The XDC project

Data Management for extreme scale computing

Daniele Cesini

daniele.cesini<at>extreme-datacloud.eu

eXtreme DataCloud is co-funded by the Horizon2020 Framework Program – Grant Agreement 777367

XDC Objectives

- 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 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

The Einfra-21-2017 Call

- (a) Support to Public Procurement of innovative HPC systems, PPI
- (b) Research and Innovation Actions for e-Infrastructure prototypes
  - 1 - Universal discoverability of data objects and provenance
  - **2 – Computing e-infrastructure with extreme large datasets**
- Service prototypes should follow common interfaces to access and analyse underlying data *collected/stored in different platforms, formats, locations and e-infrastructures […] tested against requirements of* ***very large or highly heterogeneous research data sets***.
- Funds development of **service prototypes** at **TRL6+**
  - **Bring to TRL8 and include in a unified service catalogue in 2018+**
- **Budget per proposal**: **2.5-3M€**

XDC Foundations

- **XDC take the move from**
  - the INDIGO Data management activity
  - the experience of the project partners on data-management
- **Improve already existing, production quality, Federated Data Management services**
  - By adding missing functionalities requested by research communities
  - Must be coherently harmonized in the European e-Infrastructures

INDIGO PaaS Orchestrator

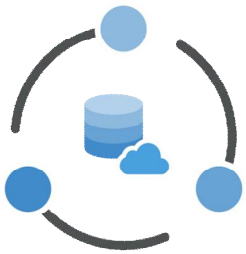

INDIGO CDMI Server

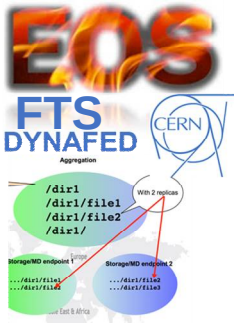

Represented research communities

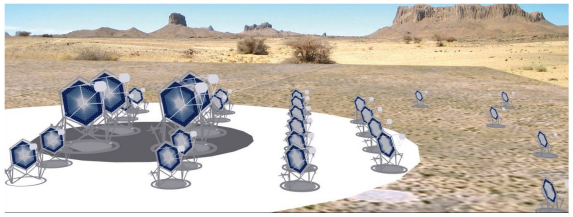

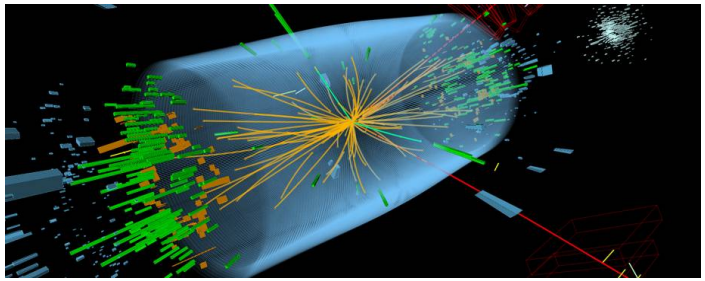

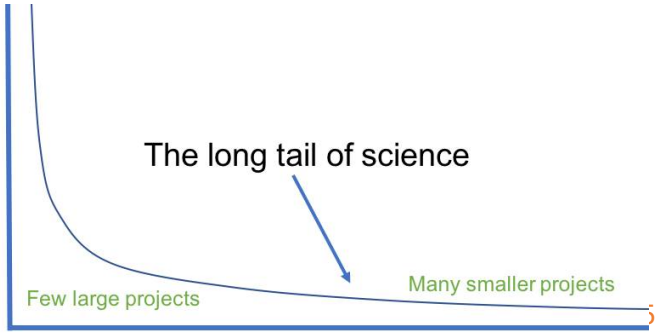

XDC Consortium

ID	Partner	Country	Represented Community	Tools and system
1	INFN (Lead)	IT	HEP/WLCG	INDIGO-Orchestrator, INDIGO-CDMI(*)
2	DESY	DE	Research with Photons (XFEL)	dCache
3	CERN	CH	HEP/WLCG	EOS, DYNAFED, FTS
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
- 7 research communities represented + EGI
- XDC Total Budget: 3.07Meuros
- XDC started on Nov 1st 2017 – will run for 27 months until Jan 31st 2020

XDC Technical 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)
  - Software lifecycle management
- Data pre-processing during ingestion
- Data management based on access patterns
  - Move to ‘glacier-like’ storage unused data, move to fast storage “hot” data
    - at infrastructure level
- Smart caching
  - Transparent access to remote data without the need of a-priori copy
- Metadata management
- Sensitive data handling
  - secure storage and encryption

Policy driven Data Management

- Intelligent & Automated Dataset Distribution
  - A typical workflow
    - Initially the data will be stored on low latency devices for fast access
    - To ensure data safety, the data will be replicated to a second storage device and will be migrated to custodial systems, which might be tape or S3 appliances
    - Eligible users will get permission to restore archived data if necessary
    - After a grace period, Access Control will be changed from “private” to “open access”
  - Data management based on access pattern

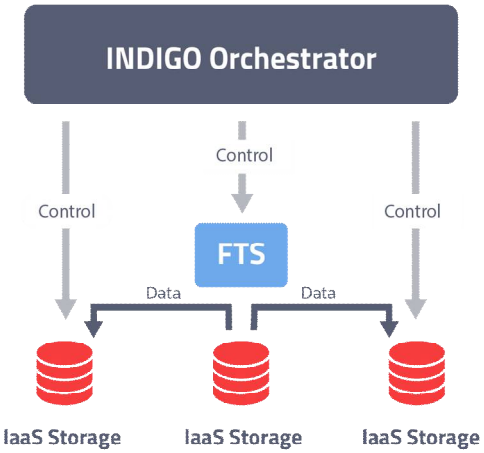

Data pre-processing

Data pre-processing during ingestion

- Automatically run user defined applications and workflows when data are uploaded
  - i.e. for Skimming, indexing, metadata extraction, consistency checks
- Implement a solution to discover new data at specific locations
- Create the functions to request the INDIGO PaaS Orchestrator to execute specific applications on the computing resources on the Infrastructure
- Implement a high-level workflow engine, that will execute applications defined by the users
- Implement the data mover to store the elaborated data in the final destination

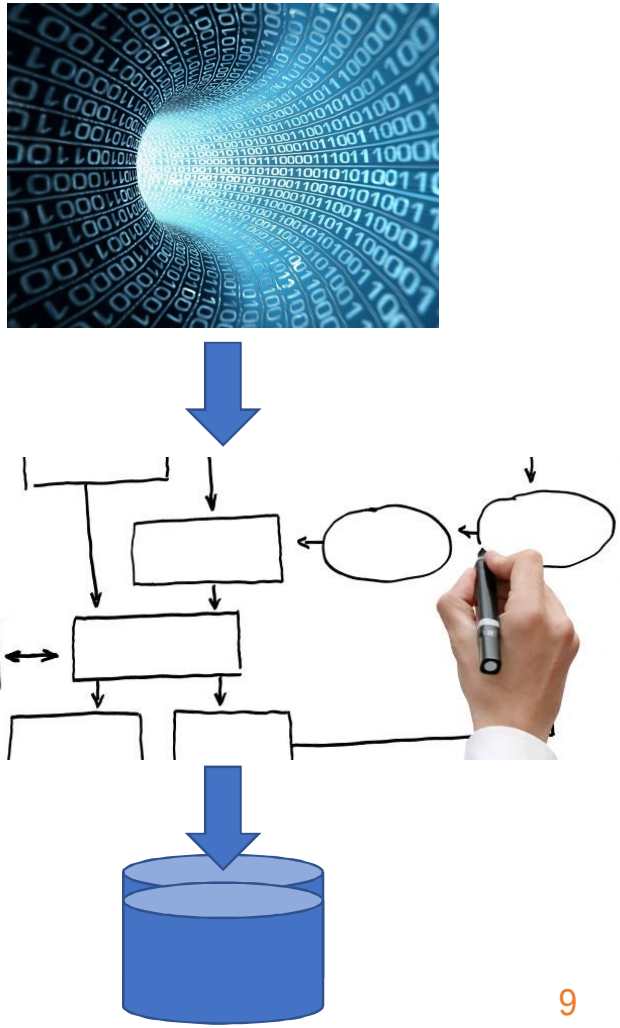

Smart caching

Smart caching

- Develop a global caching infrastructure supporting the following building blocks:
  - dynamic integration of satellite sites by existing data centres
  - creation of standalone caches modelled on existing web solutions
  - federation of the above to create a large scale caching infrastructure

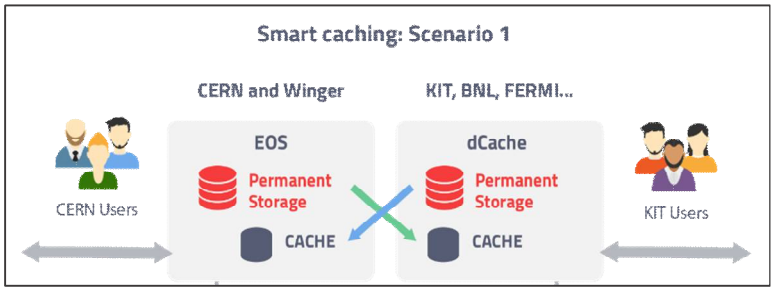

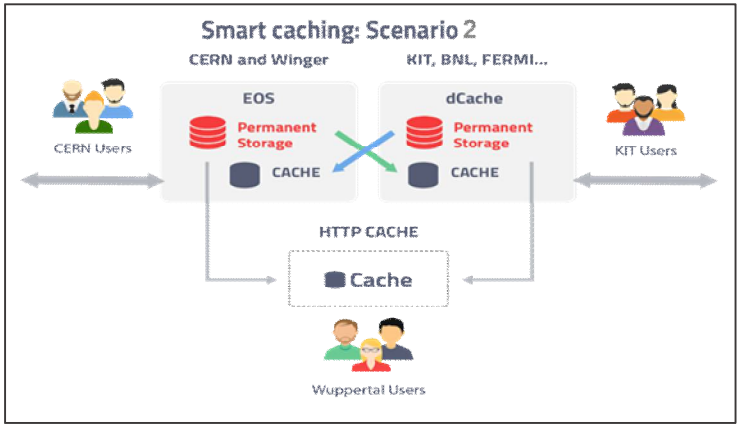

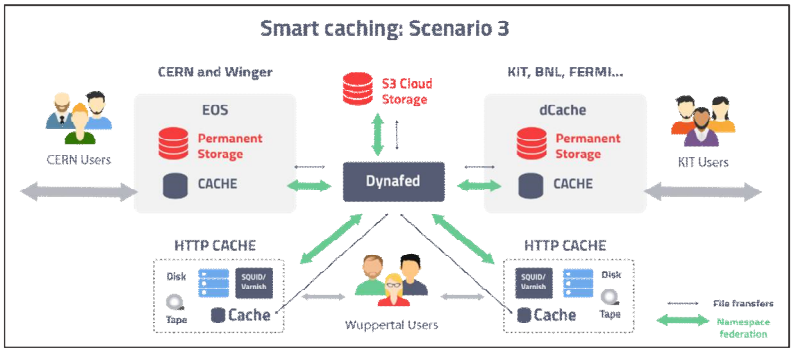

Onedata

- Is a storage federator that allows users to store, process and publish data using global data storage backed by resource providers worldwide
  - Providers deploy *Oneprovider* services near physical storage resources
  - Users use *Onezone* web interfaces
    - APIs available
    - Local mounting on users machines available
  - Storage is organized into **Zones**
    - federations of providers
    - enable the creation of closed or interconnected communities

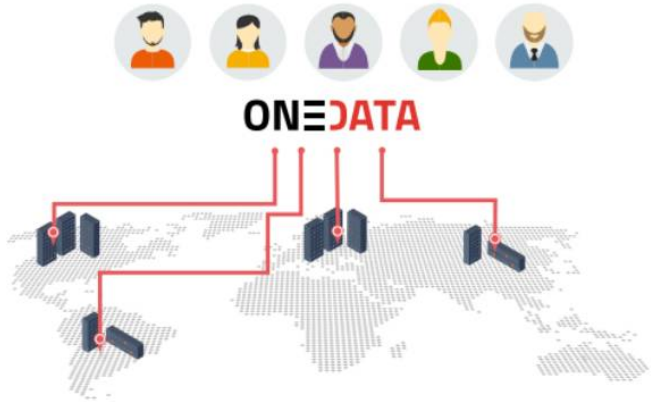

https://onedata.org

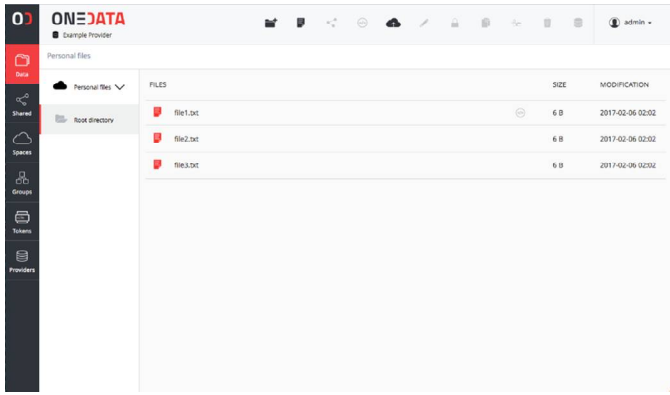

Onedata developments

- Unified data access platform at a PaaS level at the Exascale
- Advanced metadata management with no pre-defined schema
- Encryption Services and Secure Storage
- Sensitive data management and key storage within Onedata

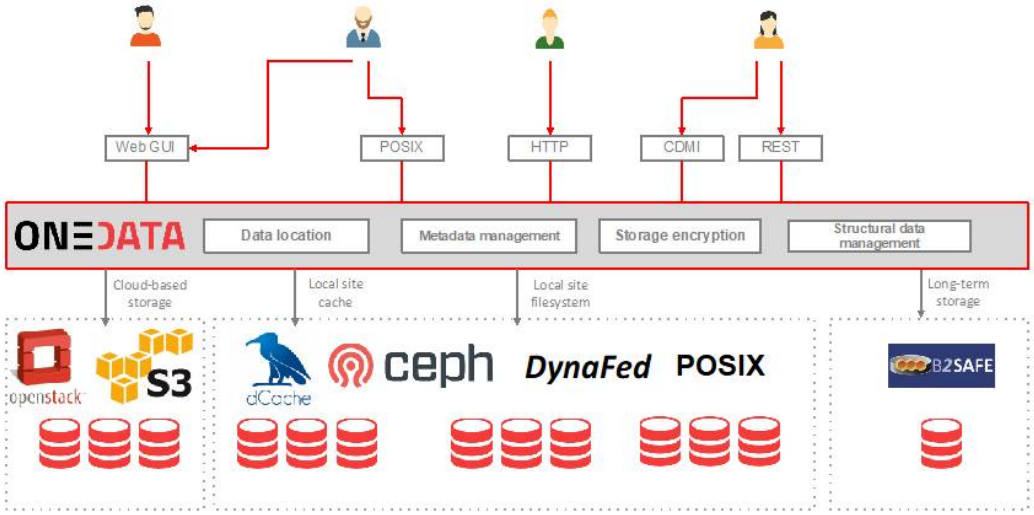

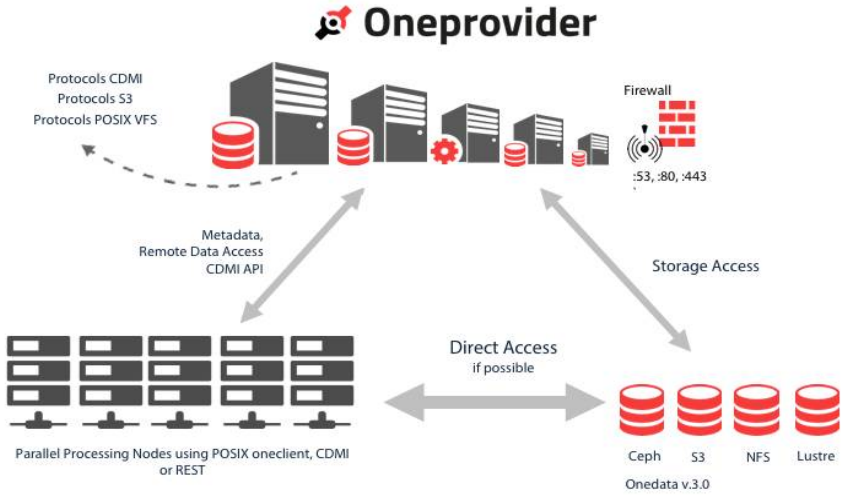

LifeWatch Use Case

- **Problem**: Life Cycle Management of data related to **Water Quality** involving heterogeneous data sources
  - Satellite, Real-time monitoring, meteorological stations.
- **Goal**: Integrate data sources and different types of modelling tools to simulate freshwater masses in a FAIR data environment
  - Use of standards like EML (Ecological Metadata Language)
- **XDC Solution**:
  - Onedata
    - Metadata management and discovery, Digital Identifier minting, storage
  - PaaS Orchestrator
    - automatic preprocessing for data harmonization and model deployment

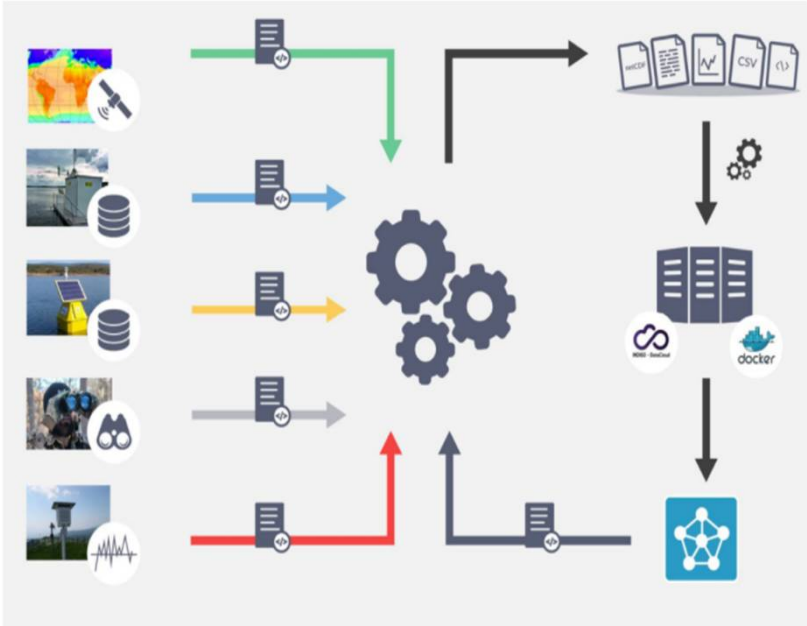

CTA Use Case

- **Problem**: Complex and Big Data management in a distributed environment. Data quality Assurance
  - The CTA distributed archive lies on the « Open Archival Information System » (OAIS) ISO standard.
  - Event data are in files (FITS format) containing all metadata.
- **Goal**: Metadata are extracted from the ingested files, with an automatic filling of the metadata database.
  - Metadata will be used for querying of archive.
  - The system should be able to **manage replicas**, tapes, disks, etc, with data from low-level to high-level
- **XDC Solutions**
  - Onedata
    - Metadata management and discoverability
  - PaaS Orchestrator + QoS

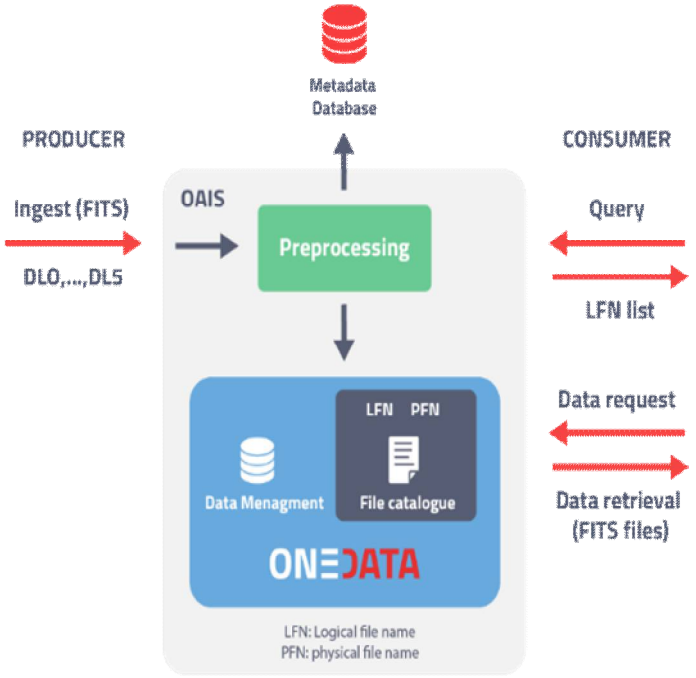

ECRIN Use Case

- **Problem**: Distributed files and data objects across different repositories. Metadata heterogeneity. Sensitive Data
- **Goal**: Single environment to make clinical trial data objects available for sharing with others. Sources are spread over
  - a variety of access mechanisms
  - several different locations
    - growing number of general and specialised data repositories
    - trial registries
    - Publications
    - the original researchers' institutions
- **XDC Solution**: Onedata
  - Metadata management and discovery
  - Secure Storage

WLCG Use Case

- **Problem**: Growing needs on storage space
  - up to 900 PB in 2027
  - Data ready to be used/exploited in a very distributed environment
- **Goal**: Reduce costs, reduce disk needs, smart data allocation.
- **XDC Solution**: Smart Caching systems
  - Data Federation
  - Multi-site storage - the “DataLake”.
  - Dynamic extension of sites to remote locations

XFEL Use Case

- **Problem**: Data Lifecycle Management for the XFEL facilities, ACL and Embargo period control
  - Online:
    - Exclusive access to cluster during beam time, only from experiment rooms. Perform preliminary, On-the-fly analysis
  - Offline:
    - Calibrated data provision for analysis on the HPC Cluster and in the OpenStack cloud environment.
- **Solution**: dChace + Orchestrator

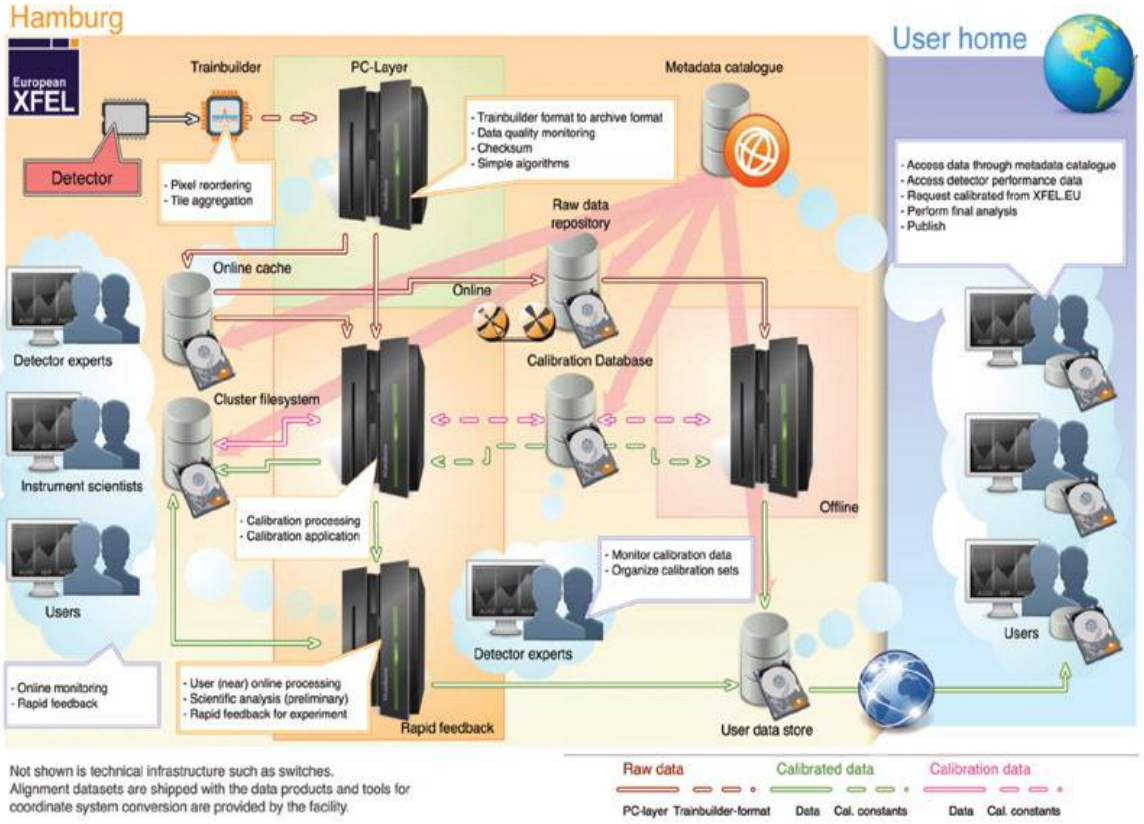

XDC high level architecture

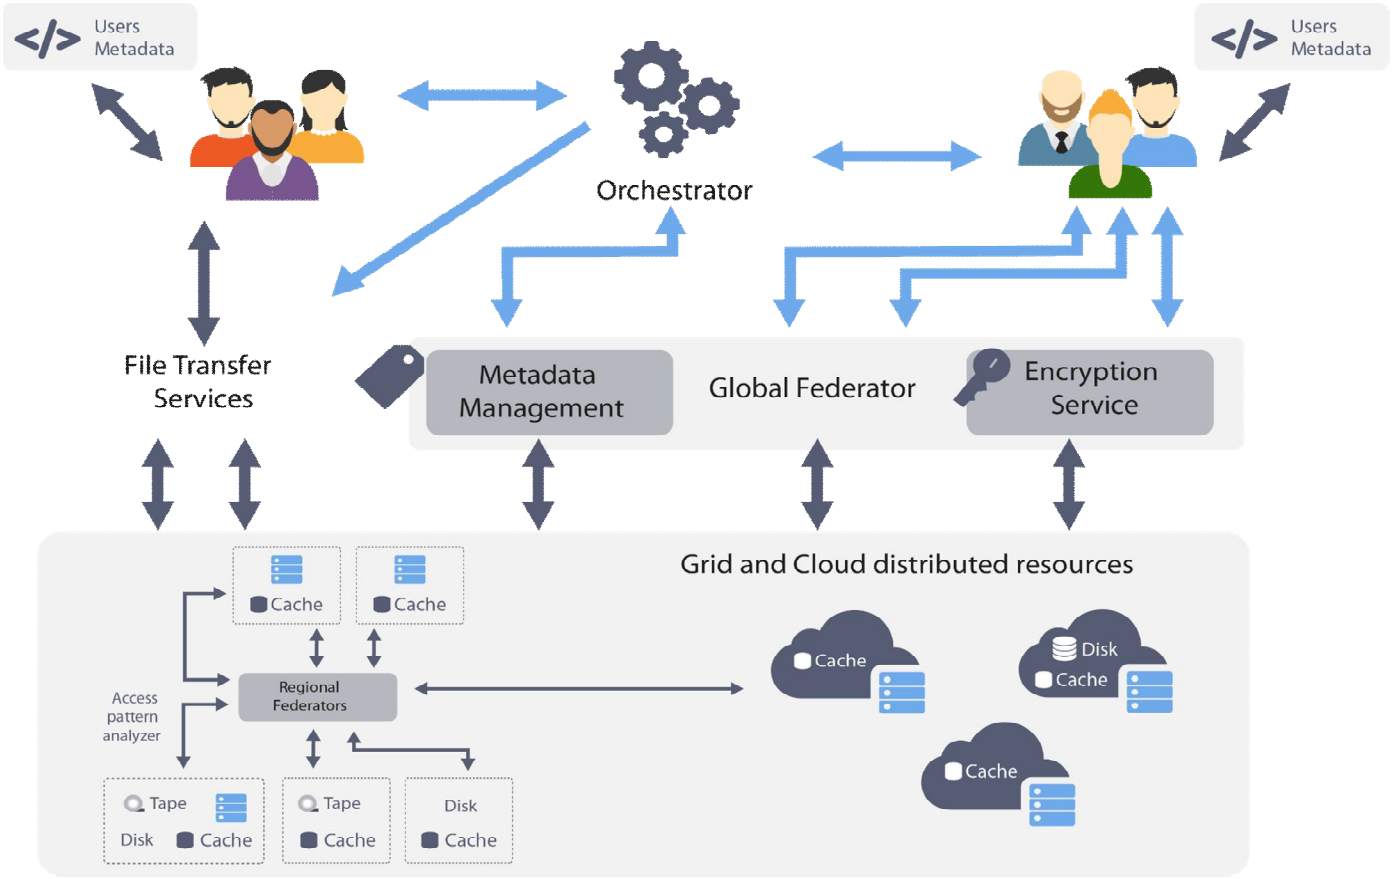

Project Status

- Started on Nov 1st 2017
- Detailed requirements collection from user communities completed
- Definition of the detailed architecture ready in May
- Creation of the Pilot Testbed started
  - Currently reserved for internal communities
  - Under discussion the possibility to open to external users
- Liaisons initiated with other DM development projects and EOSC-related initiatives
  - EOSC-Hub
  - EUDAT
  - DEEP-HybridDataCloud
  - RUCIO development team
  - All EINFRA-21 projects

The Release Plan

- Event with User Communities – Jun 18-22 2018, Santander – joint with DEEP
- XDC reference releases – 1 - Oct-Nov 2018
- XDC reference releases – 2 - Oct-Nov 2019
- Functionalities and scalability demonstrated - Jan 2020

Conclusion

- XDC has an ambitious development plan for data management services
  - We want to support very diverse use cases and requirements
- We will support as much as possible standard protocols to make the solutions as general as possible
- First release is foreseen on Oct-Nov 2018
- Sustainability of the products
  - Provide upstream to the original project all the changes developed by XDC
  - Involving the user communities in exploiting the XDC outputs in their production environments
  - Pushing XDC developments in the EOSC Service Catalogue
  - Engaging liaisons with e-infra providers and other DM development projects

XDC Contacts

- Website: www.extreme-datacloud.eu
- **@XtremeDataCloud on Twitter**
- **Mailing list: info<at>extreme-datacloud.eu**