

# The XDC project



Data Management for extreme scale computing



Daniele Cesini

XDC – Project Coordinator

INFN

[daniele.cesini@extreme-datacloud.eu](mailto:daniele.cesini@extreme-datacloud.eu)



eXtreme DataCloud is co-funded by the Horizon2020  
Framework Program – Grant Agreement 777367  
Copyright © Members of the XDC Collaboration, 2017-2020

- ✗ The eXtreme DataCloud is a software development and integration project
- ✗ Develops **scalable** technologies for federating storage resources and managing data in highly distributed computing environments
  - ➡ Focus on efficient, policy driven and Quality of Service based DM
- ✗ The targeted platforms are the current and next generation e-Infrastructures deployed in Europe
  - ➡ European Open Science Cloud (EOSC)
  - ➡ The e-infrastructures used by the represented communities
- ✗ Addresses the EINFRA-21-2017 (b)-2: “Computing e-infrastructure with extreme large datasets”
  - ➡ Deal with heterogeneous datasets
  - ➡ Bring to TRL8 and include in a unified service catalogue services and prototype at least at TRL6


# The Approach

## ✕ Improve already existing, production quality Data Management services

- ➡ By adding **missing functionalities** requested by research communities
- ➡ Based mainly on technologies provided by the partners and by the INDIGO-Datacloud project
- ➡ Must be coherently harmonized in the European e-Infrastructures



INDIGO PaaS  
Orchestrator



INDIGO CDMI  
Server

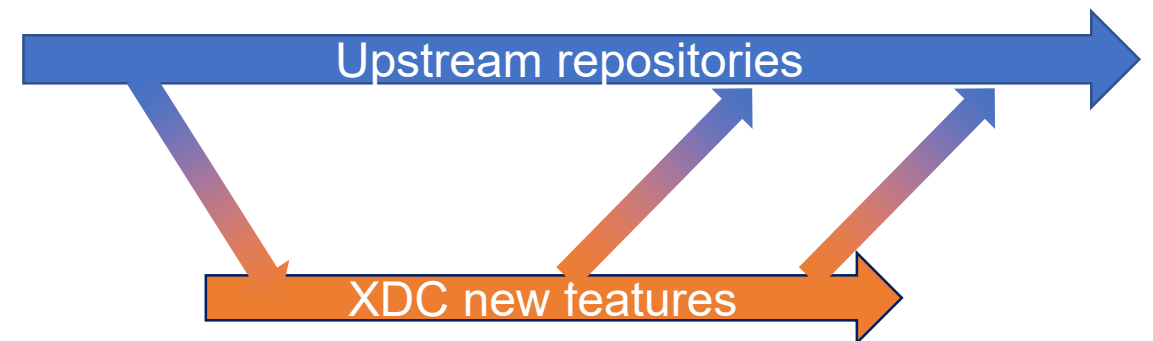


eXtreme-DataCloud Overview - Creating platform-driven  
e-Infrastructure innovation on EOSC - Athens






# XDC Approach

- ✗ The partners owning/involved in each of the tools are the main developers for that solution in XDC
- ✗ We always aim **to push back the code** in the main development tree on the **original projects**
  - This widely increase the **sustainability** of the services





# A User Driven Project



# XDC Topics

- ✗ Intelligent & Automated Dataset Distribution
  - ....→ Orchestration to realize a policy-driven data management
  - ....→ Data distribution policies based on Quality of Service (i.e. disks vs tape vs SSD) supporting geographical distributed resources (cross-sites)
  - ....→ Data lifecycle management
- ✗ Data pre-processing during ingestion
- ✗ Metadata management
- ✗ Data management based on storage events
- ✗ Smart caching
  - ....→ Transparent access to remote data without the need of a-priori copy
    - ....→ To support dynamic inclusion of diskless sites
    - ....→ To improve efficiency in multi-site storage systems and storage federations (i.e. Datalakes)
- ✗ Sensitive data handling
  - ....→ secure storage and encryption

# XDC Consortium

| ID | Partner        | Country | Represented Community           | Tools and system         |
|----|----------------|---------|---------------------------------|--------------------------|
| 1  | INFN<br>(Lead) | IT      | HEP/WLCG                        | INDIGO-Orchestrator      |
| 2  | DESY           | DE      | Research with Photons<br>(XFEL) | dCache                   |
| 3  | CERN           | CH      | HEP/WLCG                        | EOS, DYNAFED, FTS, RUCIO |
| 4  | AGH            | PL      |                                 | ONEDATA                  |
| 5  | ECRIN          | [ERIC]  | Medical data                    |                          |
| 6  | UC             | ES      | Lifewatch                       |                          |
| 7  | CNRS           | FR      | Astro [CTA and LSST]            |                          |
| 8  | EGI.eu         | NL      | EGI communities                 |                          |

- ✕ 8 partners, 7 countries
- ✕ 6 research communities represented + EGI
- ✕ XDC Total Budget: 3.07Meuros

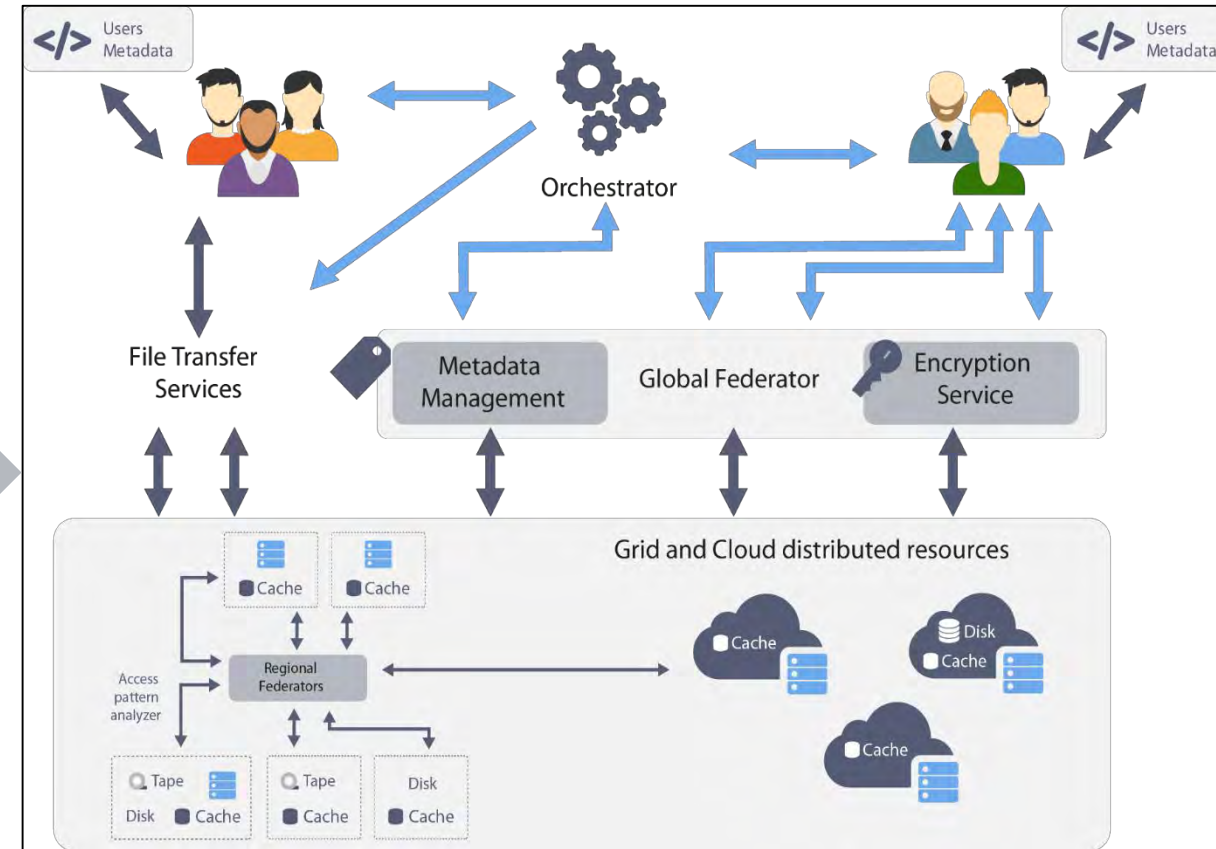
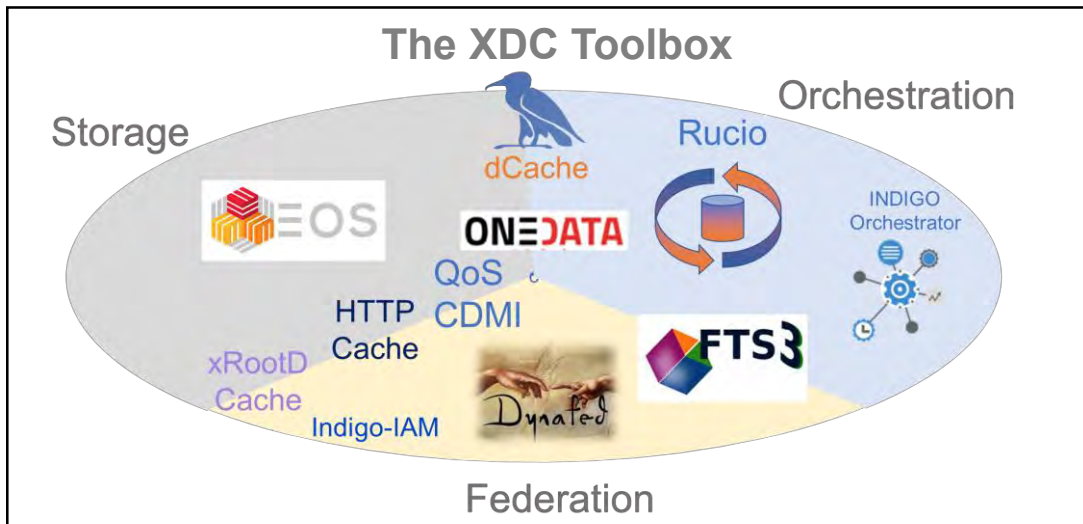




# General Architecture Definition

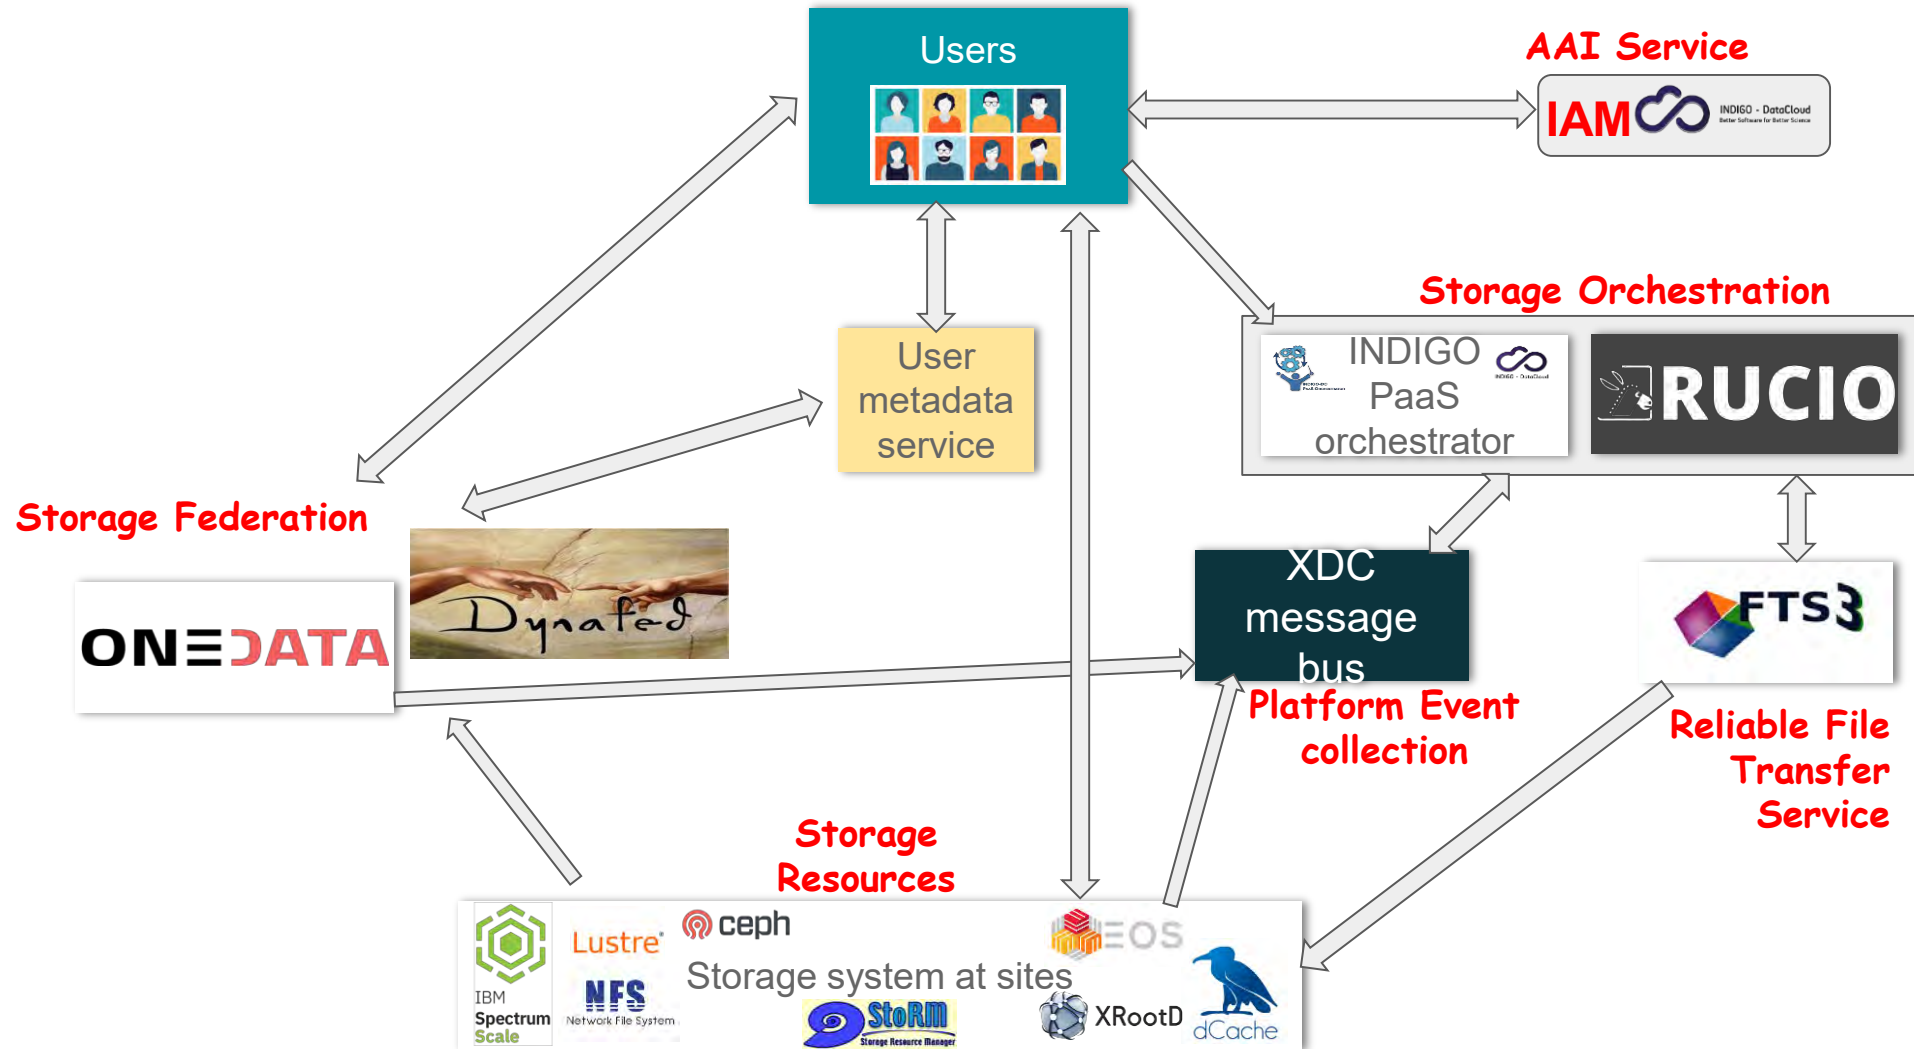
- ✗ XDC acts at all the e-infrastructure levels
  - Storage systems at sites
  - Federations of storage systems
    - regional and global
  - High level orchestration
  - User experience

- ✗ The “toolbox” was mapped in those levels to define the general architecture
  - Taking into account the user requirements





# XDC General Architecture



# Connection to external Entities



- ✗ DEEP-DataCloud
  - Connection to compute workloads
  - SQA common baseline



- ✗ DOMA (Data Organization, Management, Access)
  - DOMA “Access”
  - DOMA “QoS” (XDC Workgroup Leader)
  - DOMA “3<sup>rd</sup> Party Copy”



- ✗ ESCAPE – WP2 (XDC Task Leader)
  - Deployment of a European Data Lake



- ✗ RDA
  - Defining QoS vocabulary
  - Data Management workshop at the next plenary in Helsinki



- ✗ SNIA
  - Rendering CDMI to reflect QoS
  - XDC introduced at 2018 SNIA conference

# SQA baseline

"A set of Common Software Quality Assurance Baseline Criteria for Research Projects" - <http://hdl.handle.net/10261/160086>

## ✕ What?

- ➡ Set of conventions and recommendations for software development, aiming "to serve as a reference within the European research ecosystem related projects"

## ✕ Why?

- ➡ Enhance the visibility, accessibility and distribution of the produced software
- ➡ Promote code style standards => readability & reusability
- ➡ Reliable operation



The cover of the document features three logos at the top: Hybrid DataCloud (dosp), INISO - DataCloud (infinity symbol), and eXtreme DataCloud (eXtreme logo). The title "A set of Common Software Quality Assurance Baseline Criteria for Research Projects" is prominently displayed in the center. Below the title is an "Abstract" section stating the purpose of the document. At the bottom is a "Document Log" table.

### A set of Common Software Quality Assurance Baseline Criteria for Research Projects

#### Abstract

The purpose of this document is to define a set of quality standards, procedures and best practices to conform a Software Quality Assurance plan to serve as a reference within the European research ecosystem related projects for the adequate development and timely delivery of software products.

#### Document Log

| Issue | Date       | Comment             |
|-------|------------|---------------------|
| V1.0  | 31/01/2018 | First draft version |
| V2.0  | 05/02/2018 | Updated criteria    |

# XDC Components

## The Components



### ✂ Orchestration and Federation Components

- XDC Orchestrator
- INDIGO PaaS Orchestrator
- Flowable © (BPM)
- Rucio Data Management System



### ✂ Data Transfer and Data Federation technologies



- FTS, File Transfer Service,
- Dynafed, Data Federator, Onedata



ONE DATA

### ✂ Storage Systems

- dCache
- EOS
- StoRM





# XDC Orchestration Components



## ✕ INDIGO PaaS Orchestrator

- ➡ Based on INDIGO-DataCloud developments.
- ➡ Allows to coordinate complex deployments on hybrid clouds featuring advanced scheduling and federation capabilities
- ➡ Orchestrates compute resources and provides data-aware scheduling of jobs through data placement plugins (XDC extensions)
- ➡ Integrates with Rucio for data location and transfer orchestration (XDC developments)
- ➡ Operates with an professional BPM system. (Flowable)



## ✕ Flowable © (BPM)

- ➡ Provides a workflow and Business Process Management (BPM) platform for developers, system admins and business users



## ✕ Rucio


- ➡ Originally LHC ALTAS data management tools.
- ➡ Recently adopted by a growing number of other communities.
- ➡ Already provides interfaces to most XDC components



# XDC Transport Components

## ✕ FTS, File Transfer Service

- WLCG data transfer workhorse.
- Transfers around 1 Exabytes of WLCG data per year between hundreds of storage sites around the world.
- Performs request queueing and network shaping.
- Can be used as “micro service” or with GUI (WebFTS).
- Support X509 and token based authentication for endpoints.



## ✕ Dynafed, Data Federator

- Federates storage endpoints to a single root namespace.
- Supported Protocols: http/WebDAV, S3.
- Performs metadata prefetching.
- Provides location meta data to high level services.



# XDC Storage Components


## The Components



### dCache



- ➡ Open Source Storage system provided by DESY, Fermilab and NDGF.
- ➡ Handling 150 PBytes at more than 60 big data centers, including 7 WLCG Tier 1 centers.
- ➡ Supports industry standard data access and security protocols on top of a geo-aware multi tier storage stack.



### EOS



- ➡ Scalable storage running at CERN and elsewhere.
- ➡ Geo-aware management of hundred of PBs.
- ➡ HTTP interface.

### StoRM

- ➡ Provided by INFN/CNAF
- ➡ Engine providing multiple data transport and control protocols on top of GPFS and Lustre.

# XDC Storage Components

## ✕ INDIGO CDMI Reference Implementation

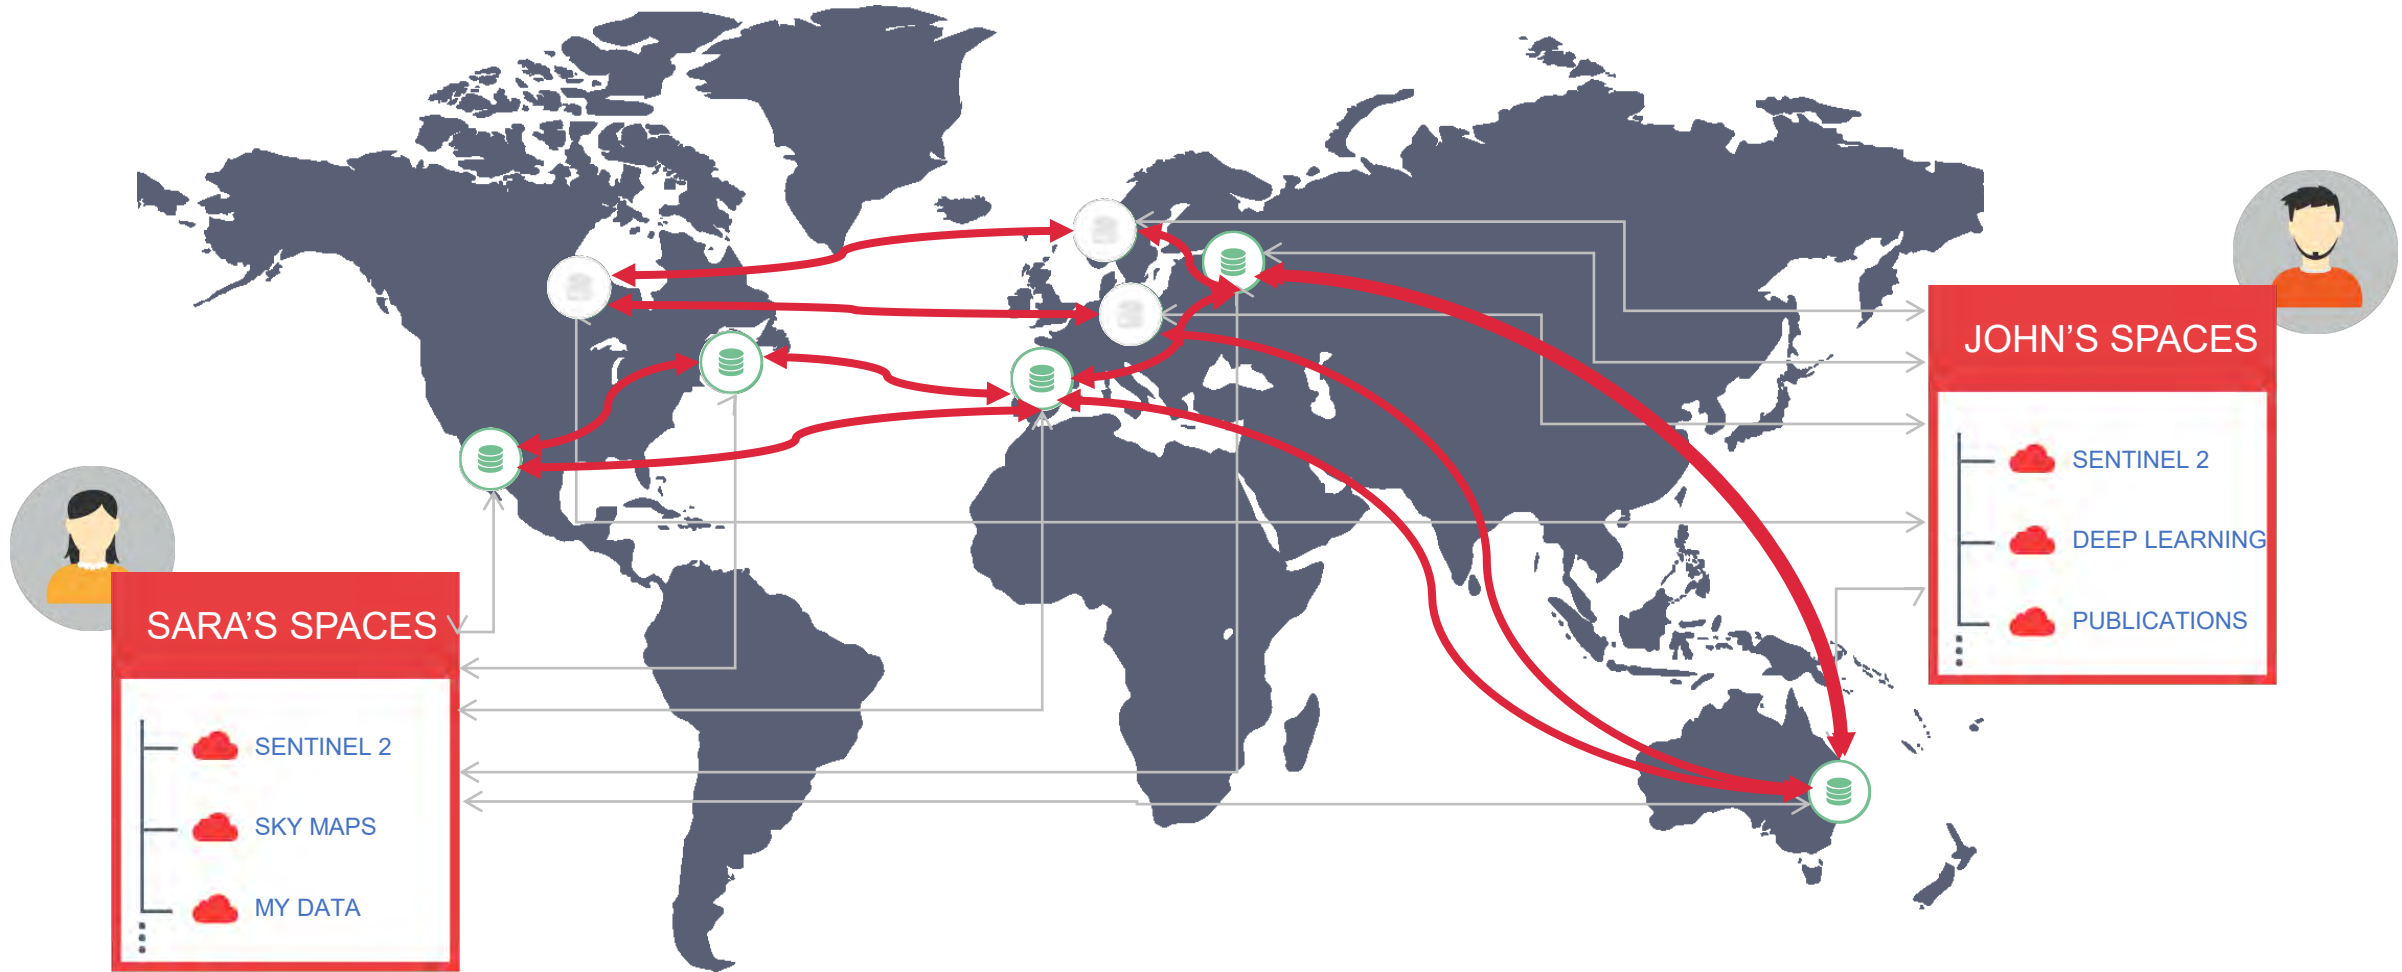
- ➡ INDIGO re-implementation of the SNIA CDMI reference implementation, now hosted by SNIA.
- ➡ Provided the CDMI protocol engine and forwards the requests to a plug-in system.
- ➡ Provides plug-ins for a REST protocol dialect as well as for CEPH and GPFS.

## ✕ XCache

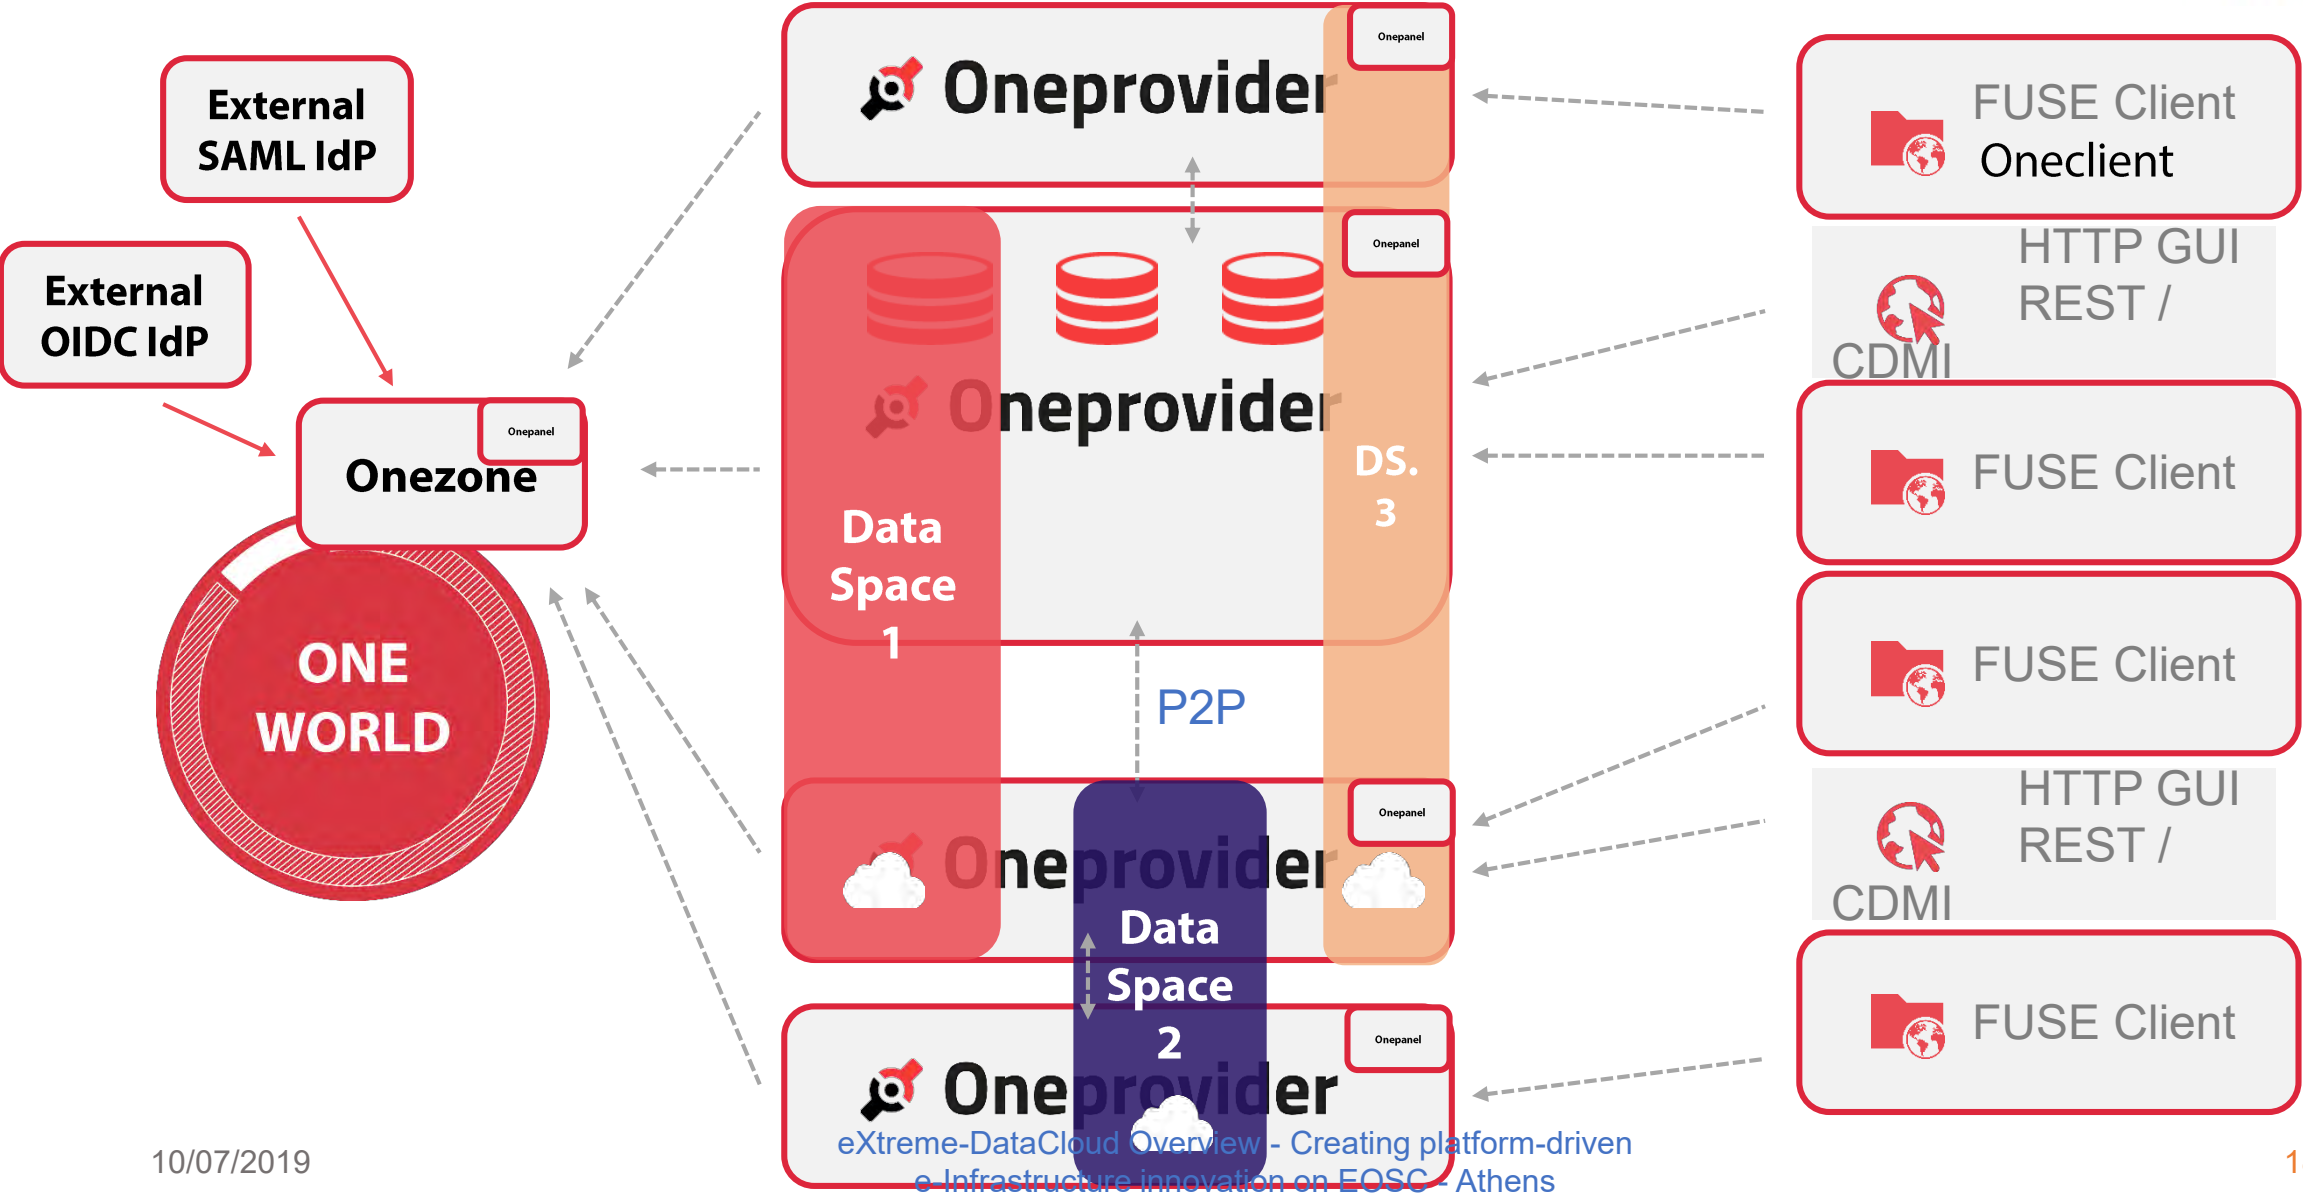
- ➡ Read-only, block-level data cache
- ➡ Deployed close to CPU to hide latency and reduce WAN traffic
- ➡ HTTP interface



# ONEDATA DISTRIBUTED DATA IN HYBRID CLOUDS



# ONEDATA SYSTEM ARCHITECTURE



# First XDC Release

## ✗ Involved tools

- CachingOnDemand
- dCache
- Dynafed
- EOS
- FTS, GFAL
- Onedata
- PaaS Orchestrator plugin
- TOSCA types & templates plugin

## ✗ Key technical highlights

- OpenIDConnect support for token based authentication
- New QoS types integration and support in dCache, FTS, GFAL
- Orchestrator integration with other components
- Performance improvements in Onedata
- Support for groups and roles in Onedata
- EOS-dCache integration
- Caching systems instantiation
- Storage events notification in dCache
- EOS caching with XCache for geographic deployment
- EOS external storage adoption

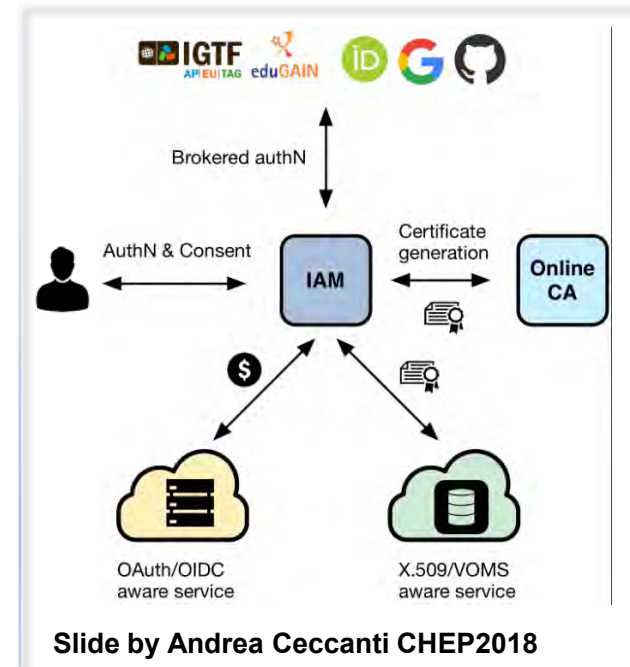
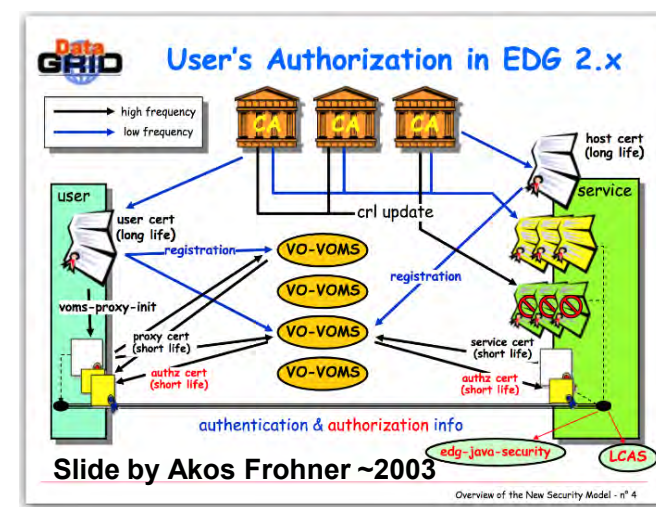


<https://releases.extreme-datacloud.eu/en/latest/releases/pulsar/index.html>

# First XDC Release

## Key technical highlights

- **OpenIDConnect support for token based authentication**
- new QoS types integration and support in dCache, FTS, GFAL
- Orchestrator integration with other components
- Performance improvements in Onedata
- Support for groups and roles in Onedata
- EOS-dCache integration
- Caching systems instantiation
- Storage events notification in dCache
- EOS caching with XCache for geographic deployment
- EOS external storage adoption



Sign Up With Google

Sign Up With Facebook

Sign Up With Twitter

LOGIN

Email

☒ Remember Me

Password

[Forgot Password?](#)



# First XDC Release

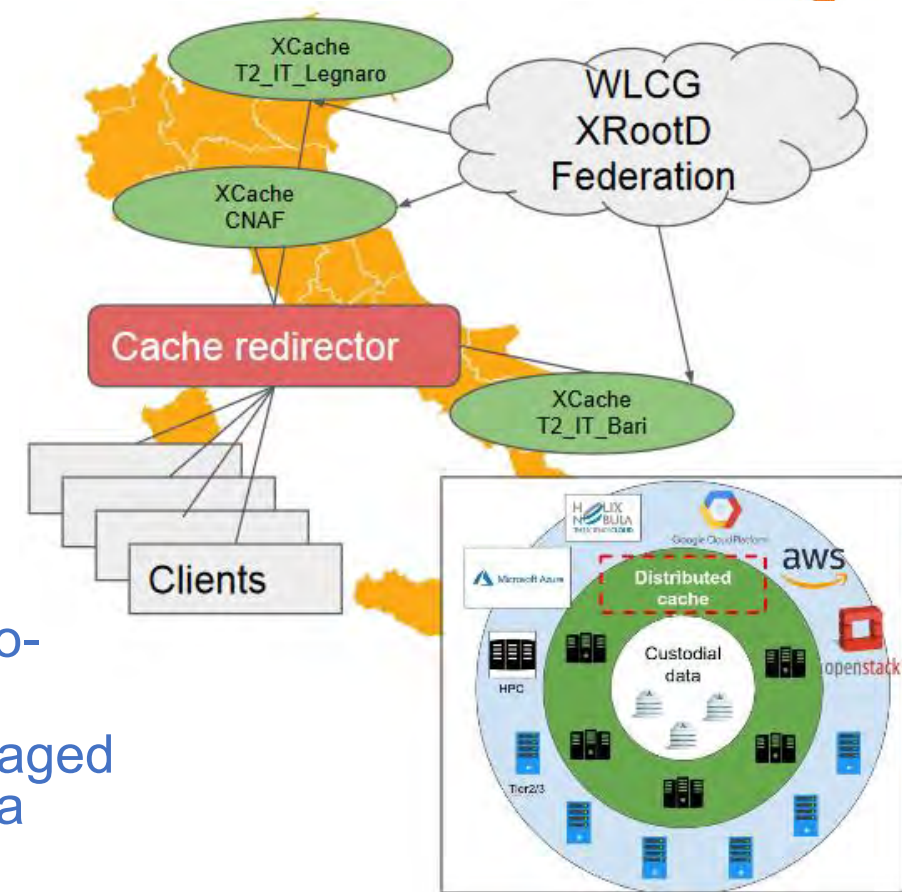
## ✕ Key technical highlights

- OpenIDConnect support for token based authentication
- new QoS types integration and support in dCache, FTS, GFAL
- Orchestrator integration with other components
- Performance improvements in Onedata
- Support for groups and roles in Onedata
- EOS-dCache integration
- **Caching systems instantiation**
- Storage events notification in dCache
- EOS caching with XCache for geographic deployment
- EOS external storage adoption

- Deployment of Geo-distributed caches
- Network of unmanaged storage for hot data
- On-demand cache resources

**Based on xRootD/xCache**

See D.Ciangottini talk on “Integration of the Italian cache federation within CMS computing model”:  
<https://indico4.twgrid.org/indico/event/8/session/23/contribution/45>




Slide © Diego Ciangottini


# First XDC Release

## Key technical highlights

- OpenIDConnect support for token based authentication
- new QoS types integration and support in dCache, FTS, GFAL
- Orchestrator integration with other components
- Performance improvements in Onedata**
- Support for groups and roles in Onedata**
- EOS-dCache integration
- Caching systems instantiation
- Storage events notification in dCache
- EOS caching with XCache for geographic deployment
- EOS external storage adoption



Onedata Transparent POSIX File System  
Processing transparently cached data - 37GBytes/sec

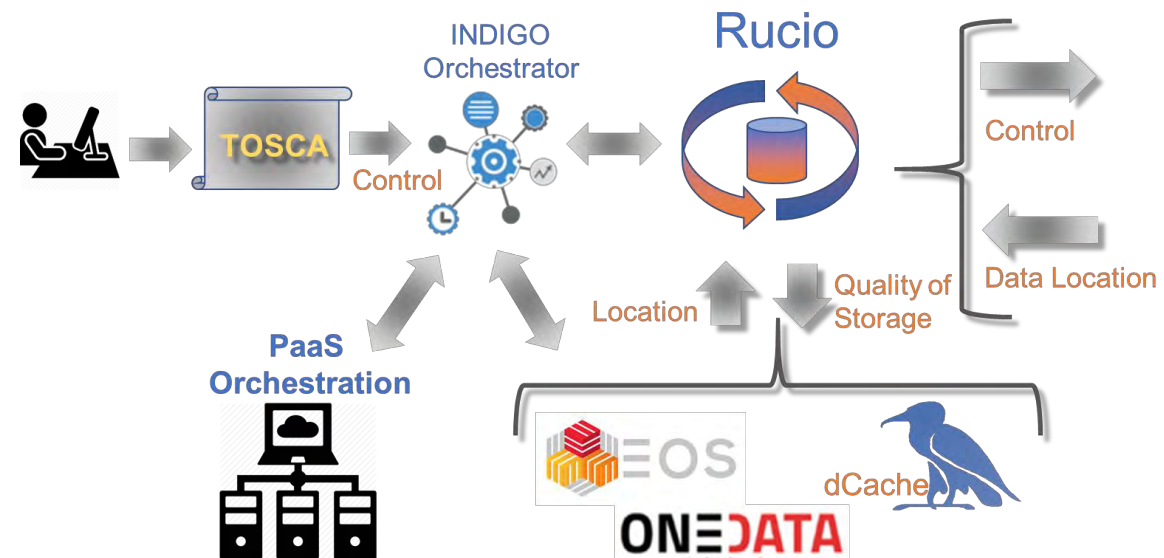


- Data Transfer Mesh
- 3 Oneproviders connected by 20+Gbit/s links
- Transfer data between all them
- Single VM Node per Provider
- Linear scalability

# First XDC Release

## Key technical highlights

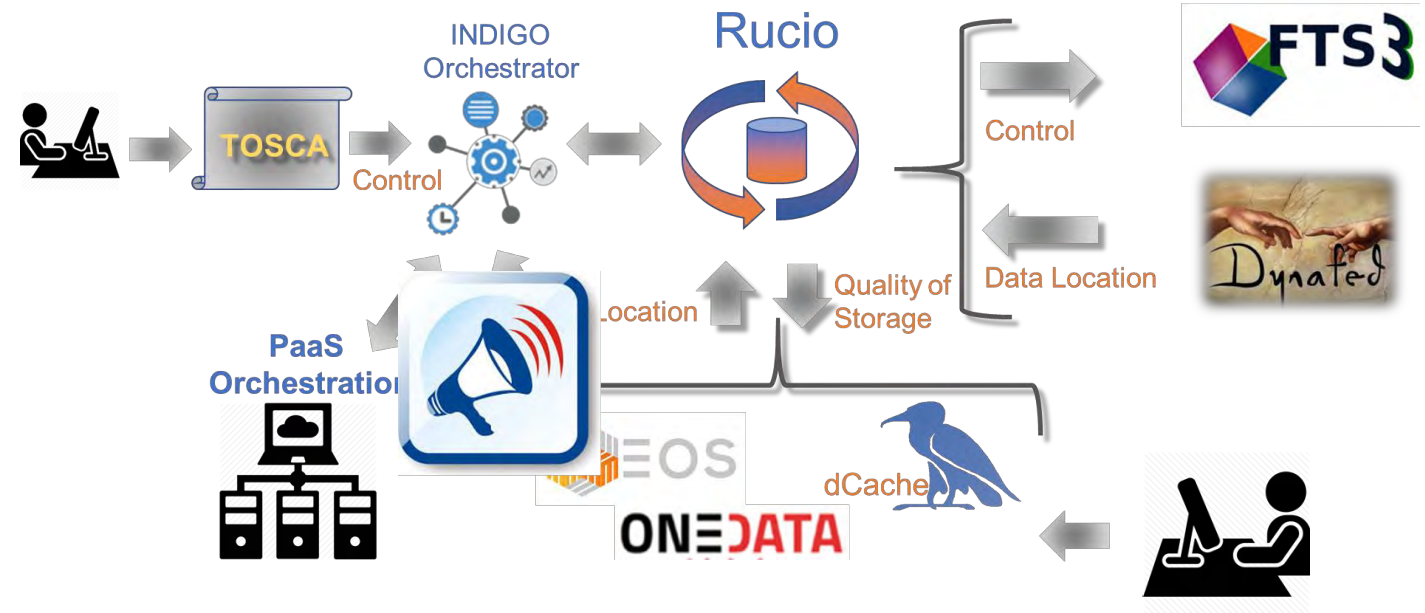
- OpenIDConnect support for token based authentication
- new QoS types integration and support in dCache, FTS, GFAL
- Orchestrator integration with other components
- Performance improvements in Onedata
- Support for groups and roles in Onedata
- EOS-dCache integration
- Caching systems instantiation
- Storage events notification in dCache**
- EOS caching with XCache for geographic deployment
- EOS external storage adoption



# First XDC Release

## Key technical highlights

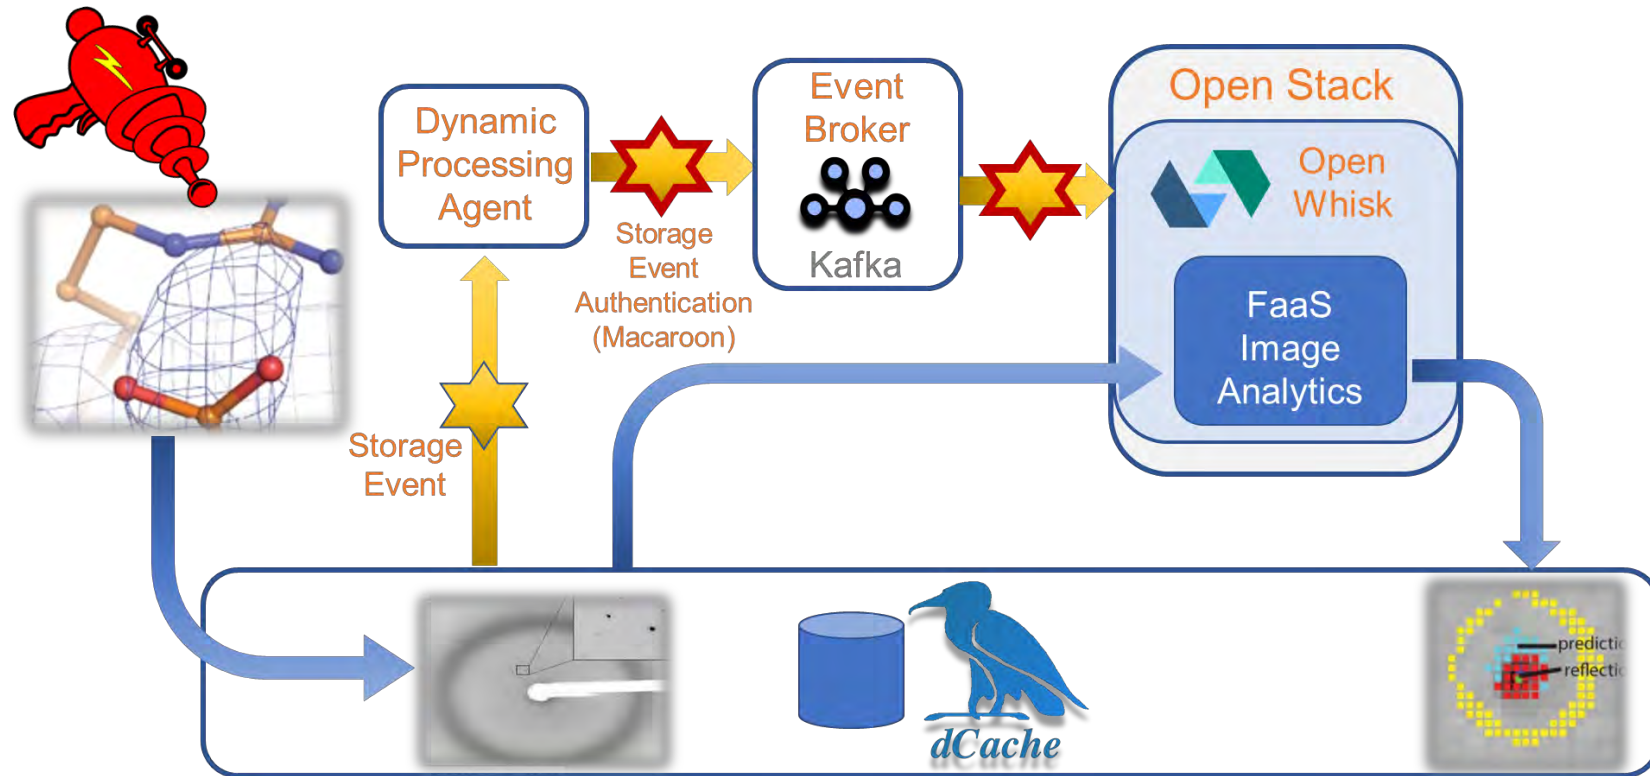
- OpenIDConnect support for token based authentication
- new QoS types integration and support in dCache, FTS, GFAL
- Orchestrator integration with other components
- Performance improvements in Onedata
- Support for groups and roles in Onedata
- EOS-dCache integration
- Caching systems instantiation
- Storage events notification in dCache**
- EOS caching with XCache for geographic deployment
- EOS external storage adoption





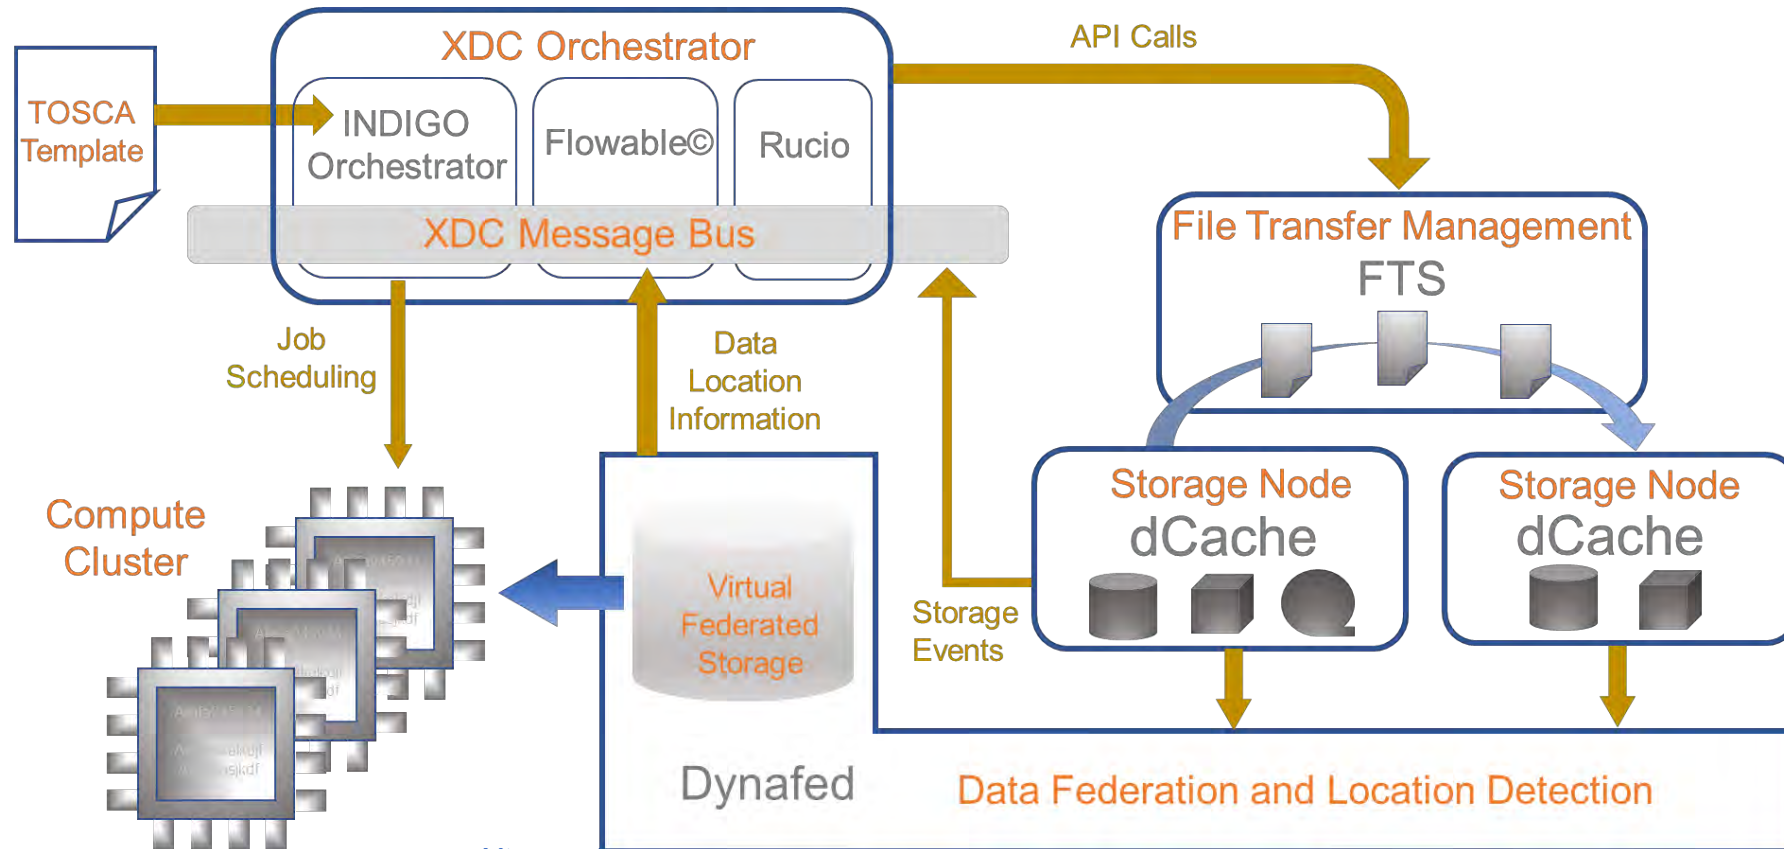
# XFEL Use Case in XDC

- ✗ The XFEL UseCase is driving the developments on storage events notifications support
  - A reference implementation is done using dCache as backend
- ✗ Refer to the Patrick's presentation:
  - <https://indico4.twgrid.org/indico/event/8/session/15/contribution/9>

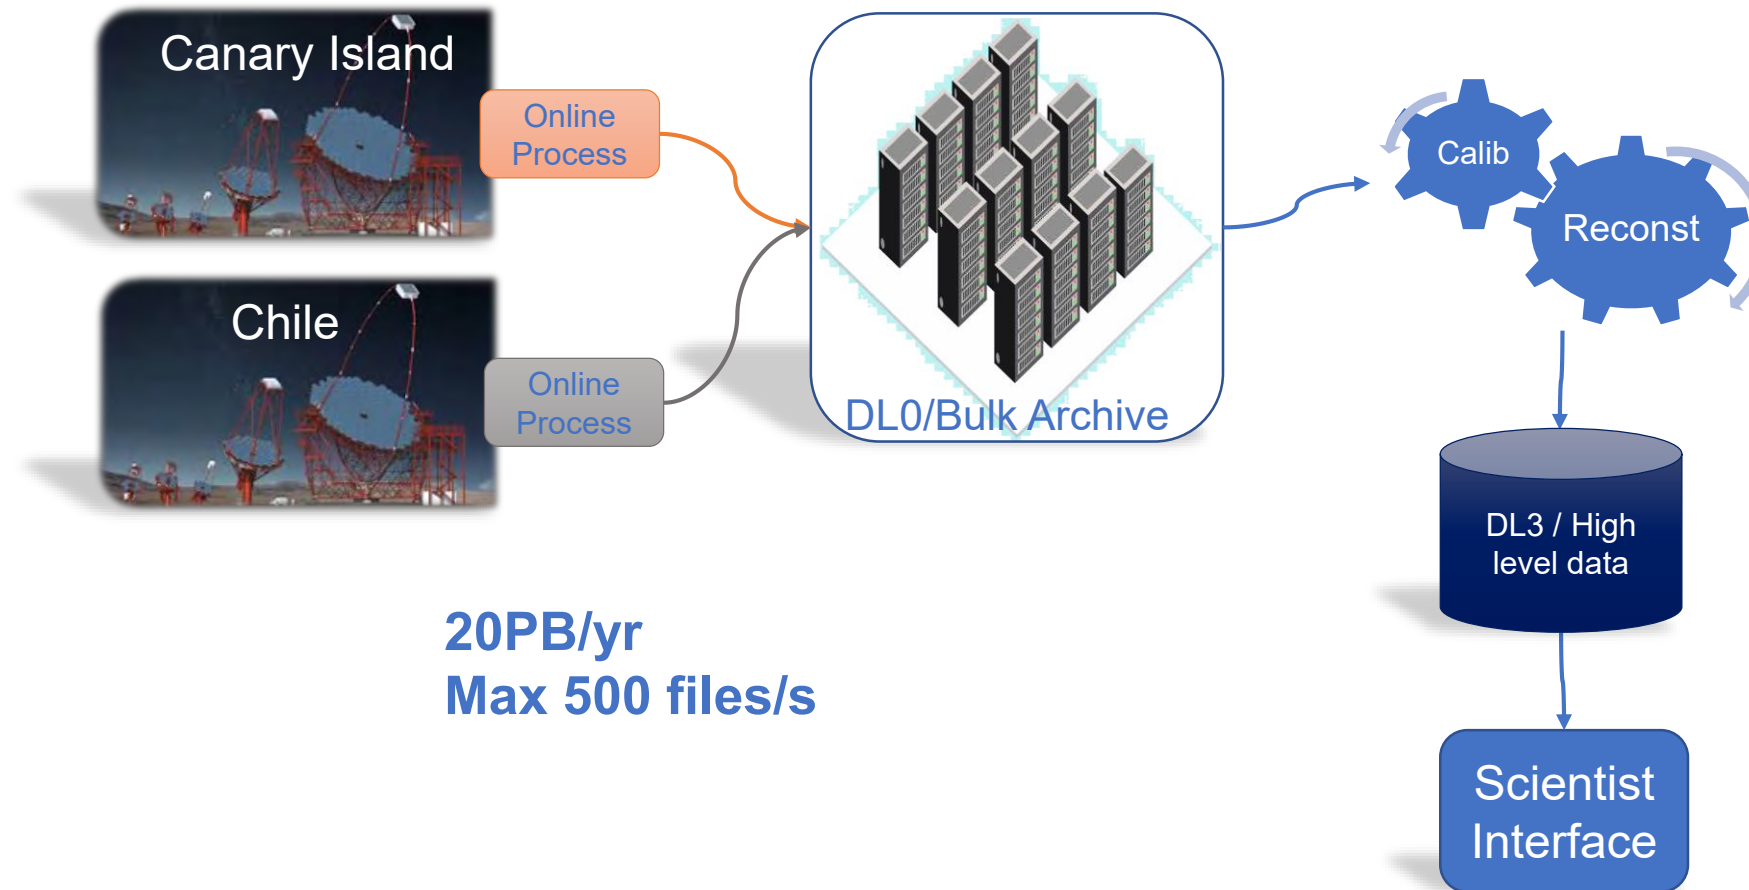


# XFEL Use Case in XDC

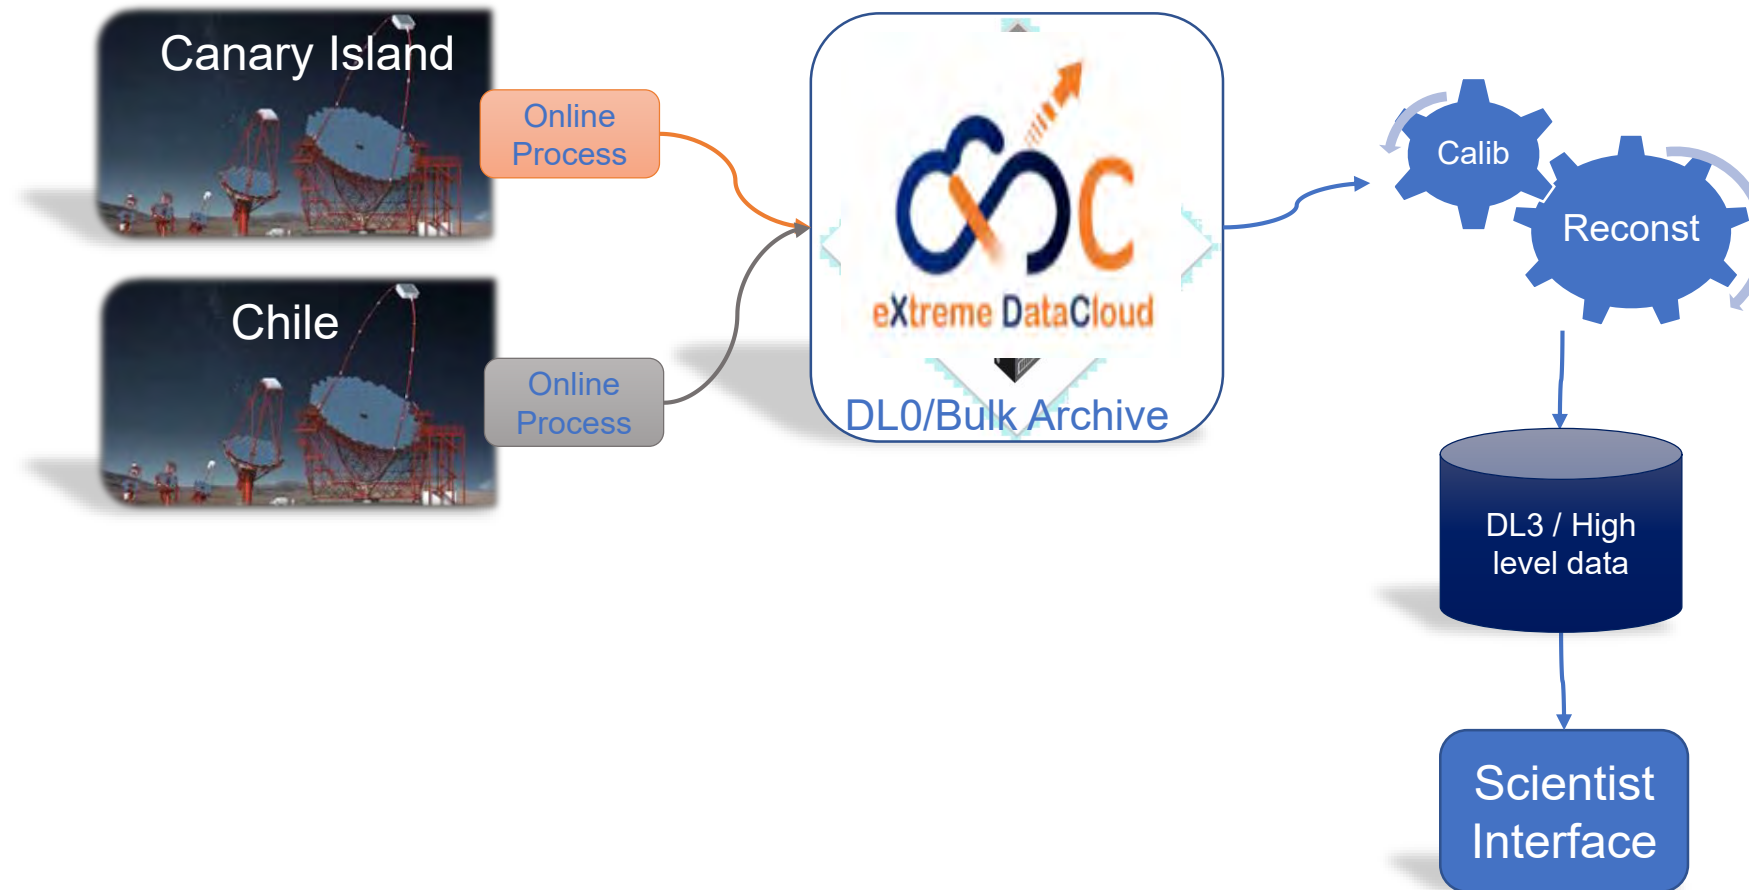
- ✗ The XFEL UseCase is driving the developments on storage events notifications support
  - A reference implementation is done using dCache as backend
- ✗ Refer to the Patrick's presentation:
  - <https://indico4.twgrid.org/indico/event/8/session/15/contribution/9>



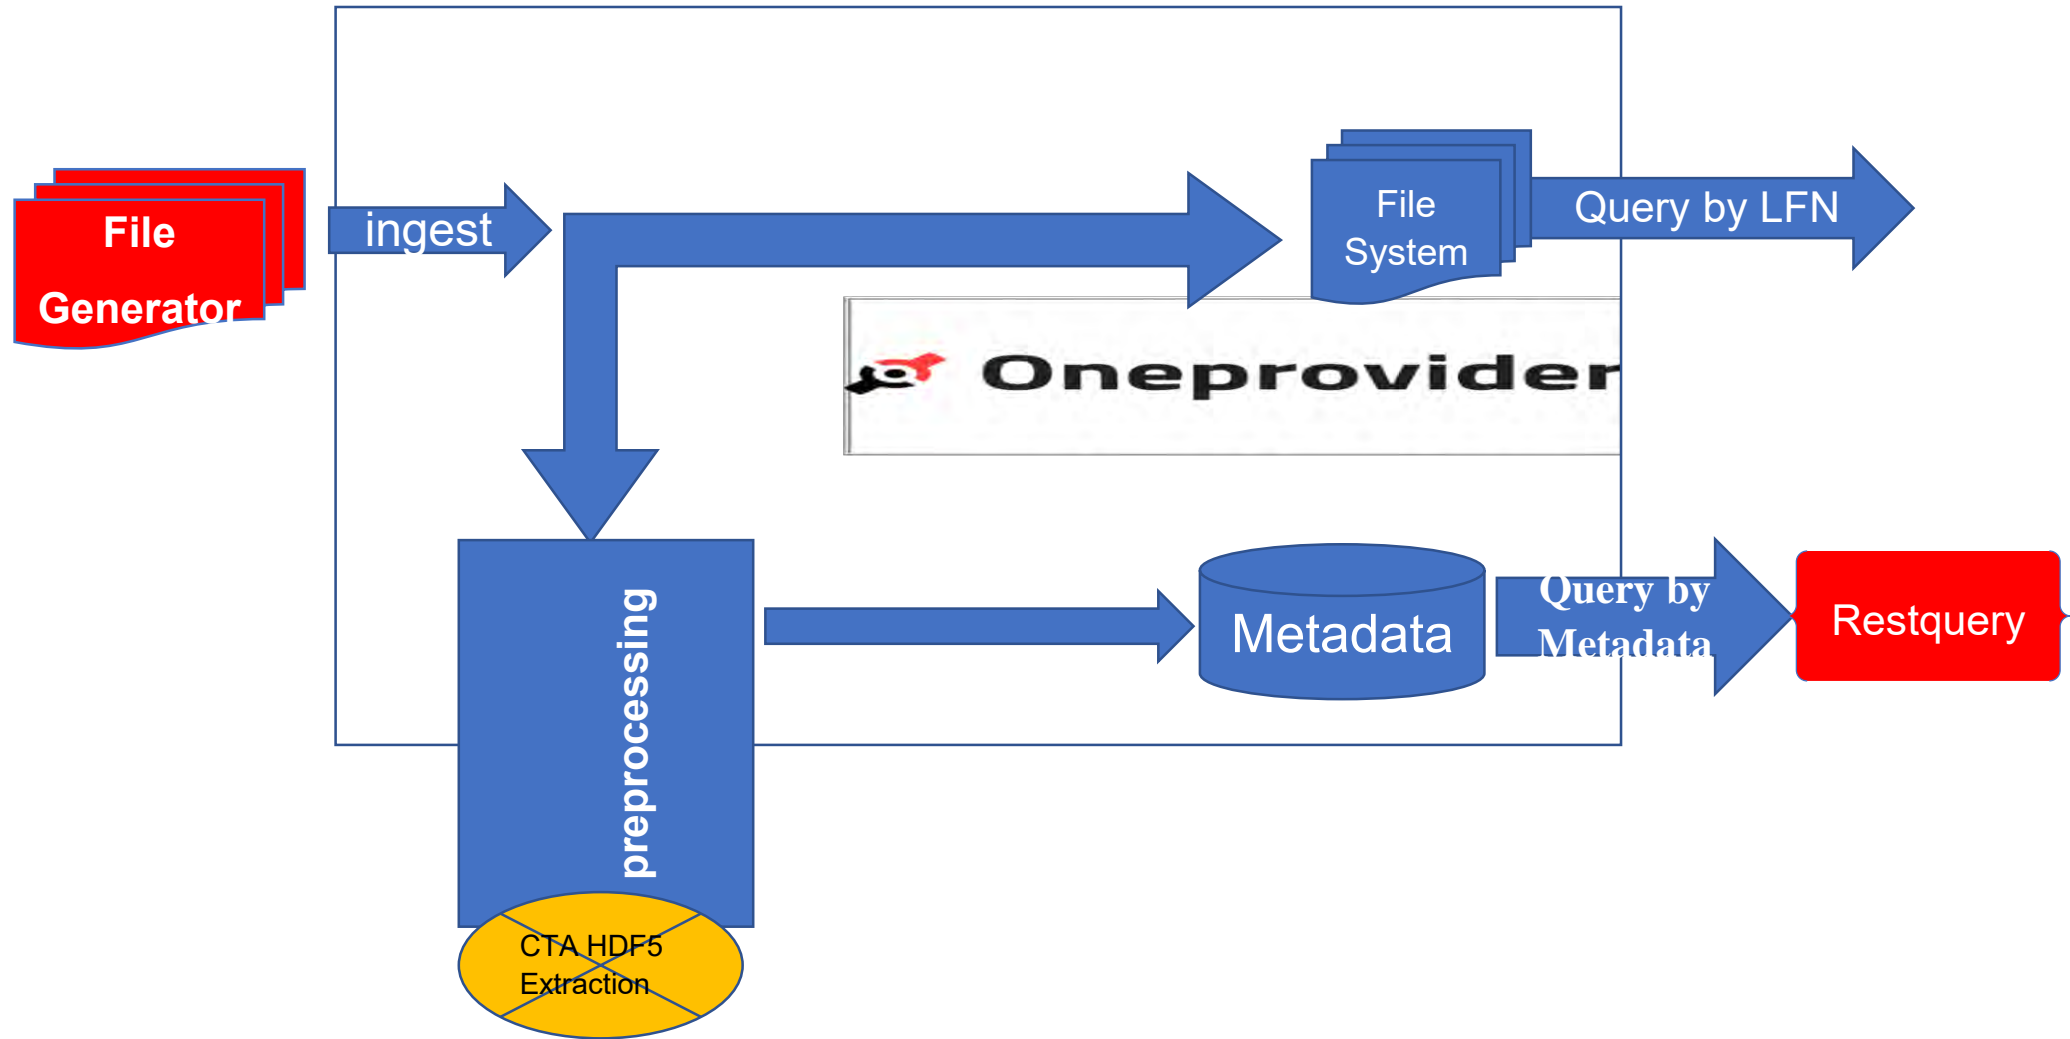
# CTA Use Case Workflow in XDC



# CTA Use Case Workflow in XDC

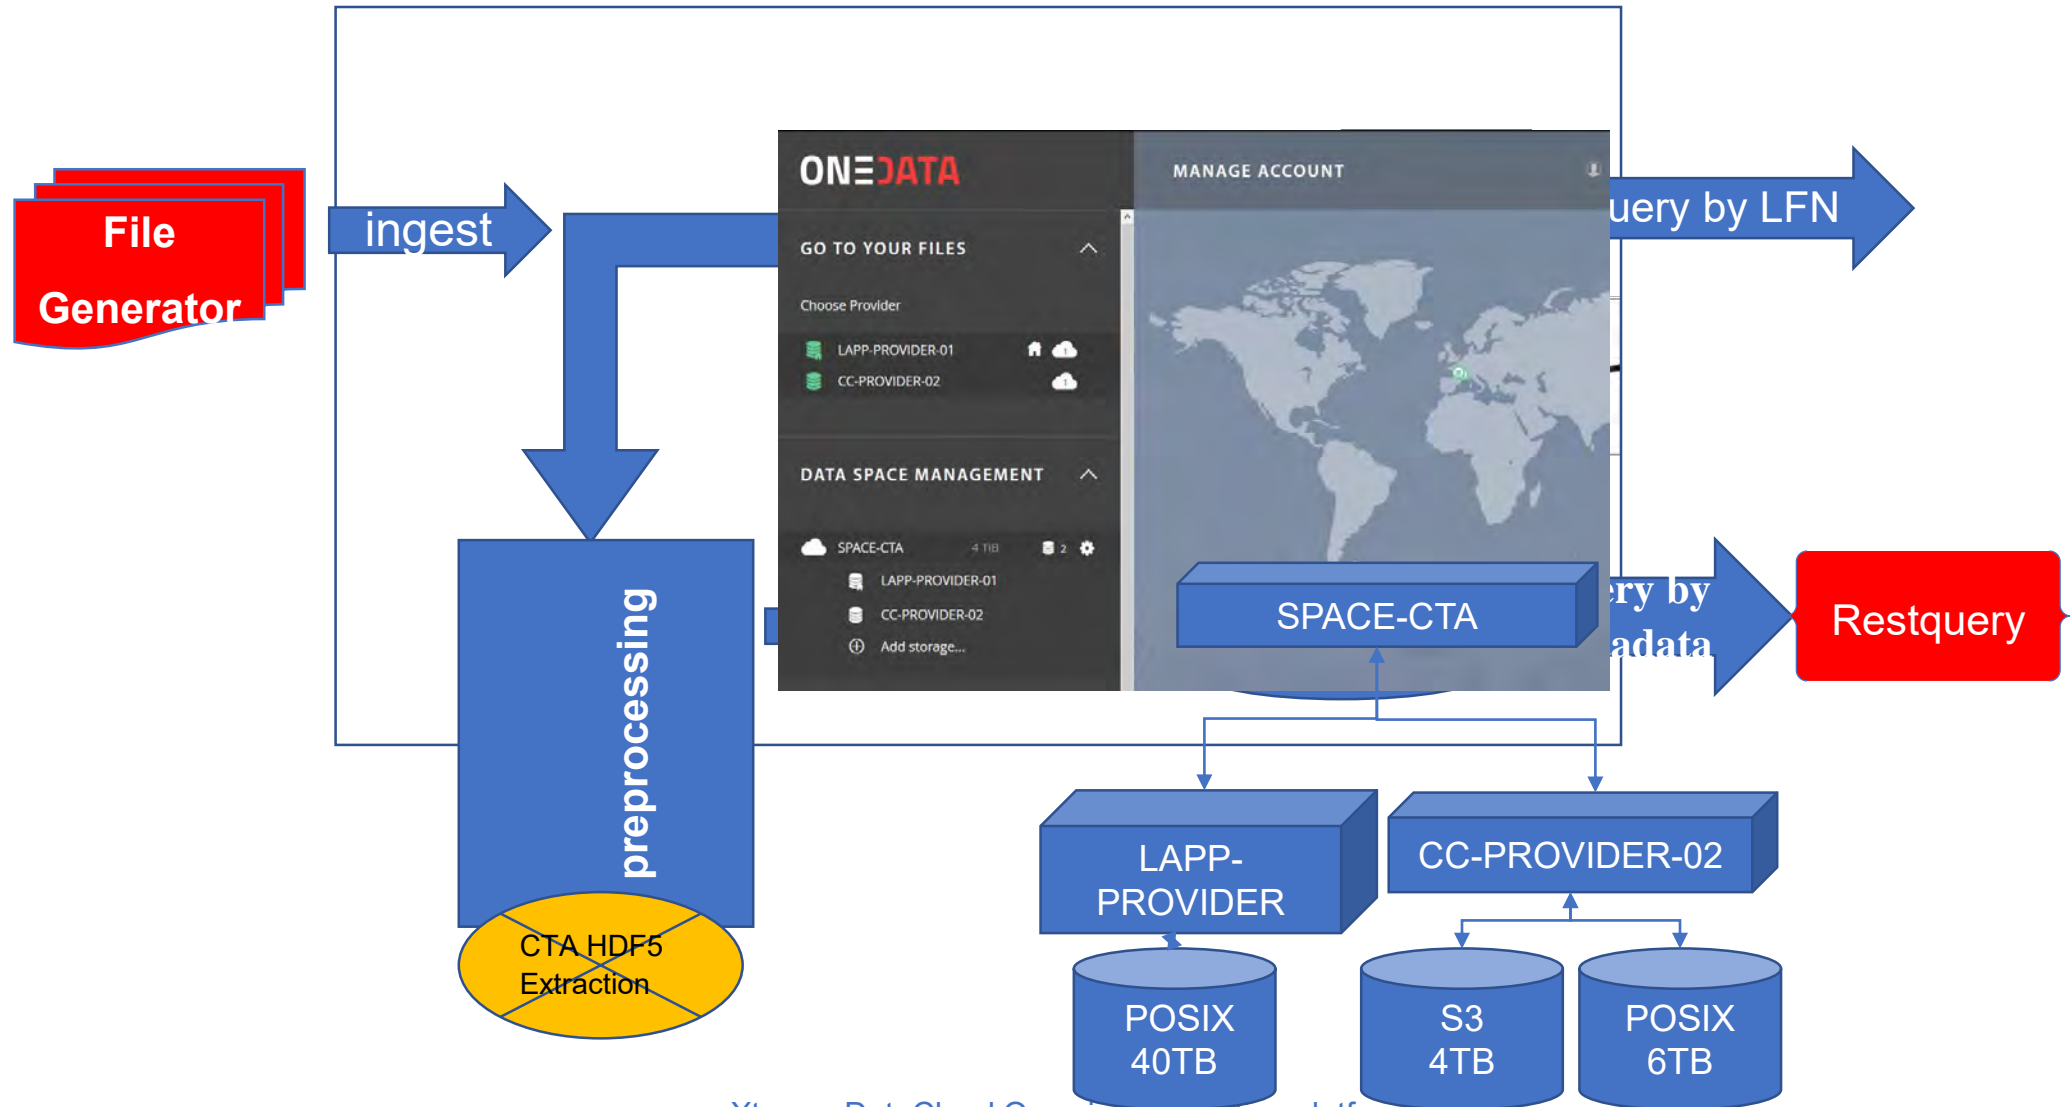


# CTA Use Case Workflow in XDC

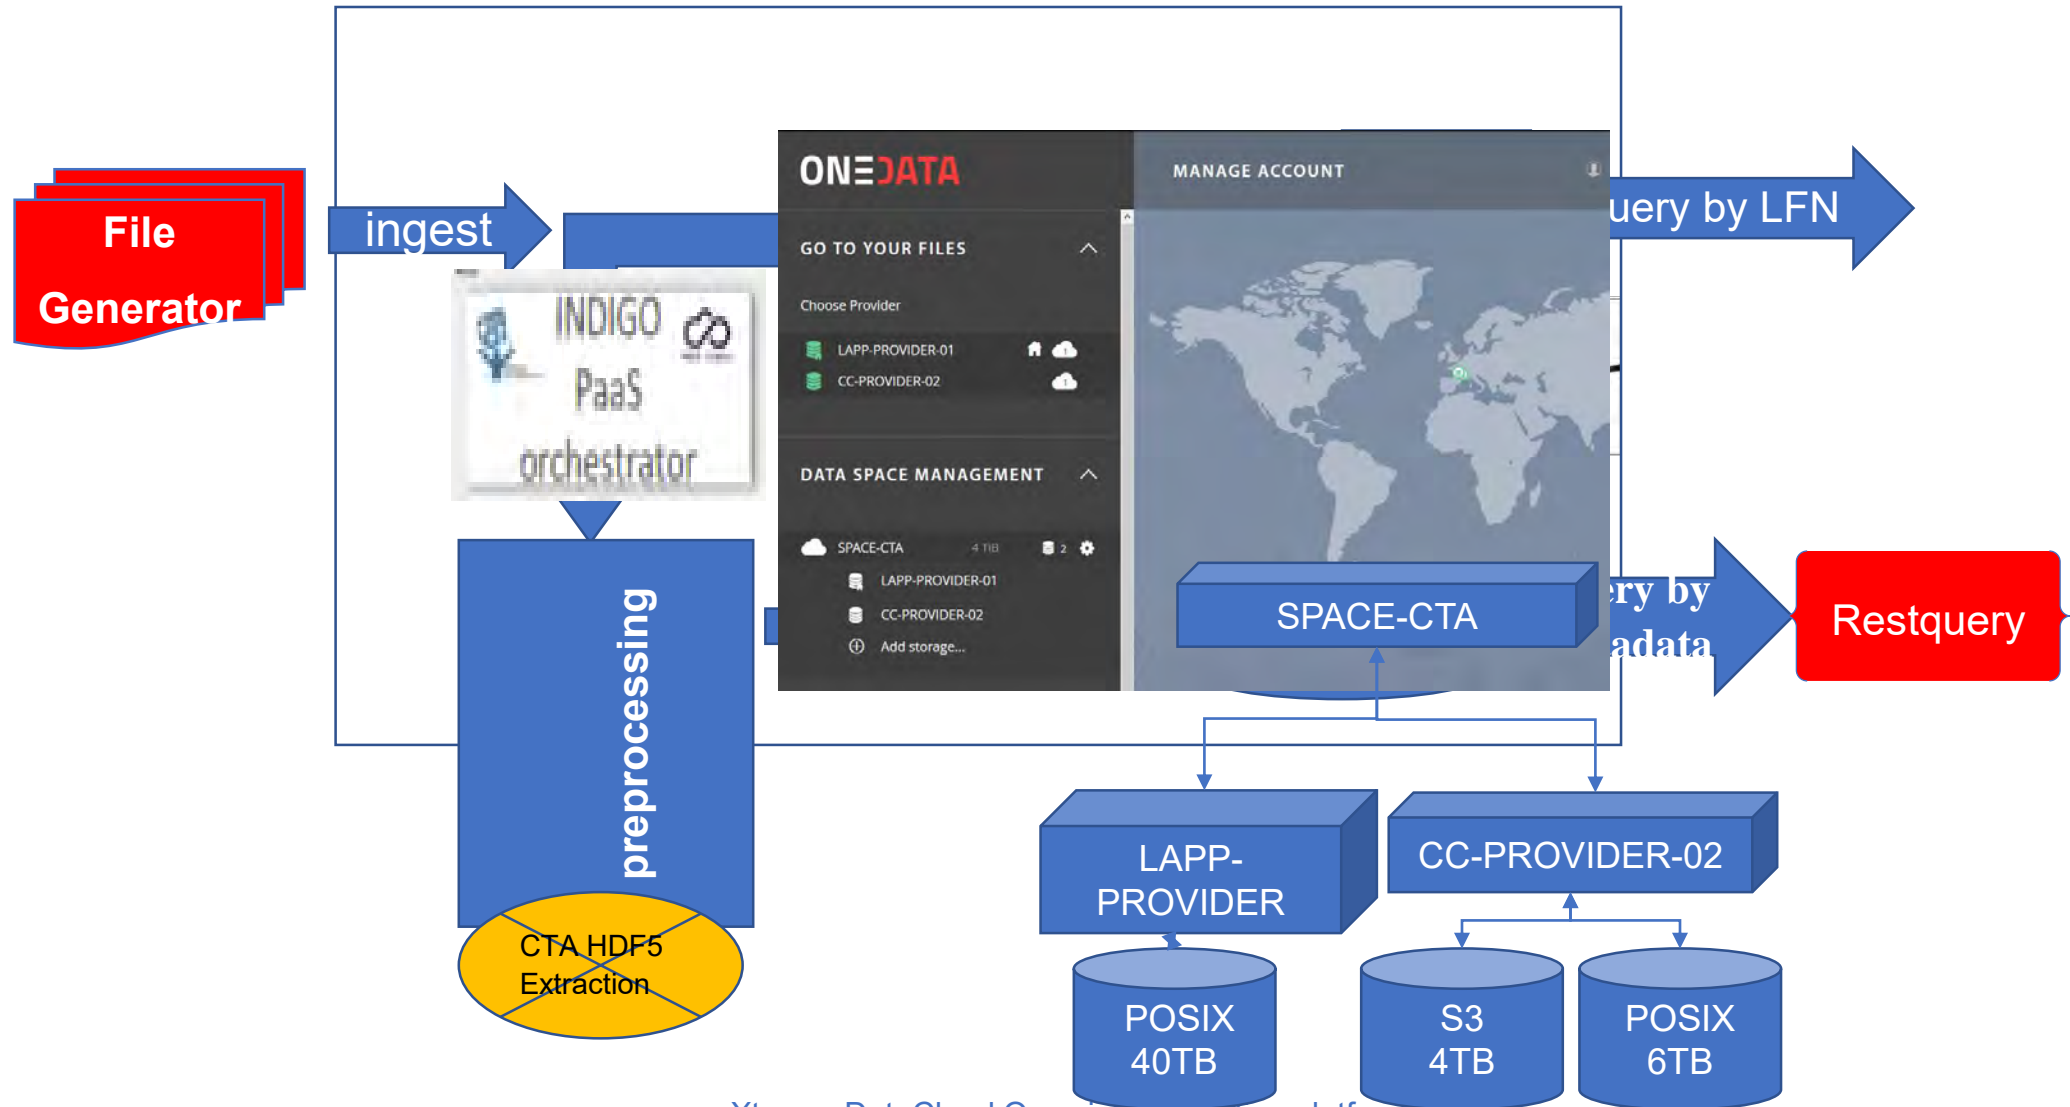




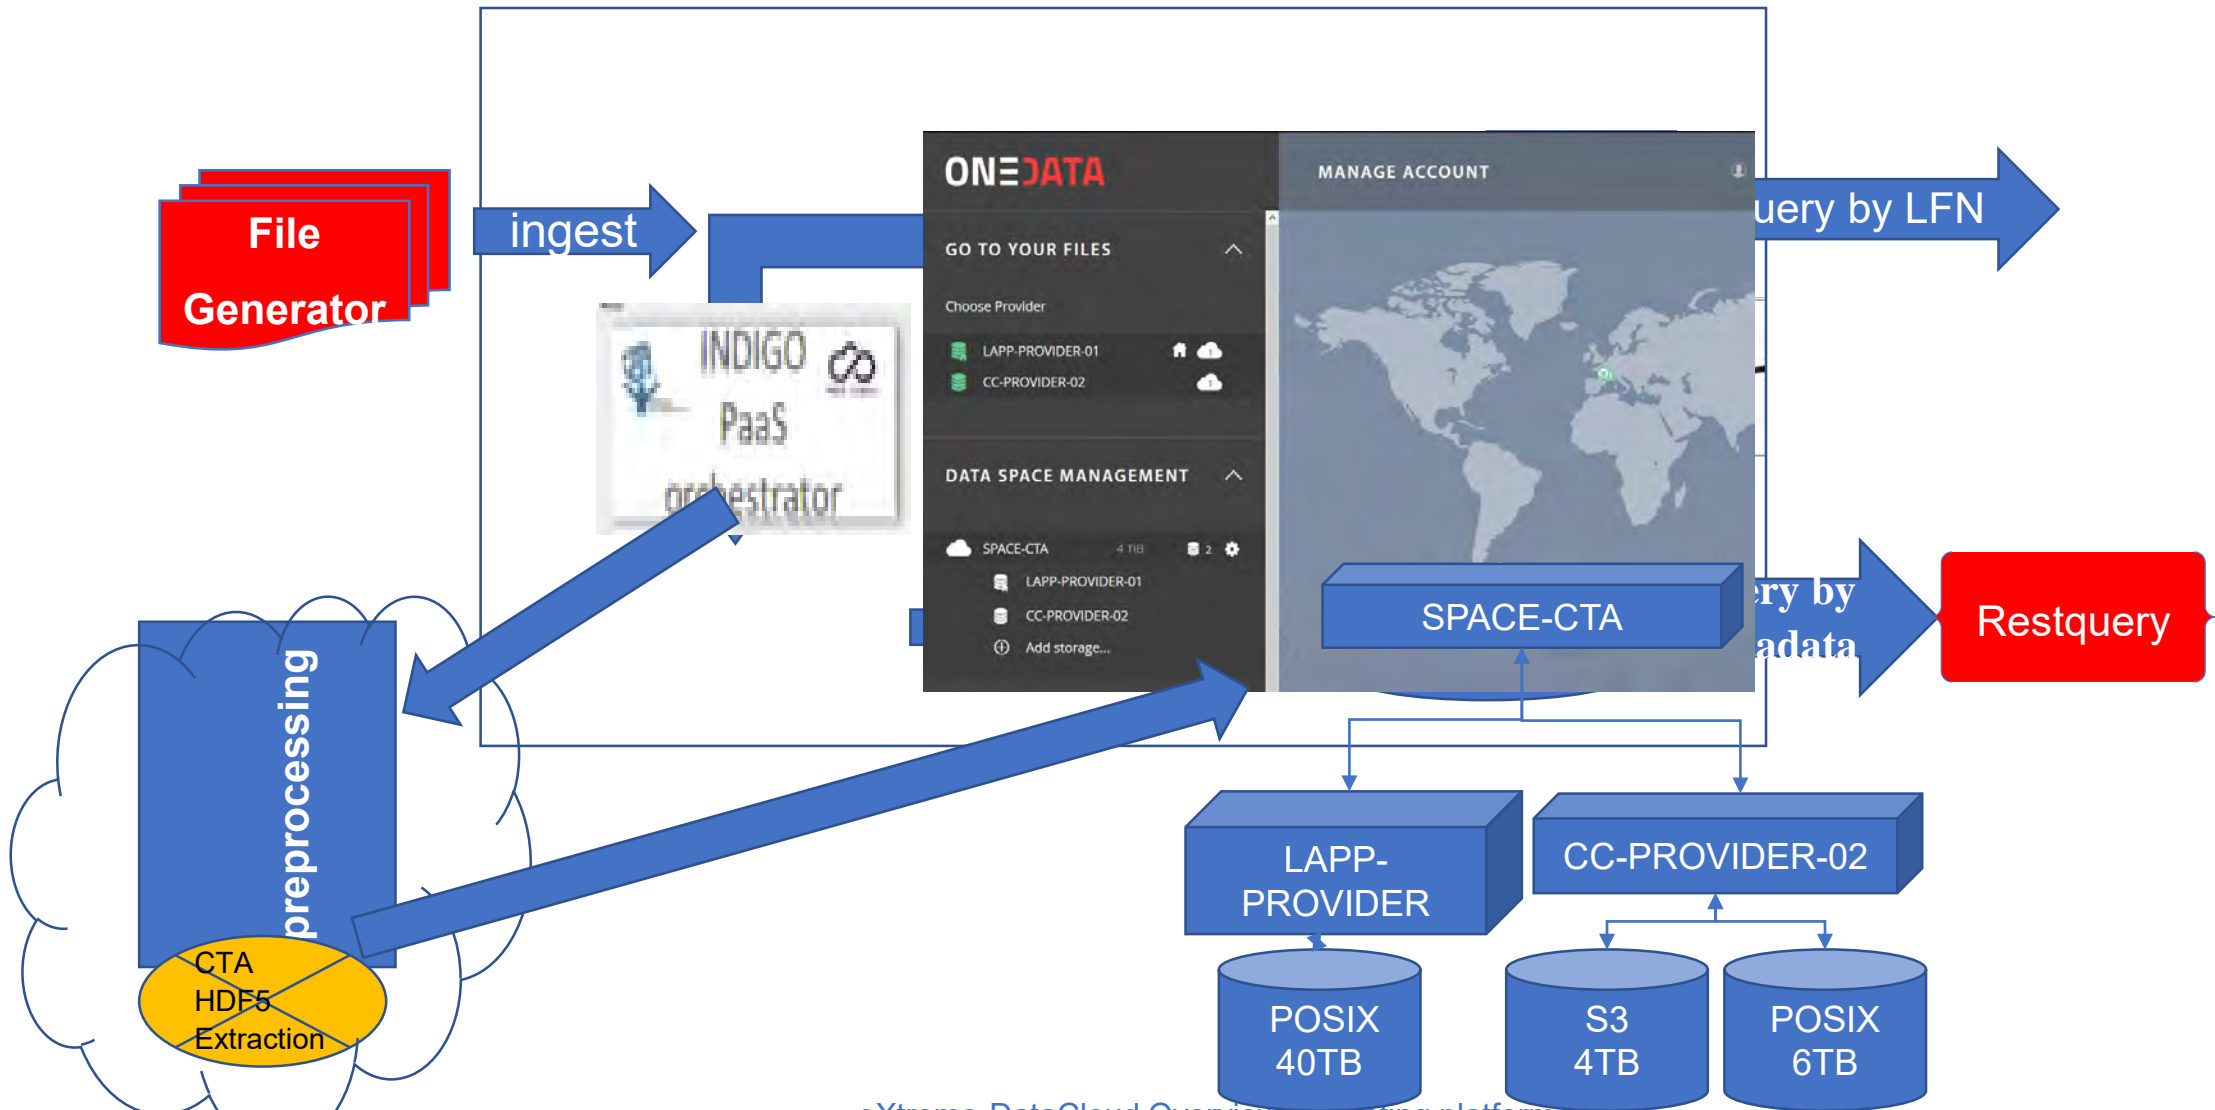
# CTA Use Case Workflow in XDC




# CTA Use Case Workflow in XDC



# CTA Use Case Workflow in XDC

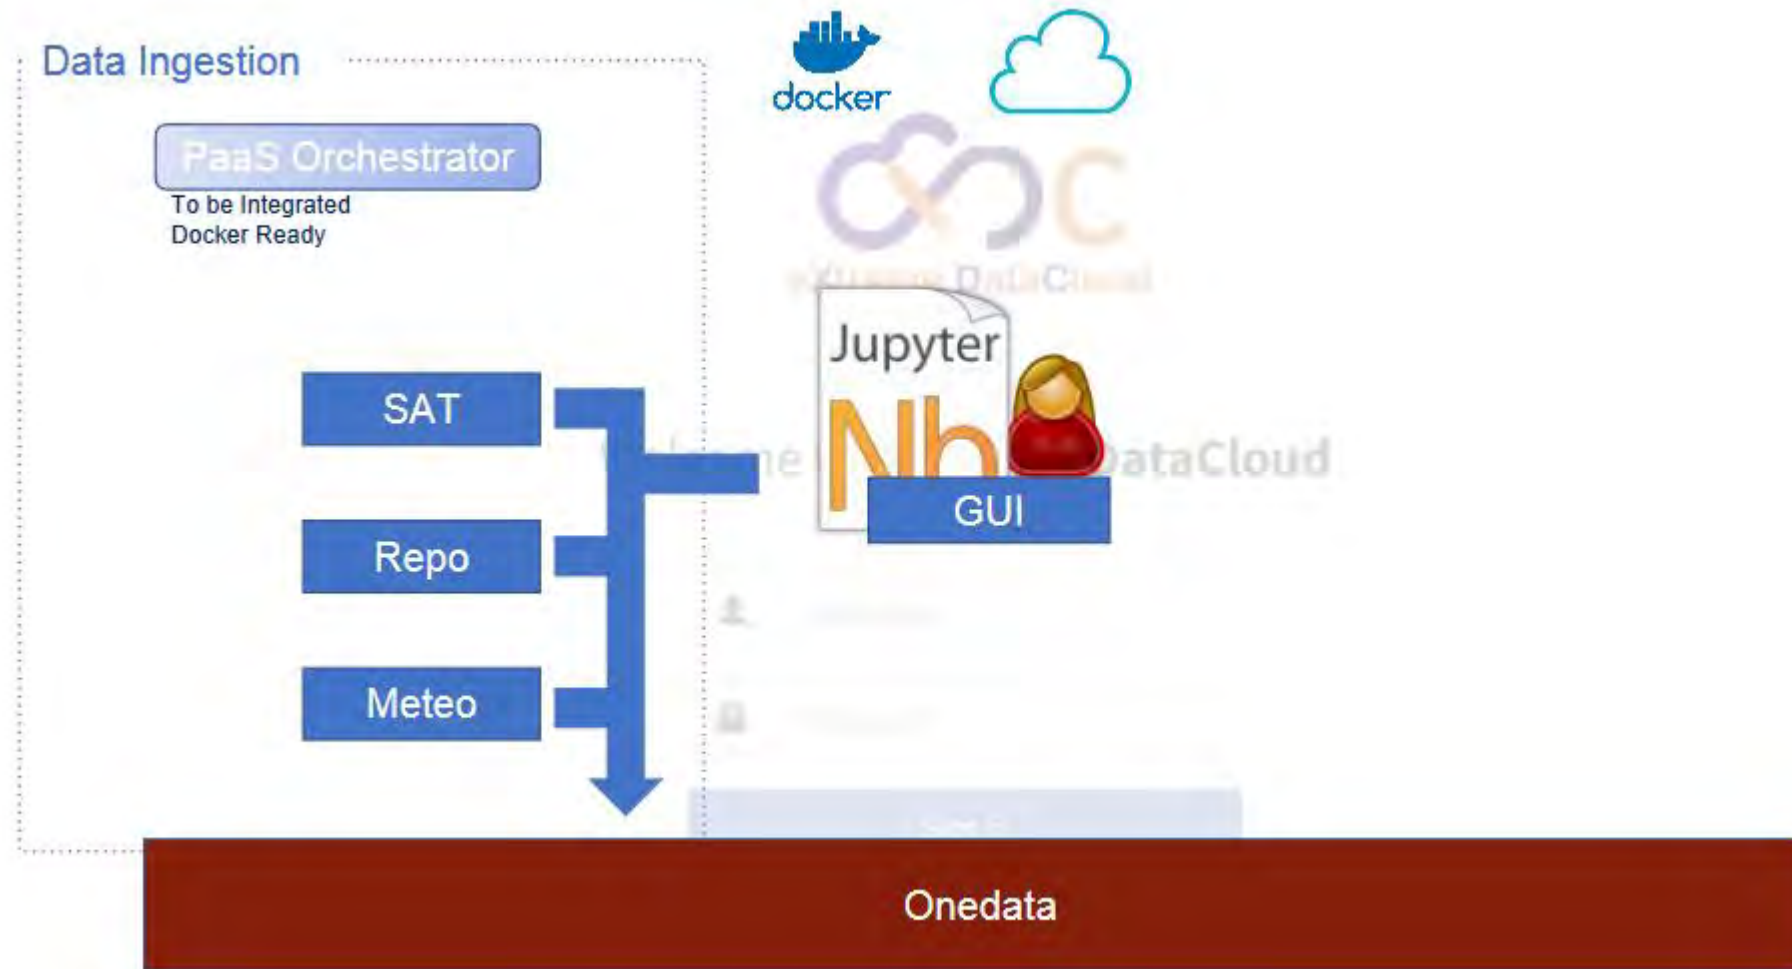


# LifeWatch Use Case in XDC



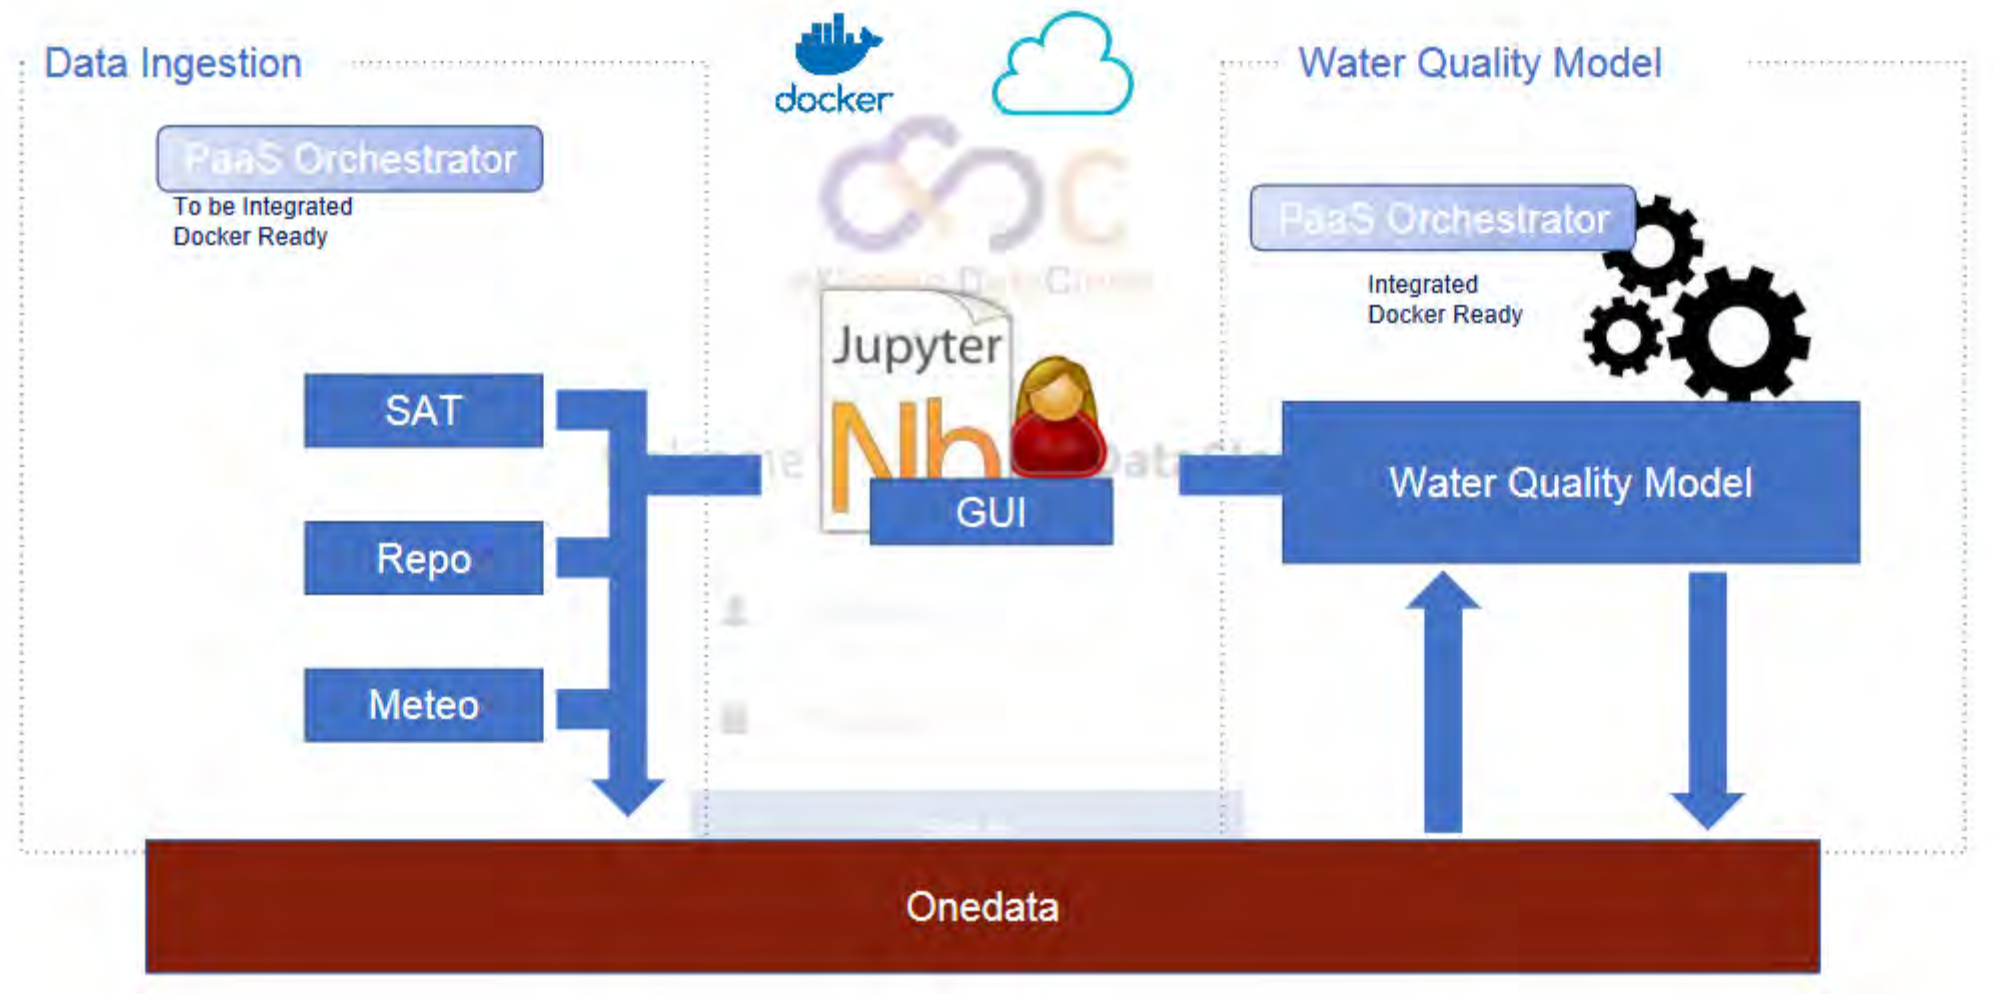
- **Objectives:** Integrate different and heterogeneous data sources: **satellite data**, **real-time monitoring system based on sensors, observations, and meteorological data** to feed the **hydrological and water quality models**, thus automating modeling and prediction of water quality.
- **XDC Services Requirements:**
  - **XDC IAM**
  - **Onedata:**
    - Onedata Attachment
    - Onedata Discovery
  - **PaaS Orchestrator**

# LifeWatch Use Case in XDC





# LifeWatch Use Case in XDC



# Use cases vs tool matrix

|                  | Problem   | Goal  | XDC services                                      | Status  |
|------------------|---|---|---|---|
| <b>Lifewatch</b> | Data Life Cycle Management of data related to Water Quality involving heterogeneous data sources              | Integrate data sources and different types of modelling tools                                 | Onedata<br>PaaS Orchestrator                      | Data Sources integrated<br>Analysis with Orchestrator |
| <b>CTA</b>       | Complex and Big Data management in a distributed environment. Data quality Assurance.                         | Integration of tools for FAIR Data Management and user access control                         | Onedata<br>QoS<br>PaaS Orchestrator               | Metadata extraction and storing/attachment            |
| <b>ECRIN</b>     | Distributed information about clinical studies and data objects across different registries and repositories. | Single environment to find data objects across repositories and registries, based on metadata | Onedata   | Sources integration.<br>Metadata harmonization        |
| <b>WLCG</b>      | Growing needs on storage space  | Reduce costs, resource aggregation, smart data allocation                                     | QoS<br>Xcache<br>EOS                              | QoS and http caching implemented                      |
| <b>XFEL</b>      | Complex Data management in a distributed and heterogeneous environment.                                       | Data lifecycle management. Processing and analytics   | QoS<br>PaaS Orchestrator<br>dCache<br>Message BUS | Orchestration configured based on events              |

# XDC Main Releases

- ✕ A second major release is foreseen before the end of the project
  - XDC Message bus implementation
  - full orchestration
  - finalize integration of RUCIO
  - secure storage in Onedata
  - finalize the ECRIN Use Case
  - complete caching reference workflows with HTTP based systems

|       | Release Date | End of Full Updates | End of Standard Updates | End of Security Updates & EOL |
|-------|--------------|---------------------|-------------------------|-------------------------------|
| XDC 1 | Jan 2018     | May 2019            | Sep 2019                | Nov 2019                      |
| XDC 2 | Oct 2019     | Apr 2020            | Jul 2020                | Sep 2020                      |

XDC products can be downloaded from XDC repositories or from each components upstream repositories after they have been pushed back

# Software quality process and testbeds



✗ XDC defined a rigorous process for ensuring software quality

- ➡ Definition of roles
- ➡ SQA Policies and Procedures
- ➡ Development testbed
- ➡ Integration testbed
- ➡ Pilot testbed
- ➡ Maintenance of Tools and Repositories

✗ Pilot Testbed is for internal use but can be opened to external communities for specific activities

## Integration and Development Testbeds

Creato da Doina Cristina Duma

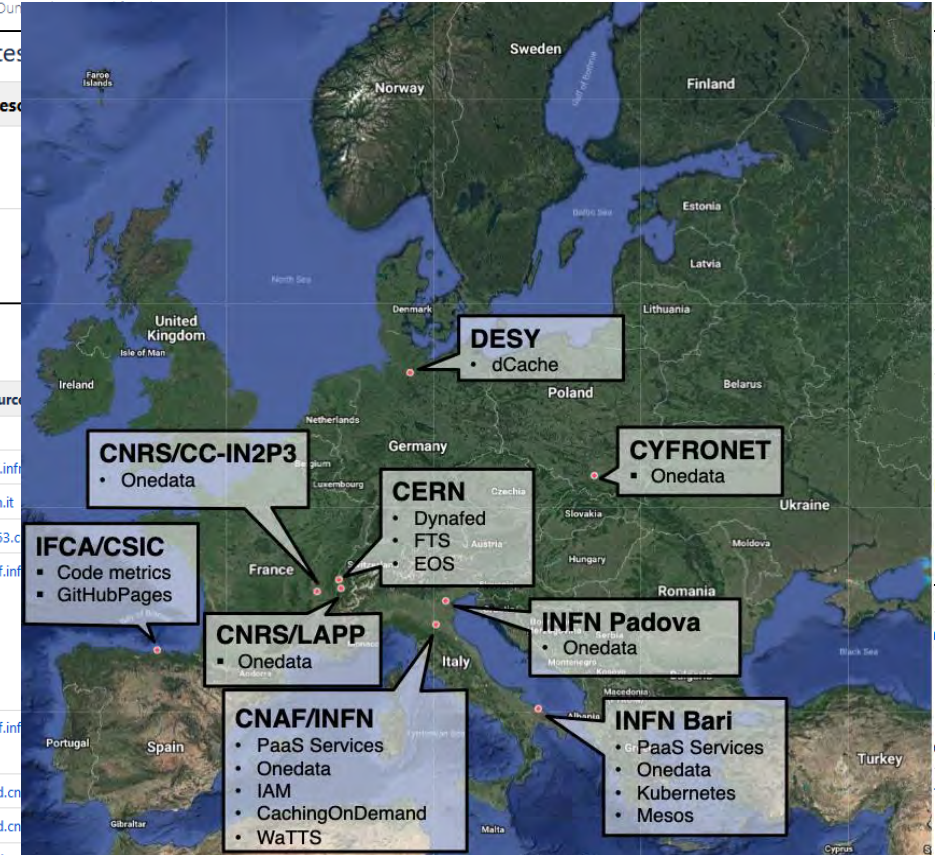
### Development testbed

| Type of resource  |
|-------------------|
| PaaS/Orchestrator |
| PaaS/CMDB         |

### Pilot Preview testbed

Creato da Doina Cristina Duma, ultima modifica il feb 19, 2019

| Type of services      | Partner     | Endpoint of service/resource  |
|-----------------------|-------------|---|
| Onedata               | CYFORNET    | <a href="https://onedata.org">onedata.org</a>   |
| Onedata (Onezone)     | INFN - CNAF | <a href="https://onezone.cloud.cnaf.infn.it">https://onezone.cloud.cnaf.infn.it</a>   |
| Onedata (OneProvider) | INFN Padova | <a href="https://one-data-01.pd.infn.it">https://one-data-01.pd.infn.it</a>   |
|                       |             | <a href="https://cloud-90-147-75-163.c">https://cloud-90-147-75-163.c</a>   |
| PaaS Orchestrator     | INFN        | <a href="https://paas-xdc.cloud.cnaf.infn.it">https://paas-xdc.cloud.cnaf.infn.it</a>   |
| PaaS SLAM             | INFN        | <a href="https://paas-xdc.cloud.cnaf.infn.it">https://paas-xdc.cloud.cnaf.infn.it</a>   |
| PaaS CMDB             | INFN        | <a href="http://paas-xdc-tools.cloud.cnaf.infn.it:8082">http://paas-xdc-tools.cloud.cnaf.infn.it:8082</a>                                   |
| PaaS CPR              | INFN        | <a href="http://paas-xdc-tools.cloud.cnaf.infn.it">http://paas-xdc-tools.cloud.cnaf.infn.it</a>   |
| PaaS Zabbix           | INFN        | <a href="http://paas-xdc-tools.cloud.cnaf.infn.it">http://paas-xdc-tools.cloud.cnaf.infn.it</a>   |
| PaaS Zabbix Wrapper   | INFN        | <a href="http://paas-xdc-tools.cloud.cnaf.infn.it:8082">http://paas-xdc-tools.cloud.cnaf.infn.it:8082</a>                                   |
| IAM                   | INFN        | <a href="https://iam.extreme-datacloud.eu/">https://iam.extreme-datacloud.eu/</a>   |
| External services     |             |   |
| VOMS                  | LIP         | <a href="https://voms01.ncg.ingrid.pt:8443/voms/vo.indigo-datacloud.eu/">https://voms01.ncg.ingrid.pt:8443/voms/vo.indigo-datacloud.eu/</a> |
| WaTTS                 | INFN        | <a href="https://tts-01.cloud.cnaf.infn.it:25554/">https://tts-01.cloud.cnaf.infn.it:25554/</a>   |



### Integration testbed

Creating platform-driven on EOSC - Athens

| Type of resource | Partner  | Endpoint of service/resource                  | Version |
|------------------|----------|---|---------|
| Onedata          | CYFORNET | <a href="https://onedata.org">onedata.org</a> | -       |



# Conclusion

- ✕ XDC is adding new functionalities to already existing, production quality, data management services
- ✕ XDC-1/Pulsar was released in January 2019
  - A step towards the complete implementation of the defined architecture
  - Research communities can already start implementing their use cases using Pulsar
- ✕ A second release is foreseen by next October
- ✕ Scalability verification is in progress and is one of the core activities in 2019
- ✕ XDC consortium members are acting as service providers to facilitate the uptake of the XDC services by the EOSC communities
  - We are involving external service providers to increase the uptake of new user communities



# XDC Contacts



Daniele Cesini  
XDC – Project Coordinator  
INFN

[daniele.cesini@extreme-datacloud.eu](mailto:daniele.cesini@extreme-datacloud.eu)

- ✗ Website: [www.extreme-datacloud.eu](http://www.extreme-datacloud.eu)
- ✗ [@XtremeDataCloud](https://twitter.com/XtremeDataCloud) on Twitter
- ✗ Mailing list: [info@extreme-datacloud.eu](mailto:info@extreme-datacloud.eu)