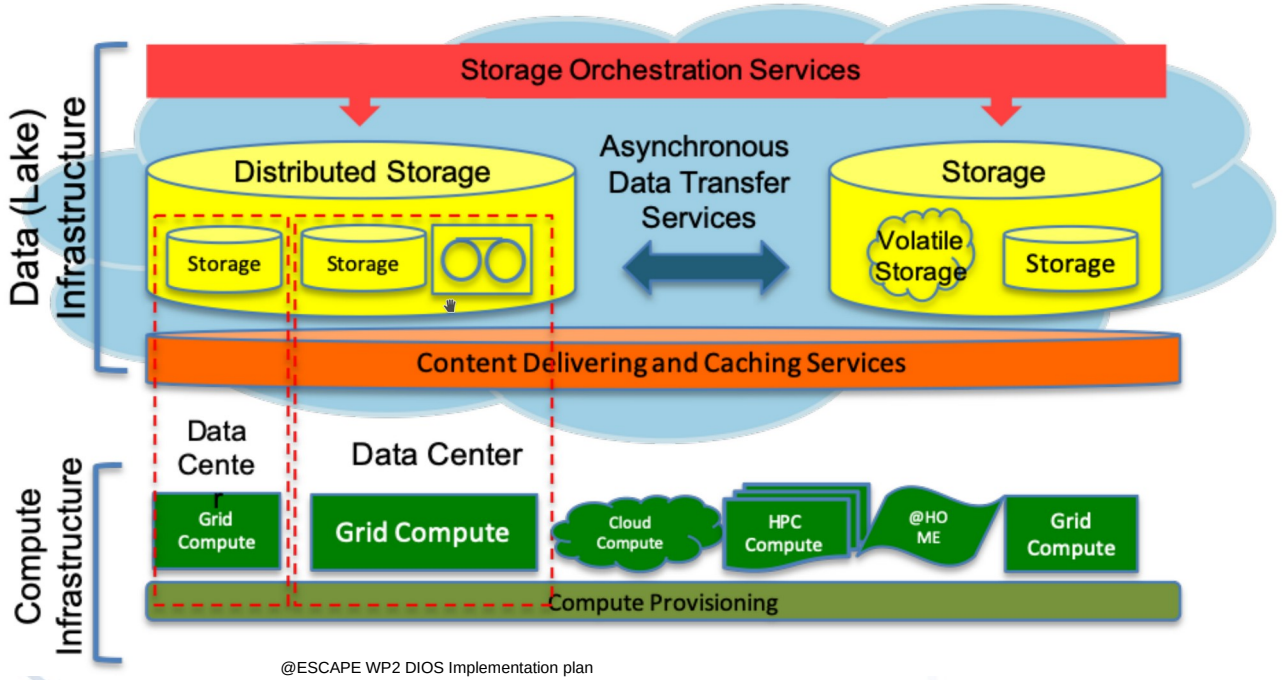

@ESCAPE WP2 DIOS Implementation plan

Funded by the European Union's Horizon 2020 - Grant N° 824064

Caching

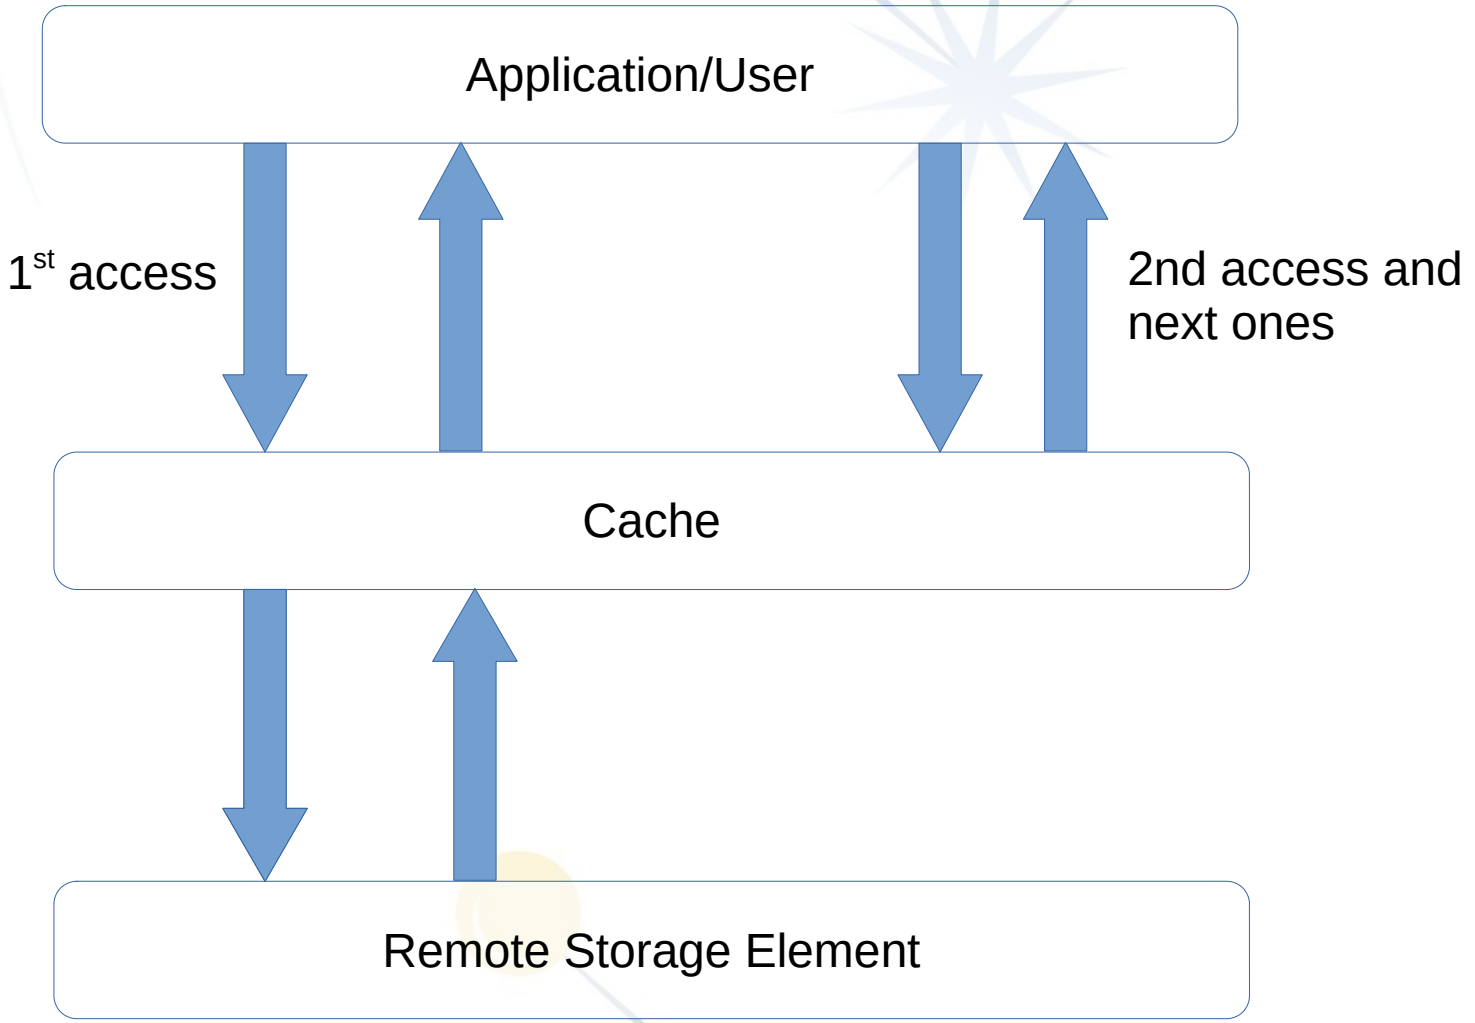

Funded by the European Union's Horizon 2020 - Grant N° 824064

Xcache

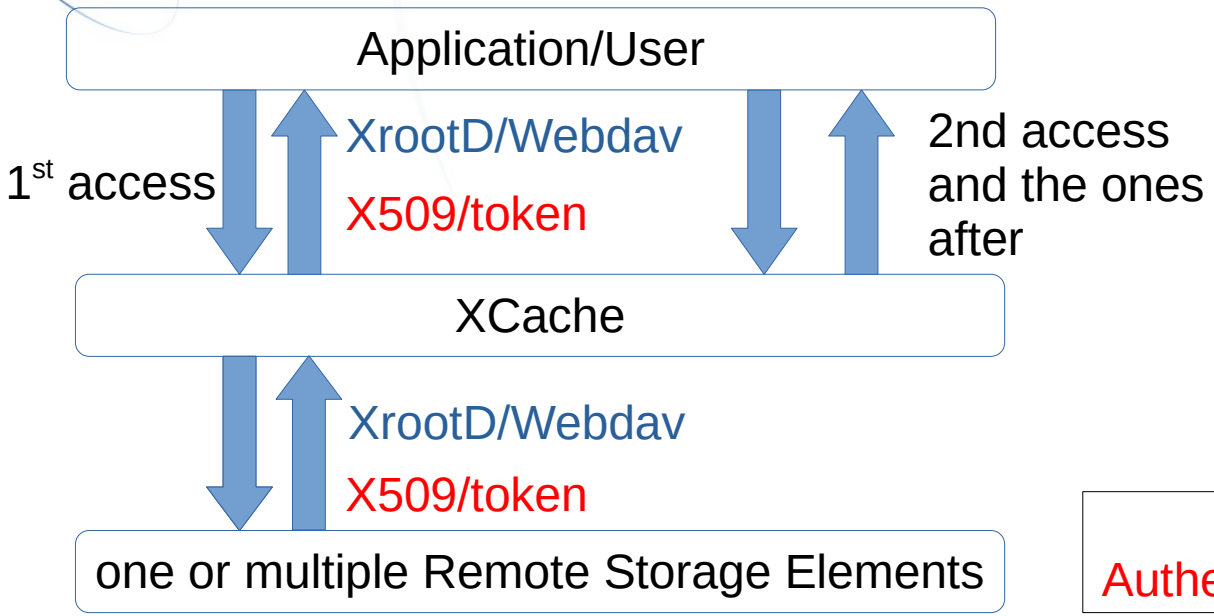

- XrootD plugin
- Authentication possible via X509/Token
- Scalable to multiple servers

Protocols
Authentication/Authorization

```
ccosvms0092:/pbs/home/p/pmusset(0)>xrdcp -f xroot://cc-xcache-escape-certificate-test.in2p3.fr:1094//root://dcache-se-doma.desy.de:1094/es
cape/wp2_rucio_testbed/desy_dcache/NESSIE/65/64/tes2_iorandom_2048MB_12.out /dev/null
[2GB/2GB][100%][==================================================][46.55MB/s]  s]
ccosvms0092:/pbs/home/p/pmusset(0)>xrdcp -f xroot://cc-xcache-escape-certificate-test.in2p3.fr:1094//root://dcache-se-doma.desy.de:1094/es
cape/wp2_rucio_testbed/desy_dcache/NESSIE/65/64/tes2_iorandom_2048MB_12.out /dev/null
[2GB/2GB][100%][==================================================][136.5MB/s]  s]
```


Funded by the European Union's Horizon 2020 - Grant N° 824064

Interest of XCache

- Data access to remote site
- For storage-less sites
  - No need for data management
  - Easier to manage than normal storage
- For CC-IN2P3
  - HL-LHC
  - LSST

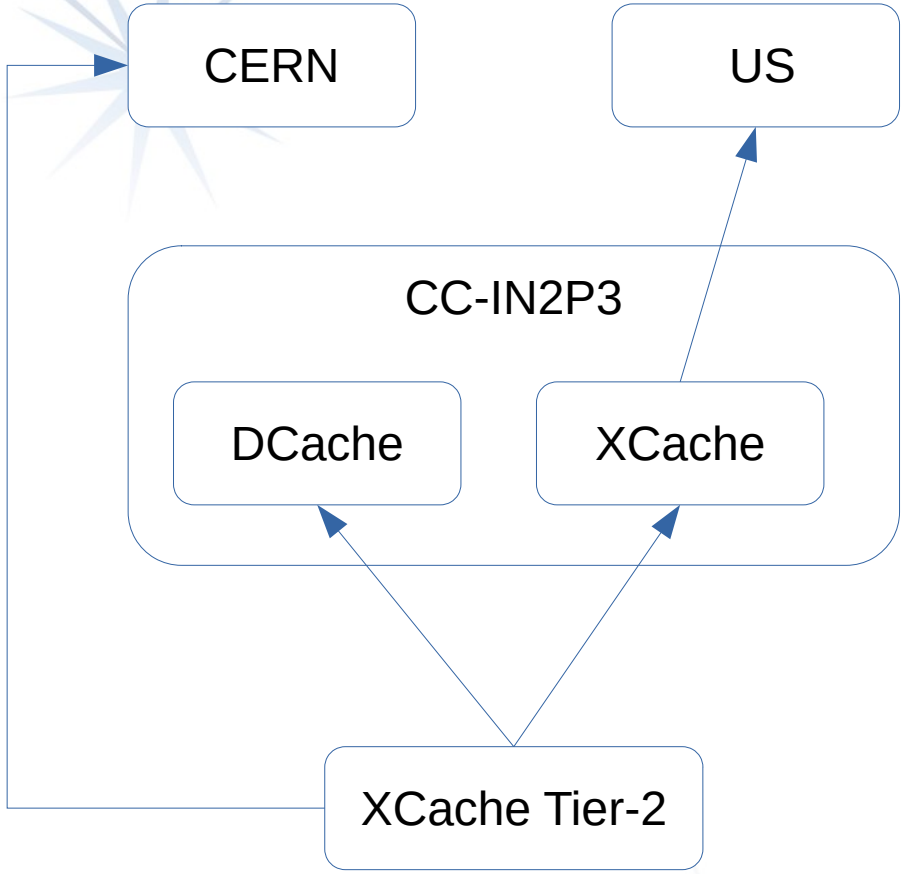

Usage Example

Funded by the European Union's Horizon 2020 - Grant N° 824064

My experience with XCache

- Different instances I installed (VM, Bare Metal)
- Setting up XCache
- Next steps

Funded by the European Union's
Horizon 2020 - Grant N° 824064

First instance

- On a standard Openstack VM
  - 8GB RAM, 4VCPUs, 80GB disks
- Installation by hand
- Had to learn some things from zero
  - XrootD : Documentation not very beginner friendly
  - Authentication/Authorization: Certificate/Token

Funded by the European Union's
Horizon 2020 - Grant N° 824064

Test first instance

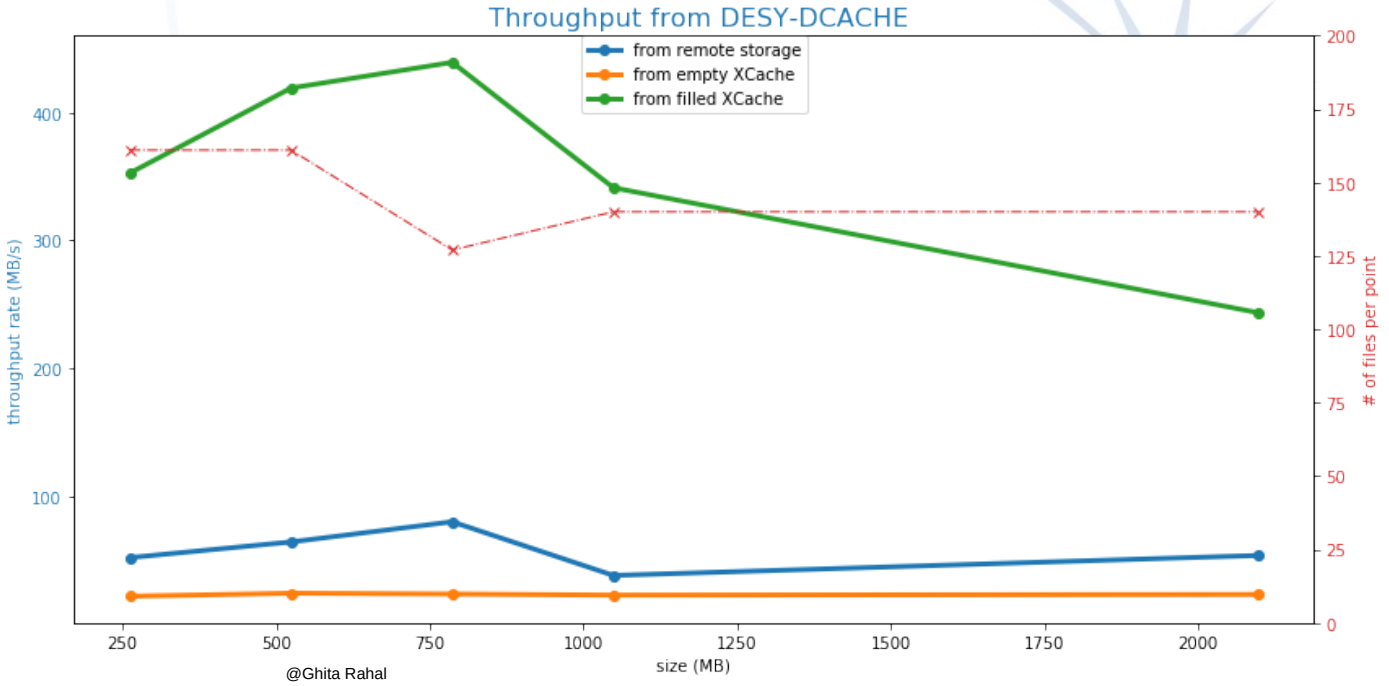

- Lower speed for bigger files
- Due to network parameter on the VM
  - Average (40MB/s), Peak (100MB/s), Burst (3GB/s)

Funded by the European Union's Horizon 2020 - Grant N° 824064

2nd Instance

- VM without network restriction
- More consistent result

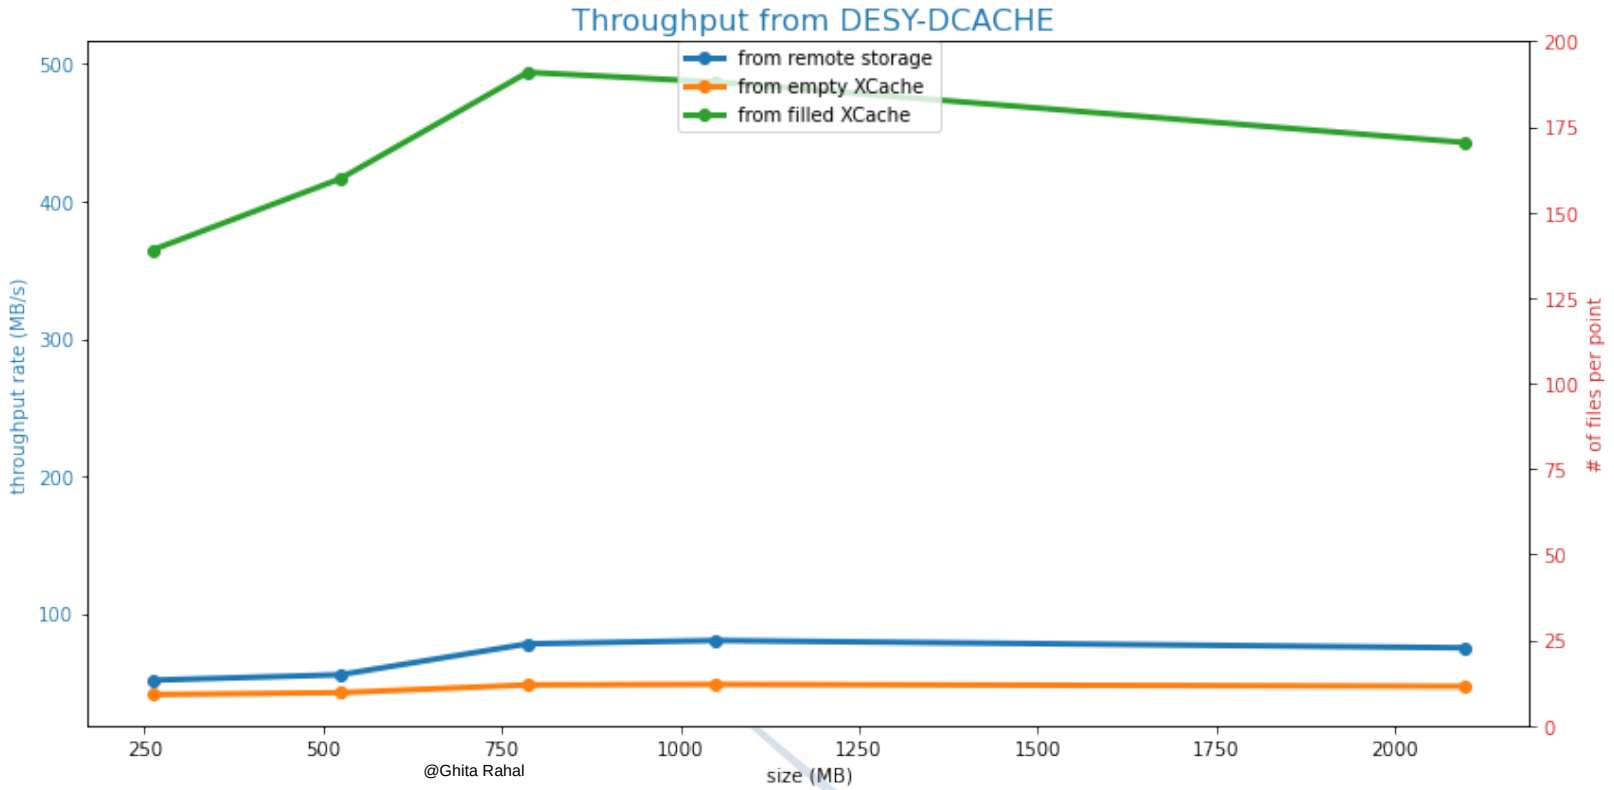

Funded by the European Union's Horizon 2020 - Grant N° 824064

Setting up XCache

- Docker
  - To have a better control on versioning of XCache and its dependencies
  - To deploy it more easily on a new machine
  - Orchestrated by docker-compose
- Ansible
  - Setup of the machine (partitions, users…)
  - WIP
- https://gitlab.in2p3.fr/CC-Escape/xcache-config

Funded by the European Union's
Horizon 2020 - Grant N° 824064

Namespace and (meta) data management

- Data on multiple disks
  - Using XRootD spaces
  - No need for hardware RAID
  - Only part of the data is lost in case of hardware failure
- Separate namespace, metadata and data
  - More operation on namespace and metadata

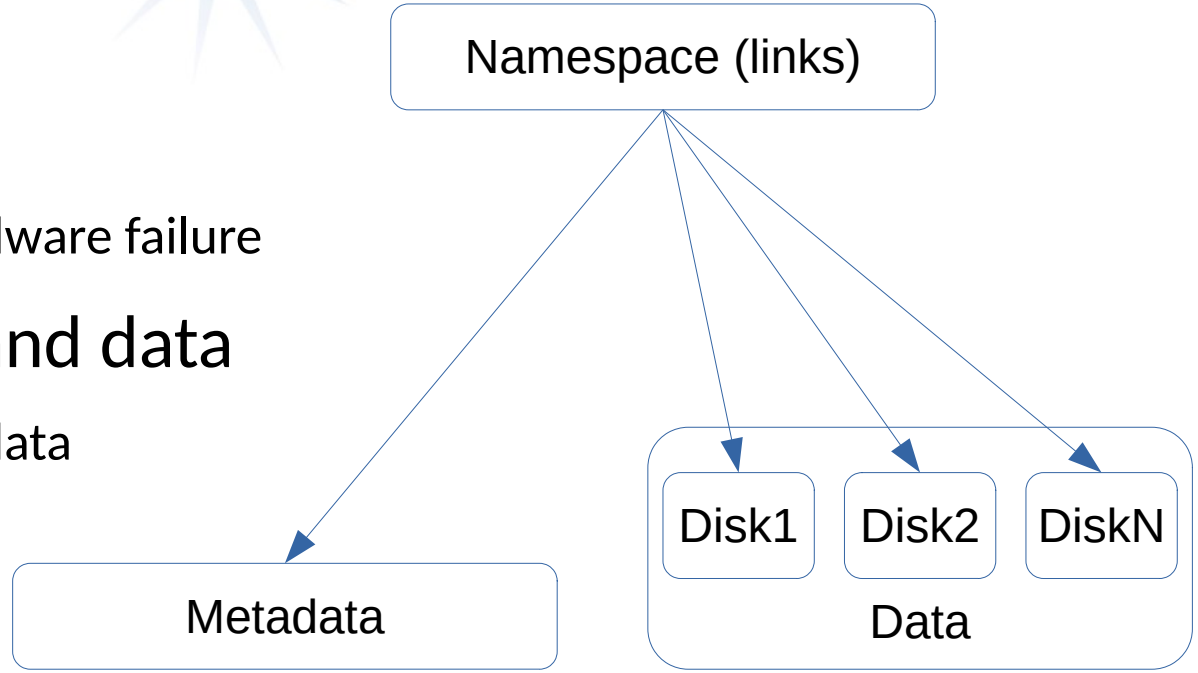


```
oss.localroot /mnt/xcache-ns
oss.space meta /mnt/xcache-meta
oss.space data /mnt/xcache-storage/data*
pfc.spaces data meta
```


Funded by the European Union's Horizon 2020 - Grant N° 824064

Bare Metal Instance

- Configuration
  - 32GB RAM
  - 24 5.5TB disks (~ 132 TB)
  - Intel Xeon E5-2609 v3
- Easy to install with experience installing on multiple VMs
- First try looks promising

Funded by the European Union's
Horizon 2020 - Grant N° 824064

Road Maps

- Go to XrootD5
  - Token authorization with Xroot protocol
  - Integrated Monitoring
- Improve documentation and configuration
- Interest for the operation (CC-IN2P3, Tier 2)
- And for the experiments/groups use case

Funded by the European Union's
Horizon 2020 - Grant N° 824064

ESCAPE
European Science Cluster of Astronomy & Particle physics ESFRI research Infrastructures

Questions ?

Funded by the European Union's
Horizon 2020 - Grant N° 824064