


The eXtreme-DataCloud project



Data Management for extreme scale computing



Daniele Cesini

info@extreme-datacloud.eu

www.extreme-datacloud.eu



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

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

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
- ☛→ **TRL 6+ → TRL8** (as requested by the H2020 call)

XDC Consortium



ID	Partner	Country	Represented Community	Tools and system
1	INFN (Lead)	IT	HEP/WLCG	INDIGO-Orchestrator
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	EGL.eu	NL	EGL communities	

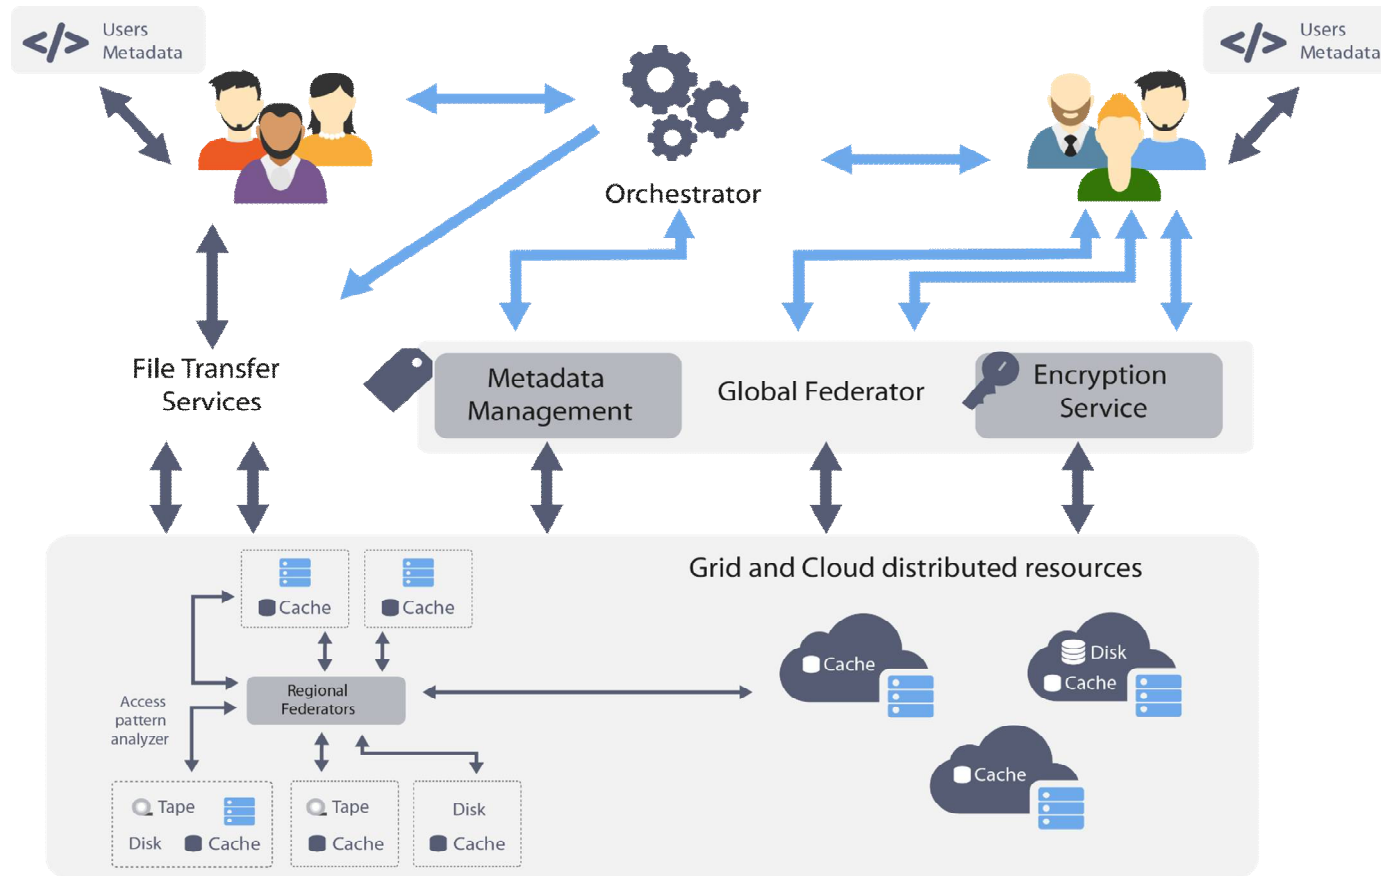


- ✗ 8 partners, 7 countries
- ✗ 7 research communities represented + EGI
- ✗ XDC Total Budget: 3.07Meuros
- ✗ XDC started on Nov 1st 2017 – will run for 27 months until Jan 31st 2020

XDC Technical Topics

- ✘ Intelligent & Automated Dataset Distribution
 - ☛→ Orchestration to realize a policy-driven data management
 - ☛→ Data distribution policies based on Quality of Service **supporting geographical distributed resources** (cross-sites)
 - ☛→ QoS beyond disk vs tape: availability, reliability, durability, latency
 - ☛→ Data lifecycle management
- ✘ Data (pre)processing during ingestion
 - ☛→ Storage Notifications based on events
- ✘ Smart caching
 - ☛→ Transparent access to remote data without the need of a-priori copy
- ✘ Metadata management
- ✘ Data management based on access patterns
 - ☛→ Move to 'glacier-like' storage unused data, move to fast storage "hot" data
 - ☛→ at infrastructure level
- ✘ Sensitive data handling
 - ☛→ secure storage and encryption

XDC high level architecture




XDC Toolbox

Storage




Rucio



Orchestration

INDIGO
Orchestrator



QoS
CDMI

HTTP
Cache

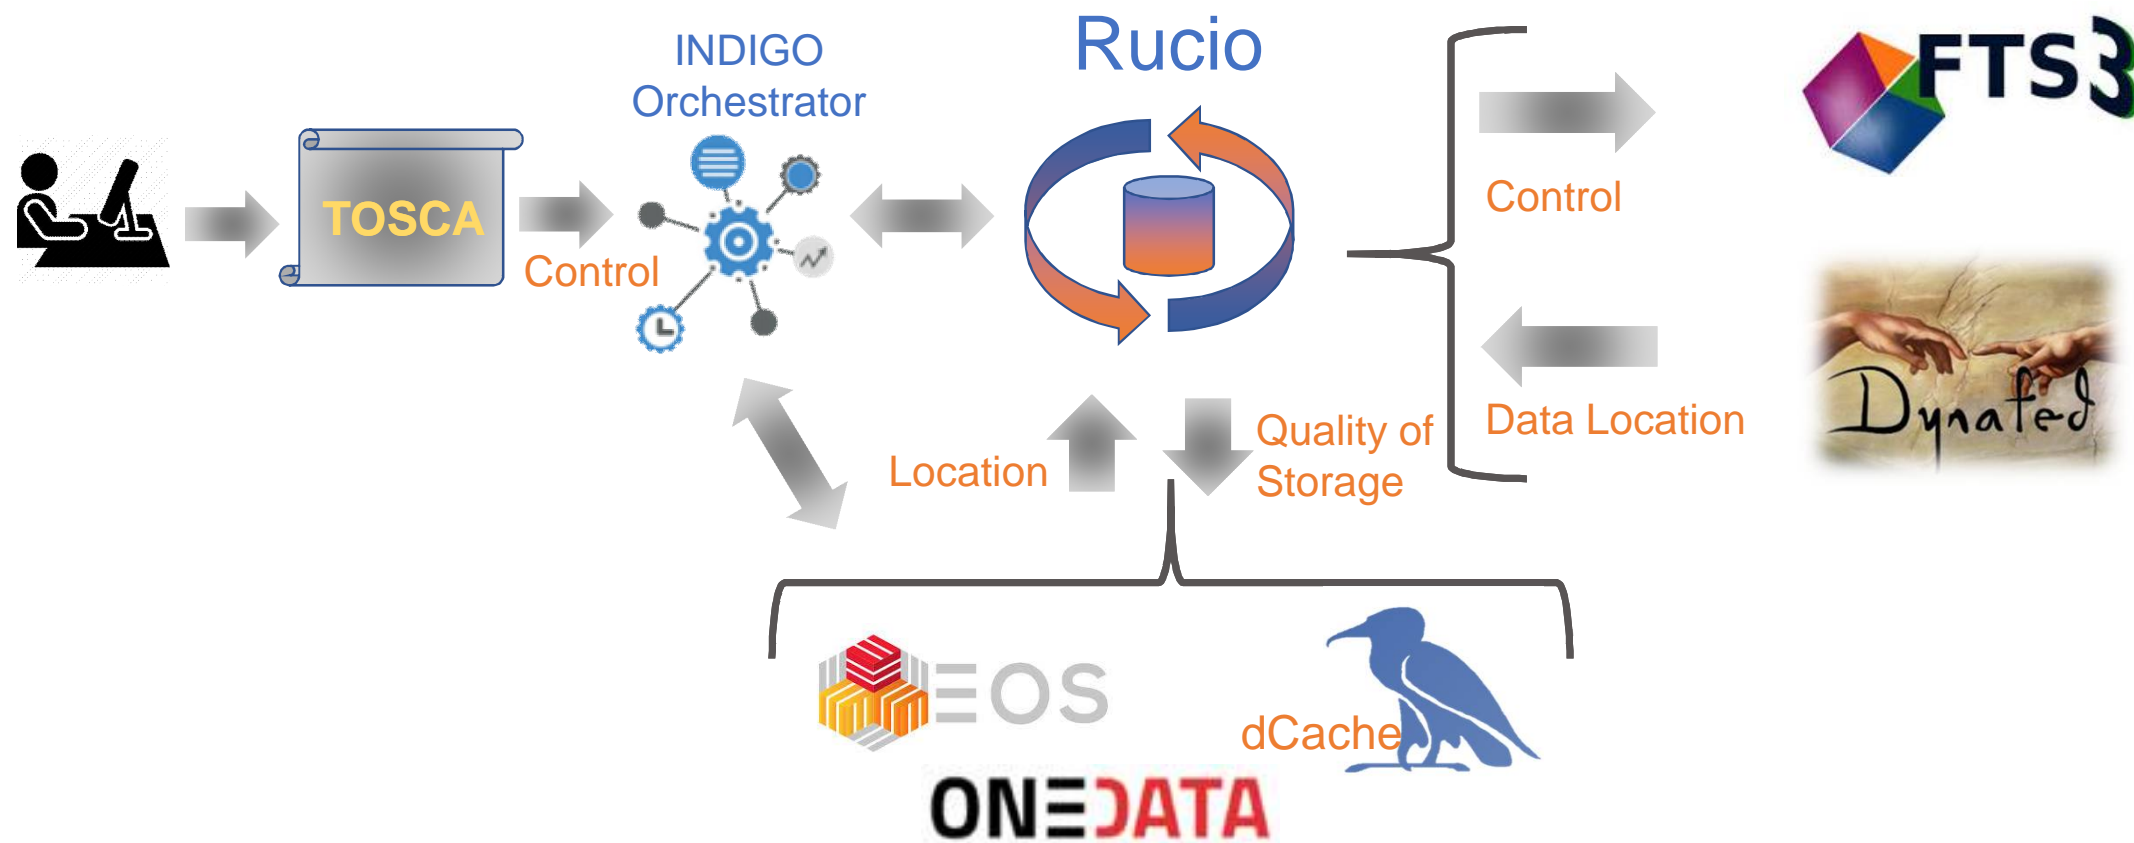
xRootD
Cache

Indigo-IAM

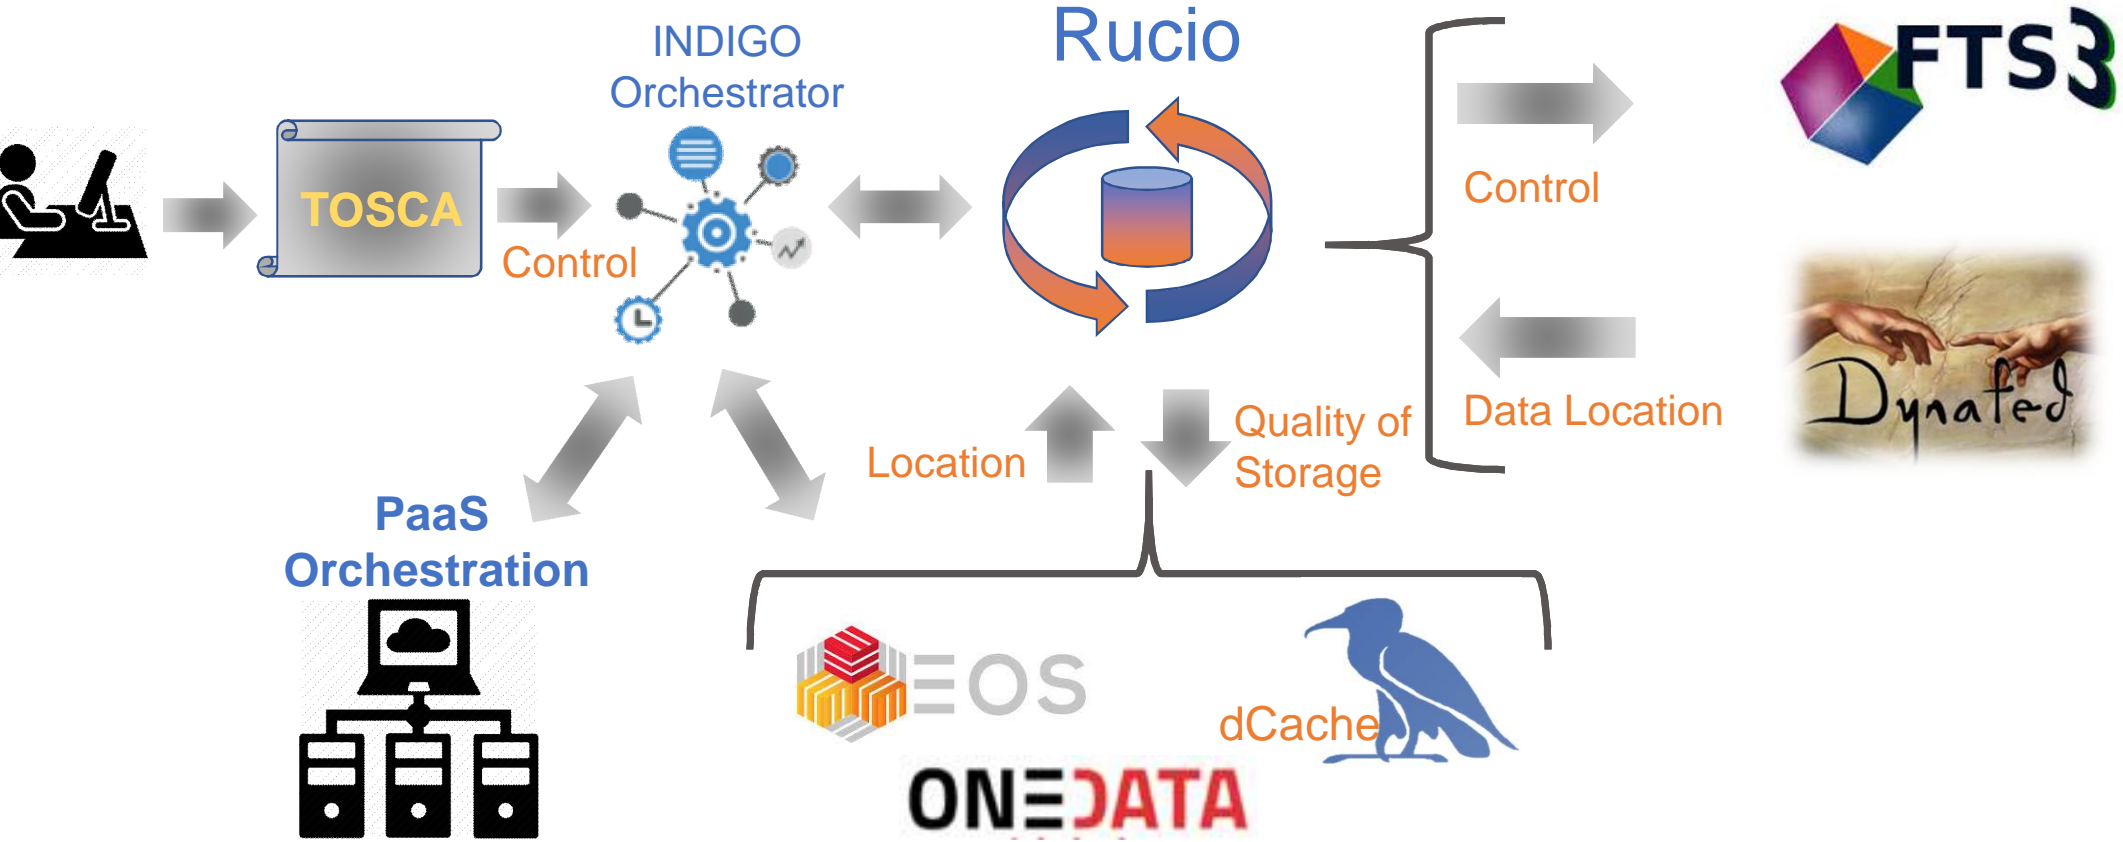


Federation

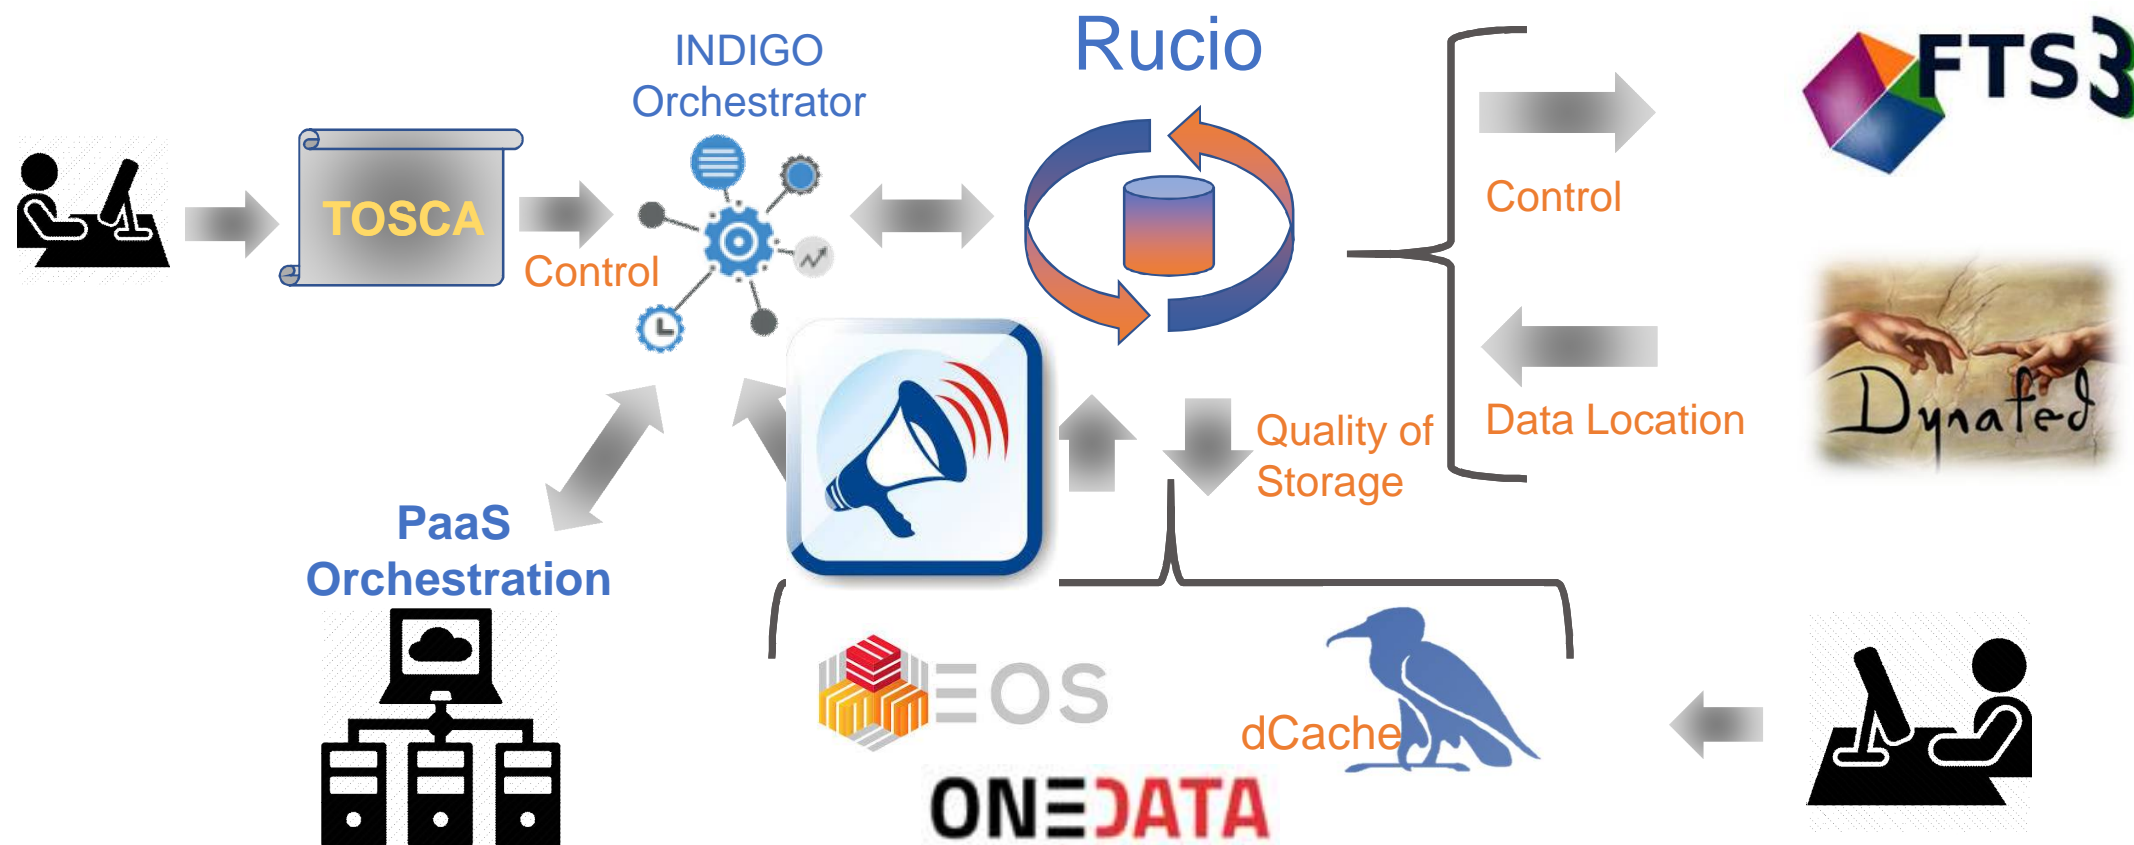
Expected Control Flow for Orchestration and Policy Driven Data Management



Expected Control Flow for Orchestration and Policy Driven Data Management

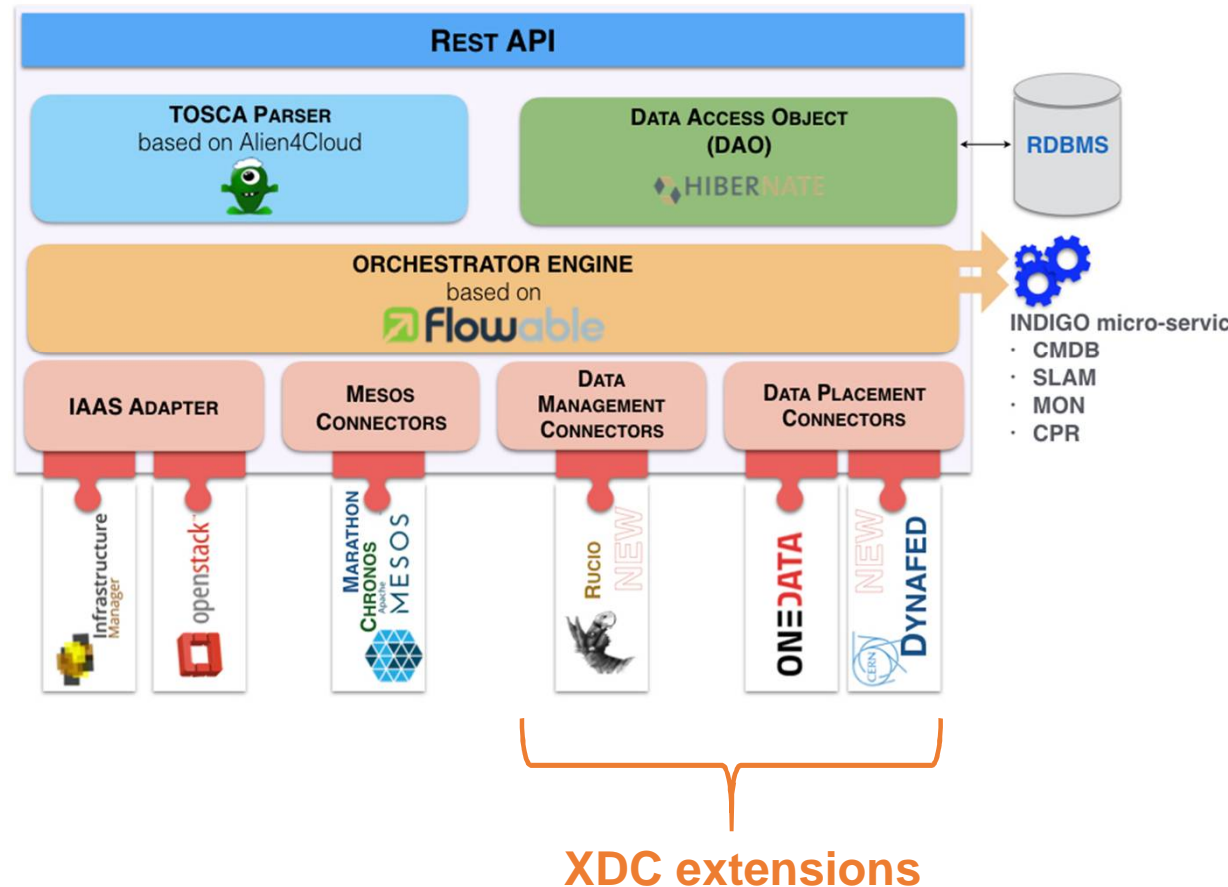


Expected Control Flow for Orchestration and Policy Driven Data Management



INDIGO Orchestrator Overview

- ✗ The Orchestrator accepts TOSCA requests
- ✗ It can run three main workflows in the PaaS:
 - Deploy
 - Undeploy
 - Update
- ✗ The Orchestration layer is being extended in order to address the new requirements:
 - move data between distributed storages
 - specify different QoS for replicas
 - launch and monitor user defined processing jobs at ingestion time



Quality of service in storage















- ✘ Modern Storage offers go beyond the WLCG model of disk and tape
 - ☛→ Often the actual cloud storage technology is not even known to the end user or might change over time
 - ☛→ Mostly only quality attributes are known, like Glacier:
 - ☛→ Durability (Retention Policy): 99.999999999%
 - ☛→ Access Latency : a) 5 minutes, b) 5 hours, c) 12 hours

- ✘ XDC aims to provide
 - ☛→ help to standardize the different attributes
 - ☛→ a prototype of rendering those attributes through a network protocol
 - ☛→ a set of reference implementations to be used by the Orchestration system

- ✘ Work is based on INDIGO-DataCloud prerequisites

Work in progress for QoS

- ✗ Definition currently based on a RDA working group
- ✗ "Cloud Data Management Interface" chosen as control protocol.
 - ➡ Defined by SNIA.
 - ➡ INDIGO acknowledged by SNIA as contributor to the reference implementation.
- ✗ Implementing the defined API into GFAL, dCache, EOS and StoRM

	Name	Access Latency [ms]	Number of Copies	Storage Lifetime	Location	Storage type	Available Transitions
	disk	100	1		DE	 Processing	tape, disk+tape
	disk+tape	100	2		DE	 Processing	tape
	DiskAndTape	50	3	20 years	DE	 Processing	
	DiskAndTape	50	2		IT	 Processing	
	DiskOnly	50	3	20 years	DE	 Processing	
	DiskOnly	50	1		IT	 Processing	
	profile1	10	3	20 years	DE	 Processing	

From INDIGO-DataCloud

- Access Latency [ms]
- Number of Copies
- Storage Lifetime
- Location
- Available Transitions

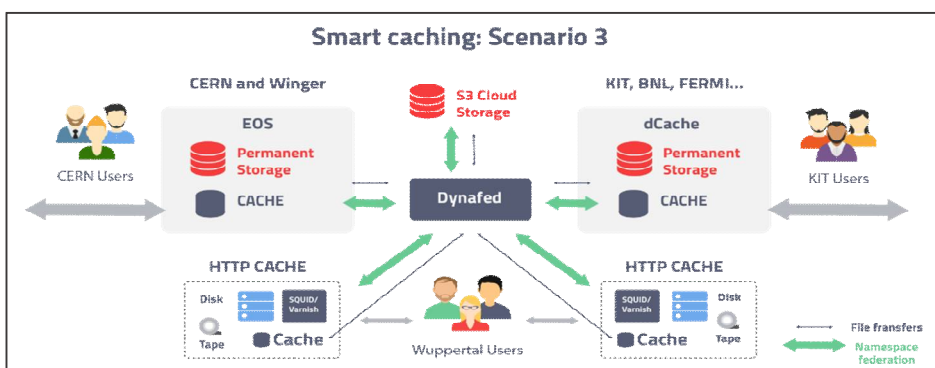
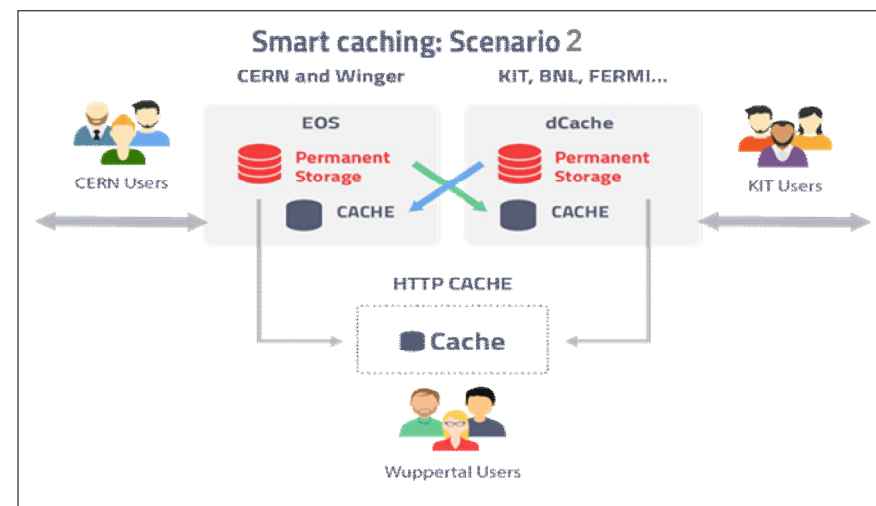
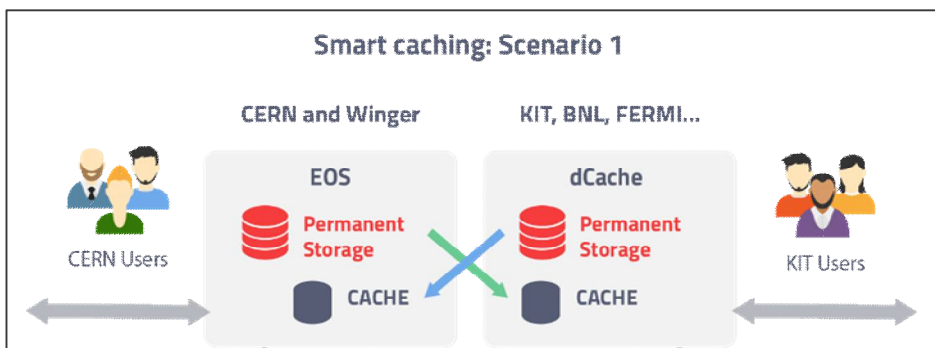
dCache storage events

- ✘ Support has been added for storage events
 - ☛→ Allow non-dCache agent to get notified that something of interest happen inside dCache
- ✘ Two mechanisms to receive events
 - ☛→ Kafka
 - ☛→ exposes dCache internal events (currently) without any security. This is intended for trusted (or otherwise tightly integrated) services.
 - ☛→ Server Sent Events (SSE)
 - ☛→ generic interface that requires authentication. Out-of-the-box, it is available to all dCache users (there's no anonymous event subscription).

Smart caching

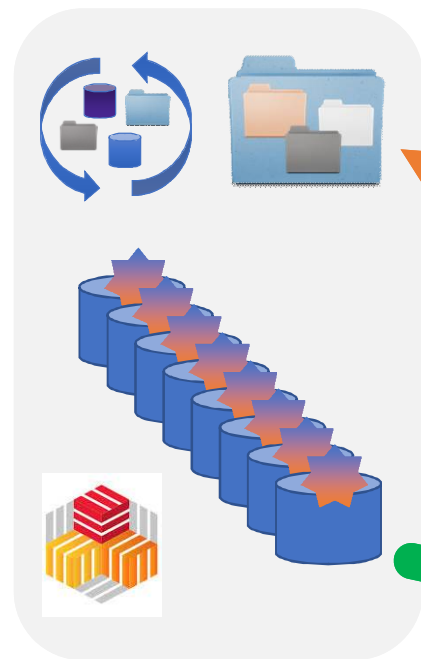
Smart caching

- ➔ 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 **http and xrootd** solutions
 - ➔ federation of the above to create a large scale, **regional caching infrastructure**



EOS-dCache integration with cached access

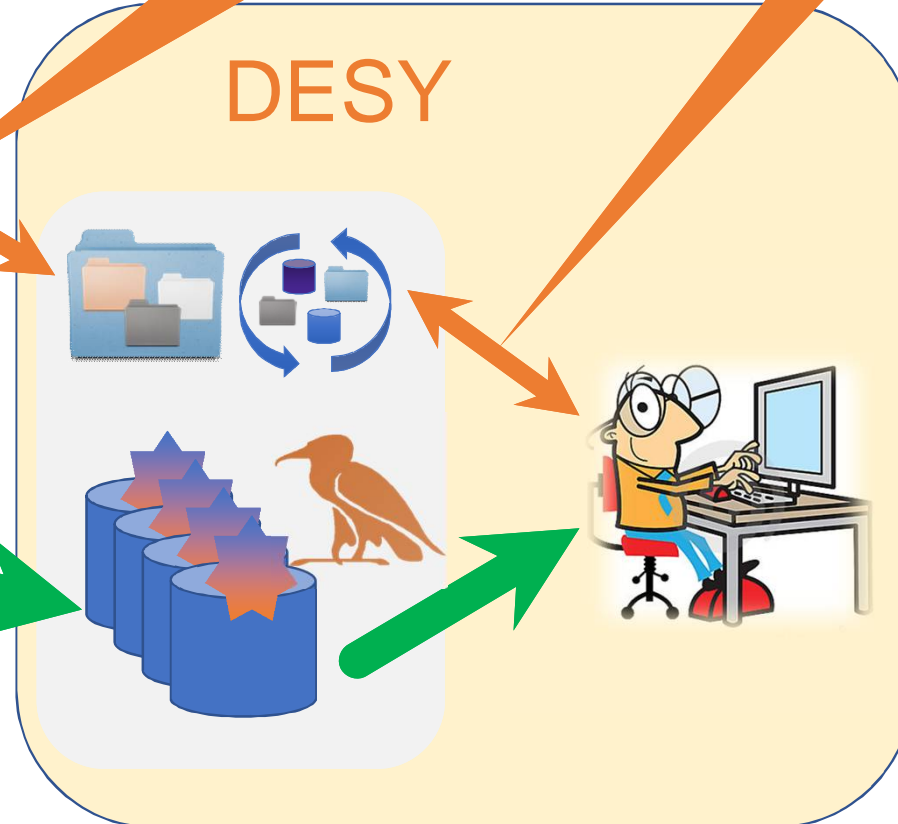
CERN



Sync Namespace

Request

DESY



Advantages

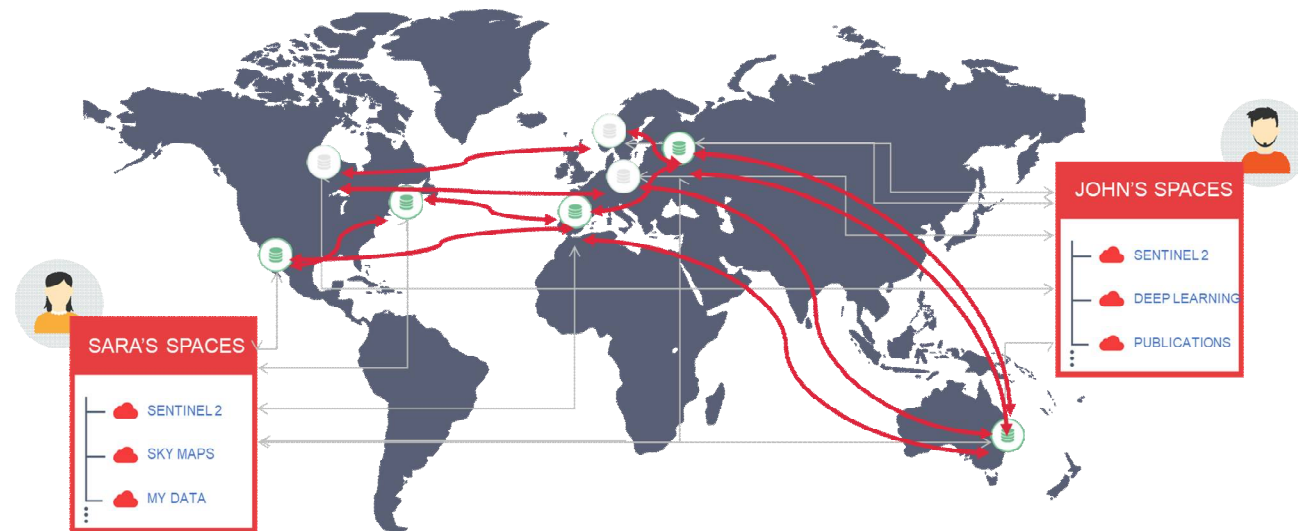
- ✗ Same software stack as we currently have at the sites.
- ✗ After the data has been transferred to the local storage system, a name space entry has been created locally and the data is available at the local site independently of the remote network link and the availability of the central service.

Data in Hybrid Cloud Environments: Onedata

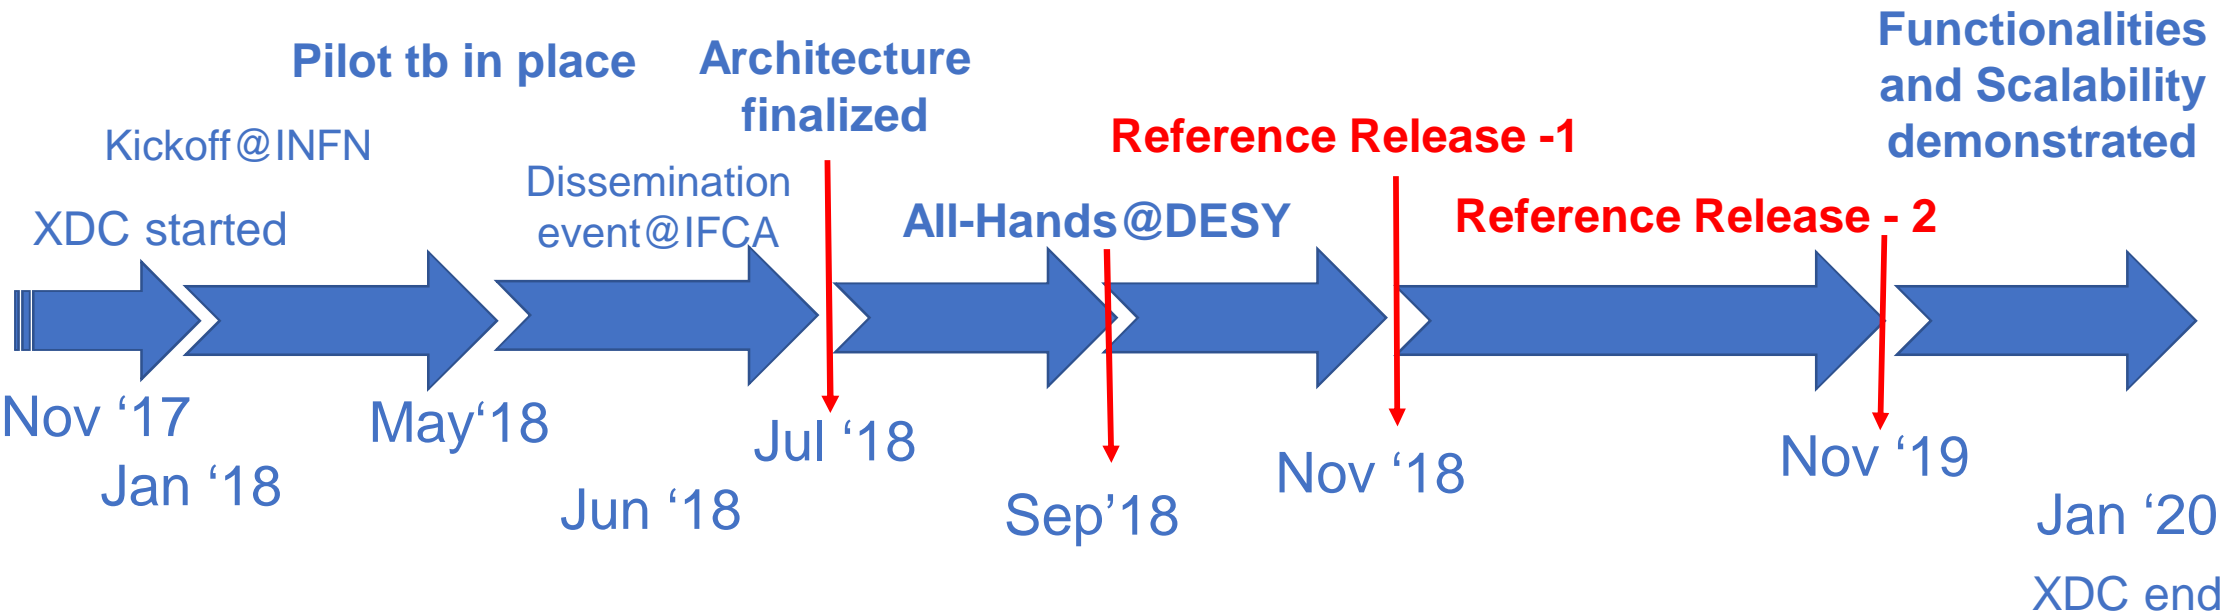
✗ ONEDATA is a storage federator that allows to use resources backed by providers worldwide

- 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>



XDC Schedule



XDC – Release 1

Release date 31/10/2018



Component	XDC functionalities
EOS	<ul style="list-style-type: none">- Caching with xcache for geographic deployment: Xcache deployed at a remote centre to accelerate its local CPU- External storage adoption (Through an S3 or a WebDAV interface)- External data adoption (Data already present on a system described above can be incorporated into EOS)
dCache	<ul style="list-style-type: none">- New QoS types integration, aggregated QoS for storage federations- OpenIDConnect support in dcache_view- dcache storage events (SSE notifications): Allow non-dCache agent to get notified that something of interest happen inside dCache
FTS+gfal	<ul style="list-style-type: none">- QoS support: can now accept a QoS job- OpenIDConnect support- QoS in gfal (gfal with basic cdmi client) – python bindings available
Dynafed	<ul style="list-style-type: none">- Integration of OIDC authentication
Caching infrastructure based on xCache or http	<ul style="list-style-type: none">- deployment receipts for geographically distributed caches (via xcache)- deployment receipts for scalable local caches (via xcache and http)
PaaS Orchestrator	<ul style="list-style-type: none">- Implementation of Dynafed plugin – interaction via INDIGO IAM oauth2 token- Enhancement of ONEDATA plugin
ONEDATA	<ul style="list-style-type: none">- Performance and stability improvements- new RADOS driver

XDC Contacts



✘ Website: www.extreme-datacloud.eu


✘ [@XtremeDataCloud](https://twitter.com/XtremeDataCloud) on Twitter

✘ Mailing list: [info<at>extreme-datacloud.eu](mailto:info@extreme-datacloud.eu)

Backup slides

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

