Status of the technical activities, main issues and all hands objectives

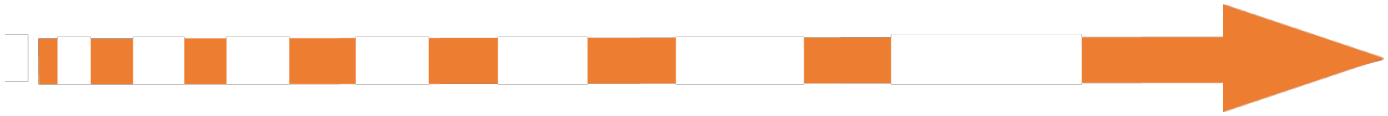

Giacinto DONVITO– INFN-Bari

donvito@infn.it

eXtreme DataCloud

Data Management for extreme scale computing

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

Outlook

- Status of the technical activities:
  - With respect to what described in the DoA
  - With respect to the schedule
- Overall view of the status of understanding of the technical architecture
- Main open issues
- All hands objective from the Joint technical activities.

XDC high level architecture

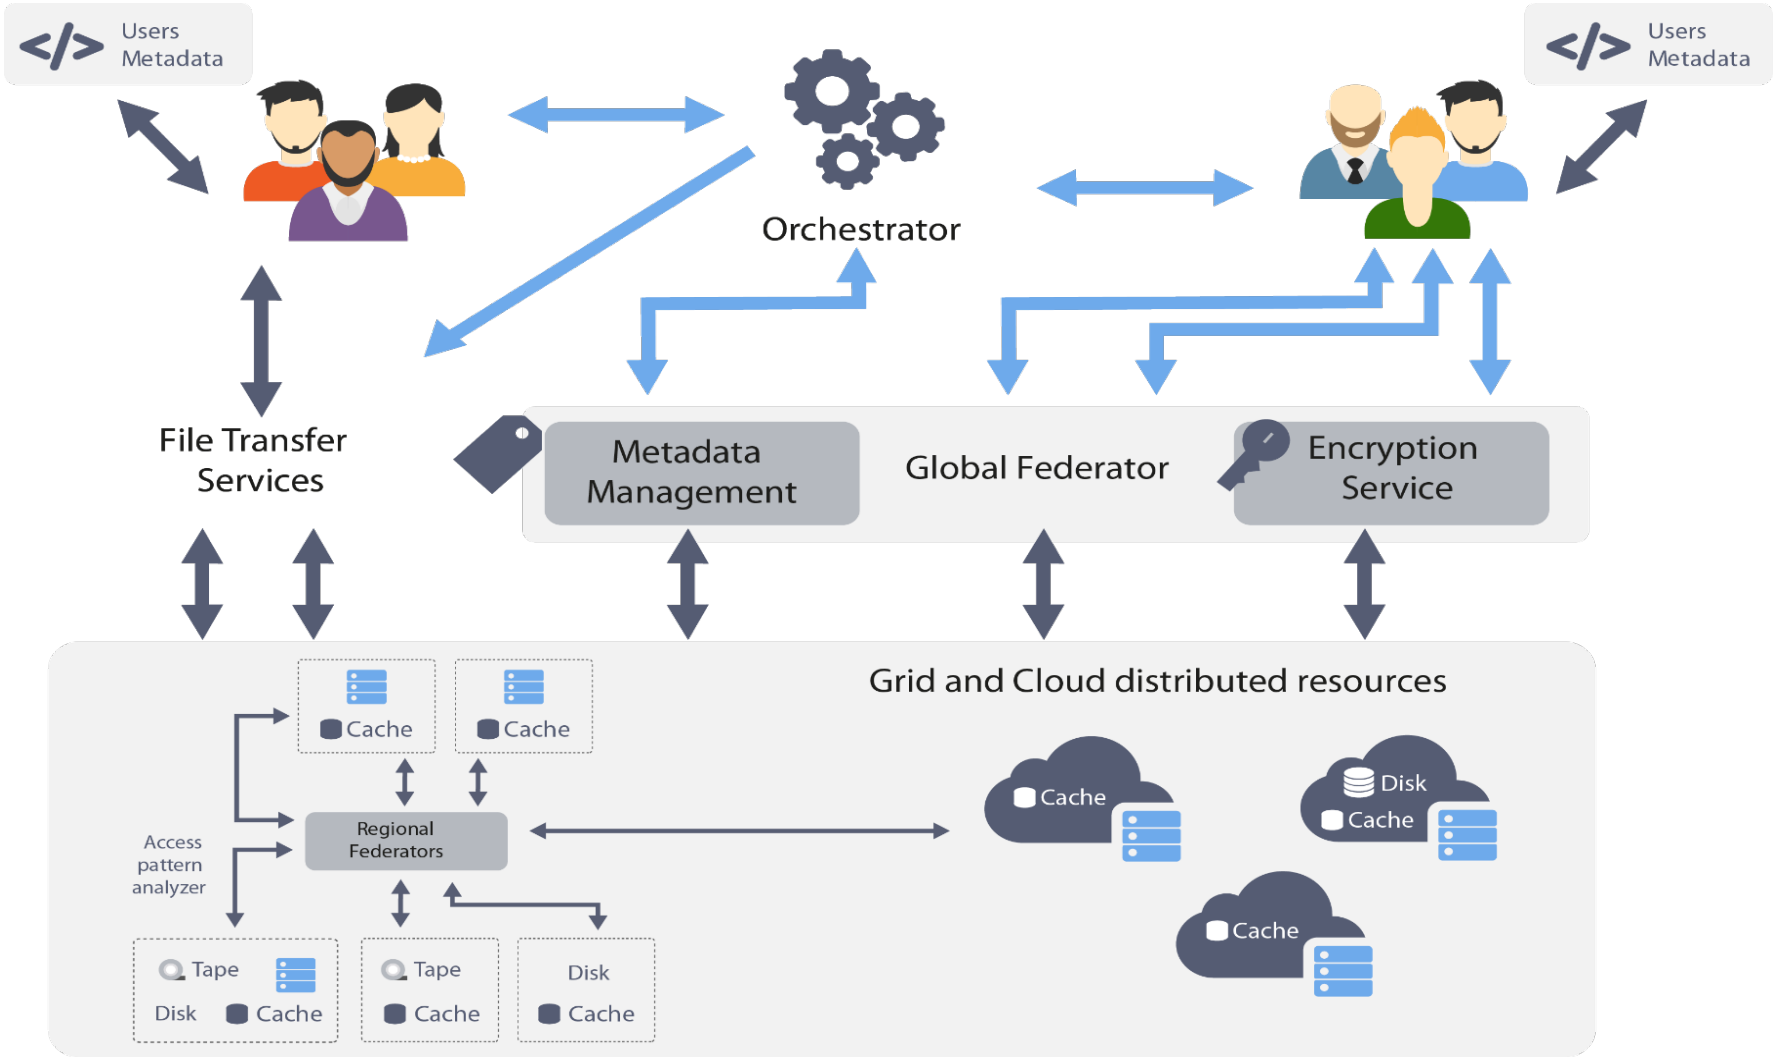

Policy driven Data Management

- 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

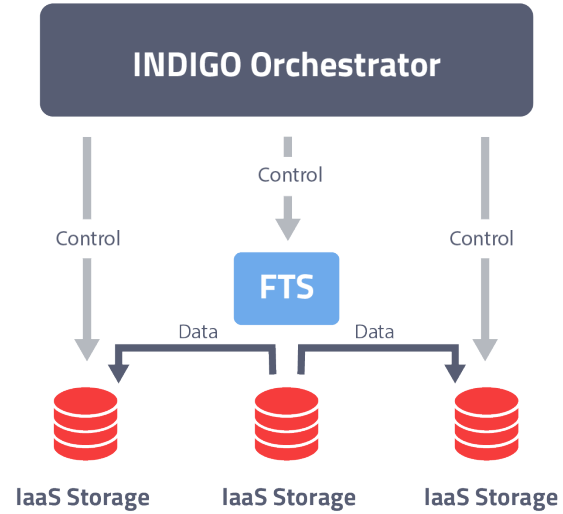

Data pre-processing

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

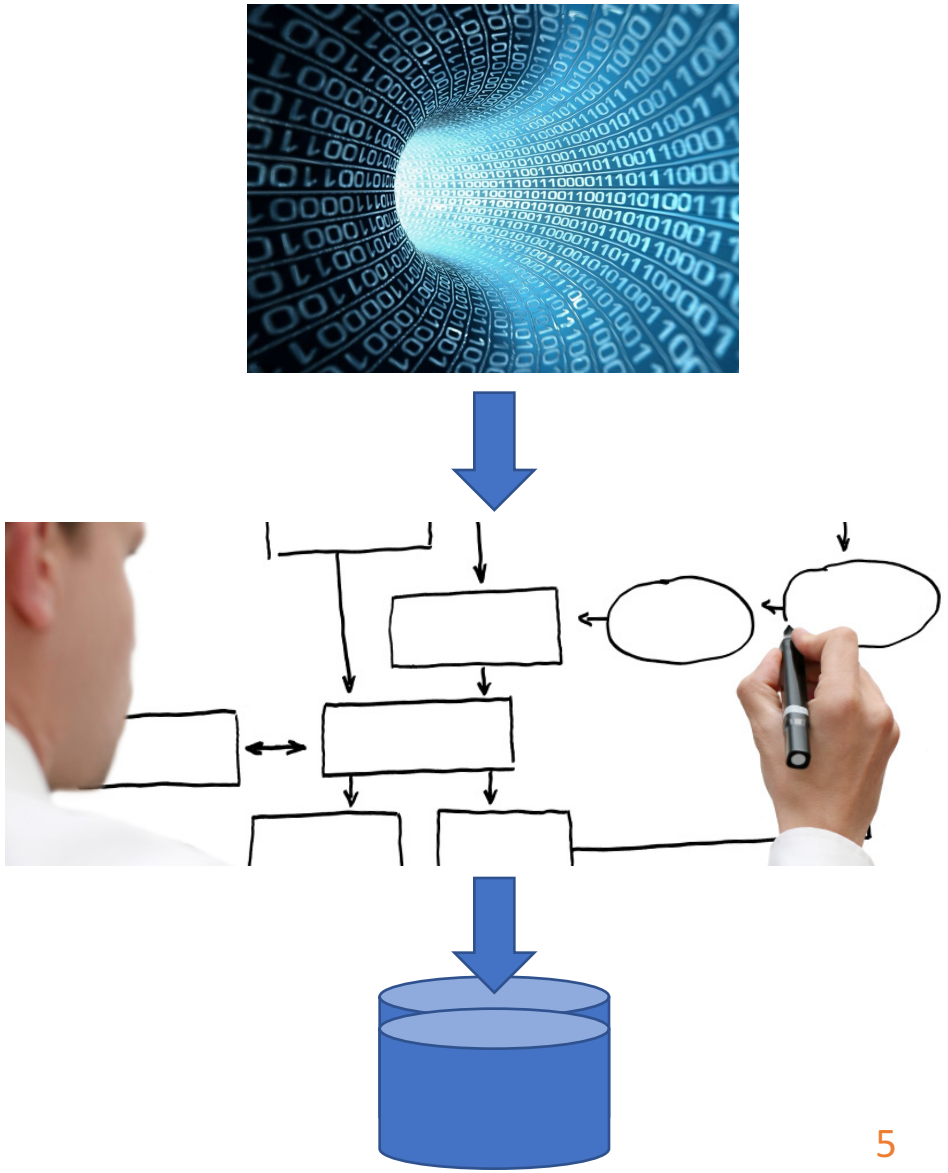

Overall architecture

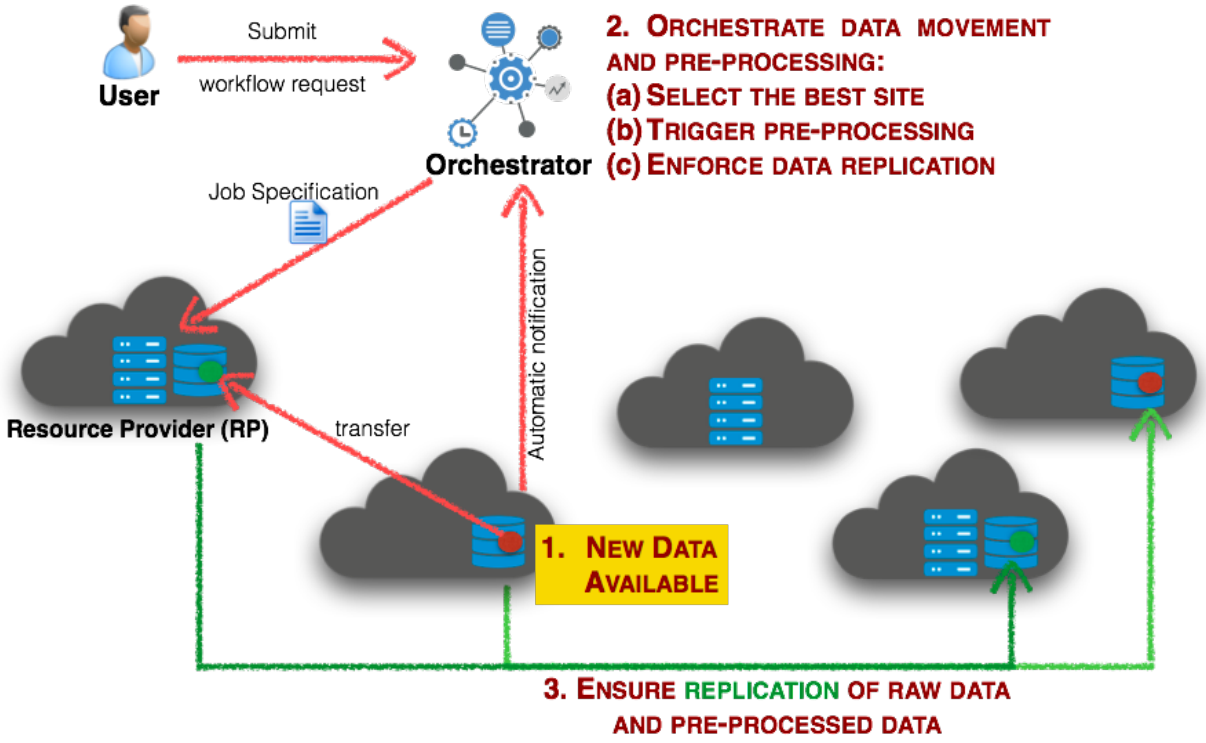

Orchestration…

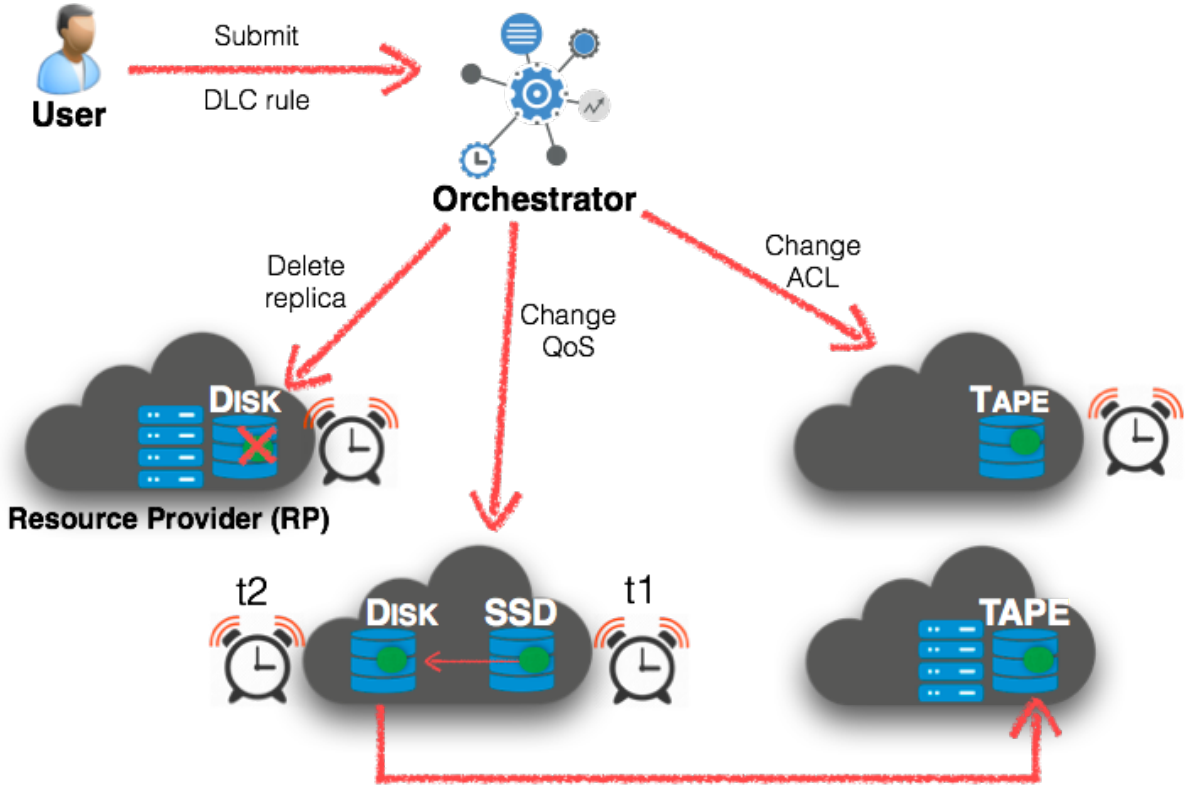

Overall architecture

QoS…

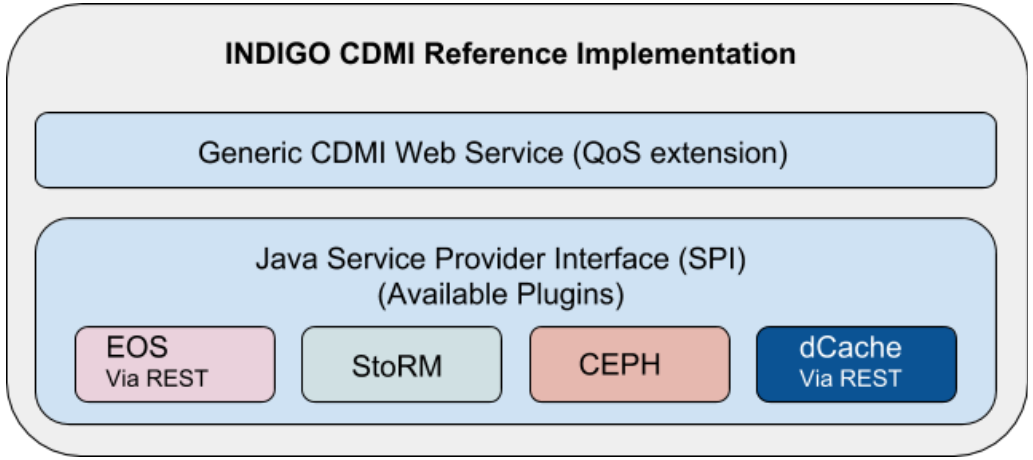

Orchestration…

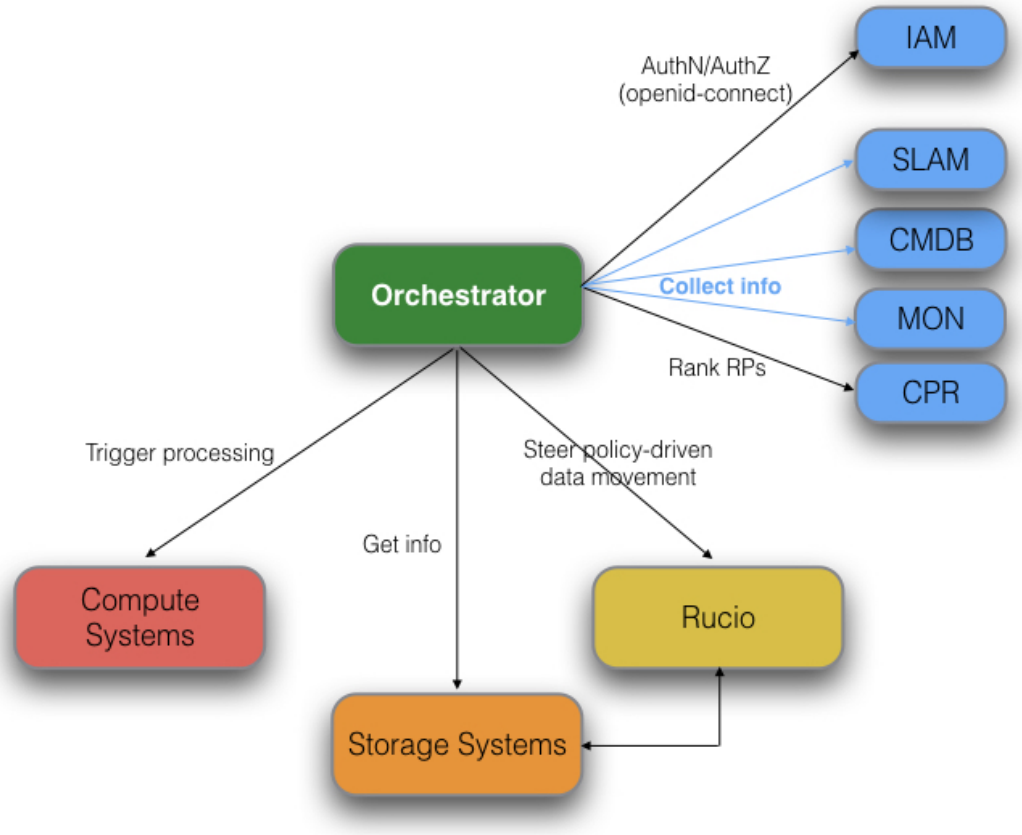

Smart caching

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

Smart caching scenarios - 1

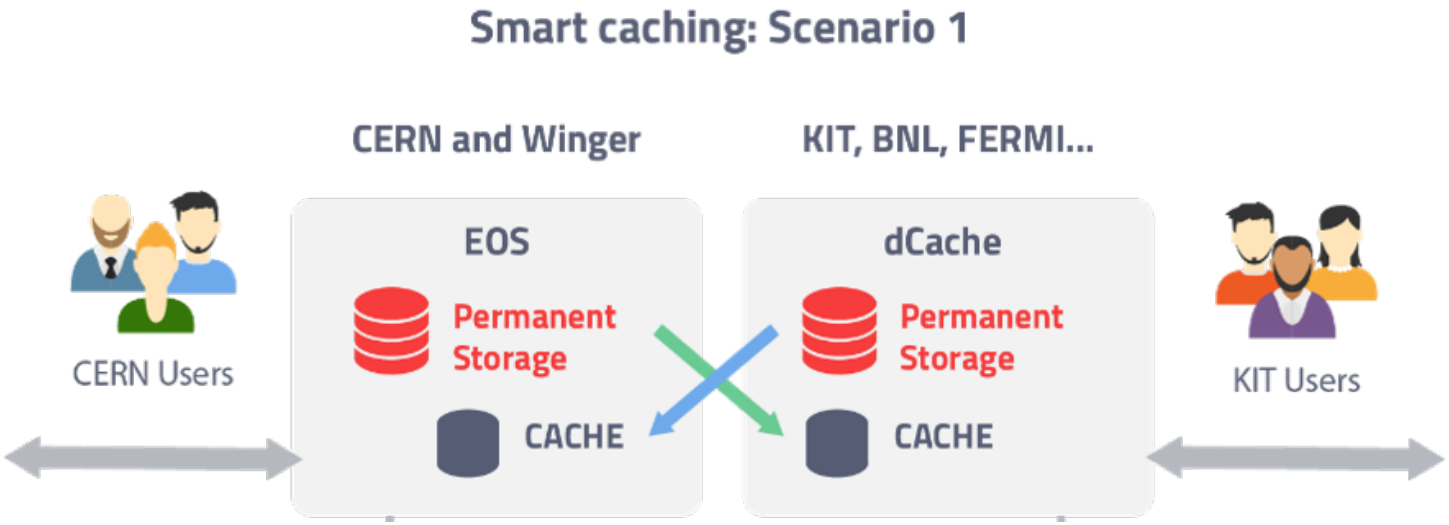

- The dynamic extension of a site to remote locations.
- Data stored in the original site should be accessible from the remote location in a "quasi"-transparent way from the clients' points of view.
- Implemented in EOS, ONEDATA and dCache using internal namespaces and algorithms. The cache is not addressable.

Smart caching scenarios – 2

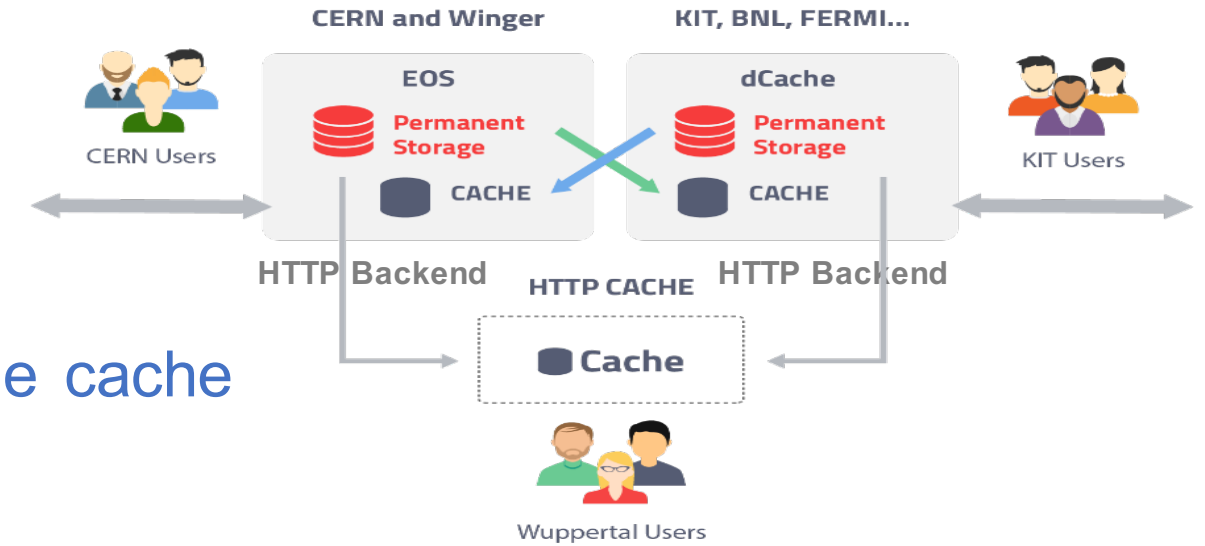

- A tactical storage set up as a stand-alone cache
  - e.g. in running squid-like services
- Clients access the cache directly
- The cache will fetch data on a miss (or at least redirect the client)
- The cache is federable, as it is directly addressable
- Cache federation at a site for scalability
- The cache namespace will be done via a federator that is not embedded into the storage systems (i.e Dynafed).

Smart caching scenarios - 3

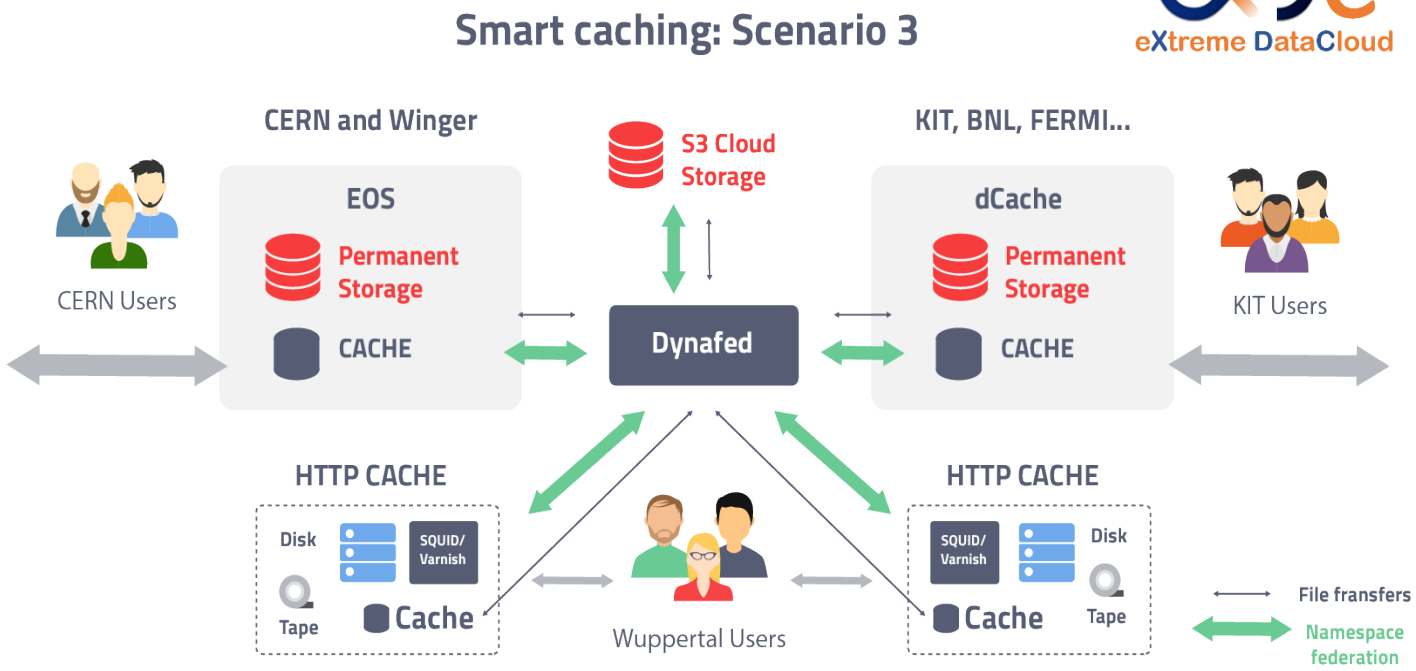

- The creation of a permanent "Virtual Data Cloud"
  - storage resources (Grid and Cloud) federated in a single namespace
  - remote data can be accessed transparently from any location without the need of explicitly copying them on the client location
- As an extension of the previous scenario, this implies the creation of a distributed and federated cache system

Overall architecture

- Caching …

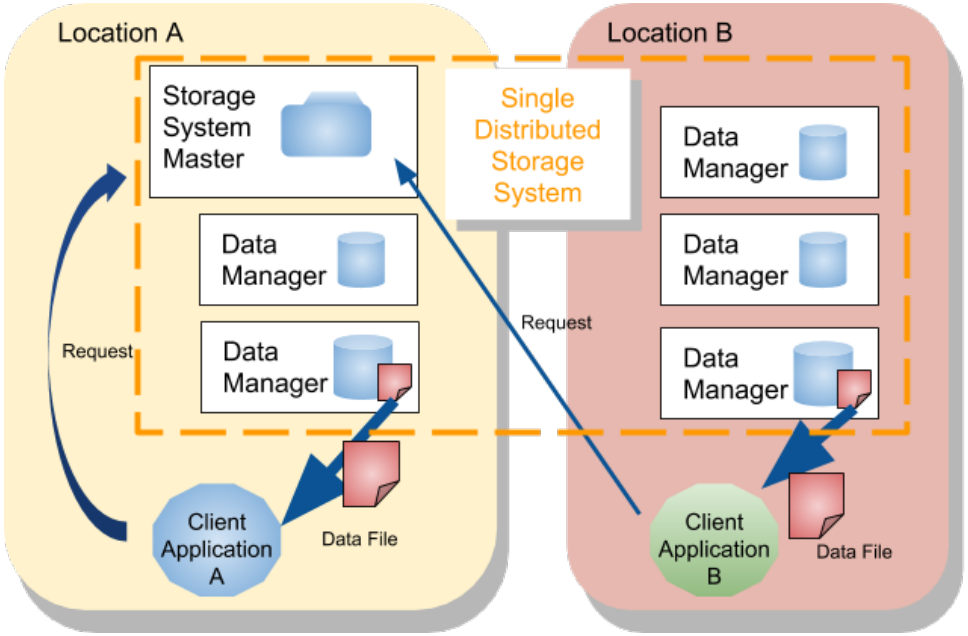

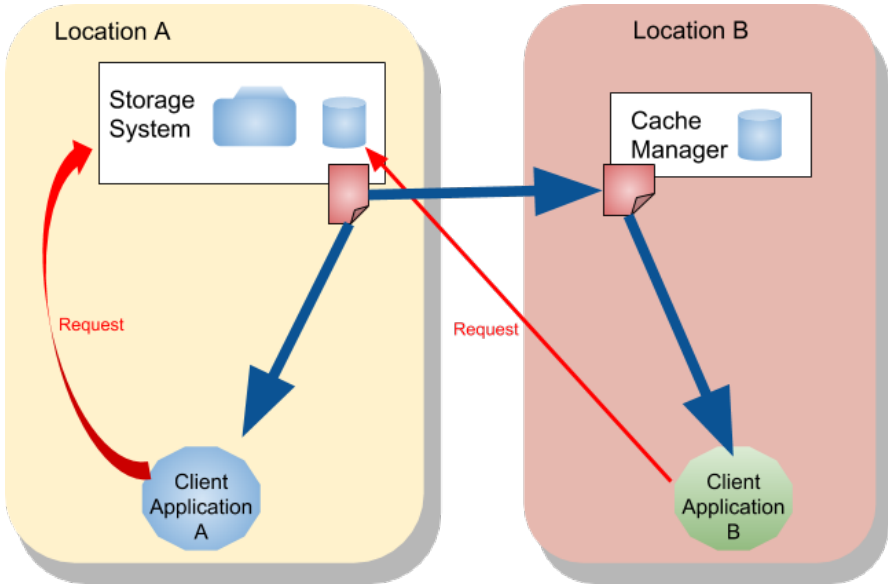

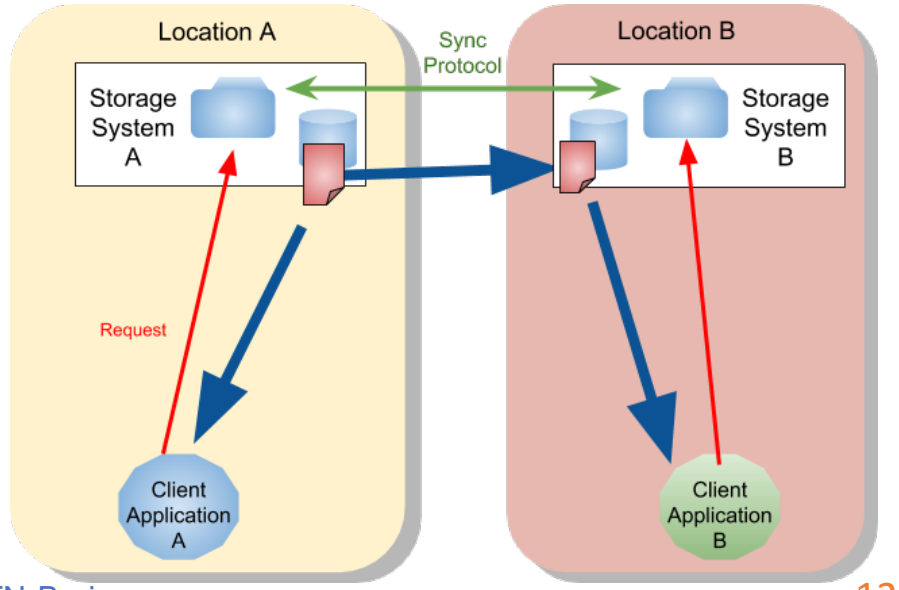

Onedata developments

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

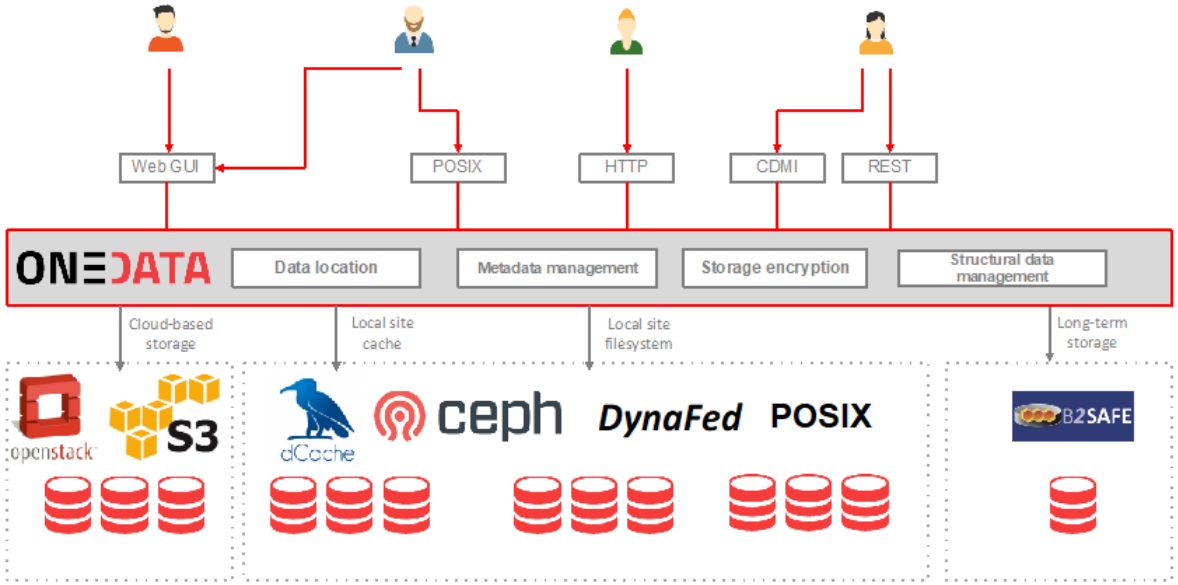

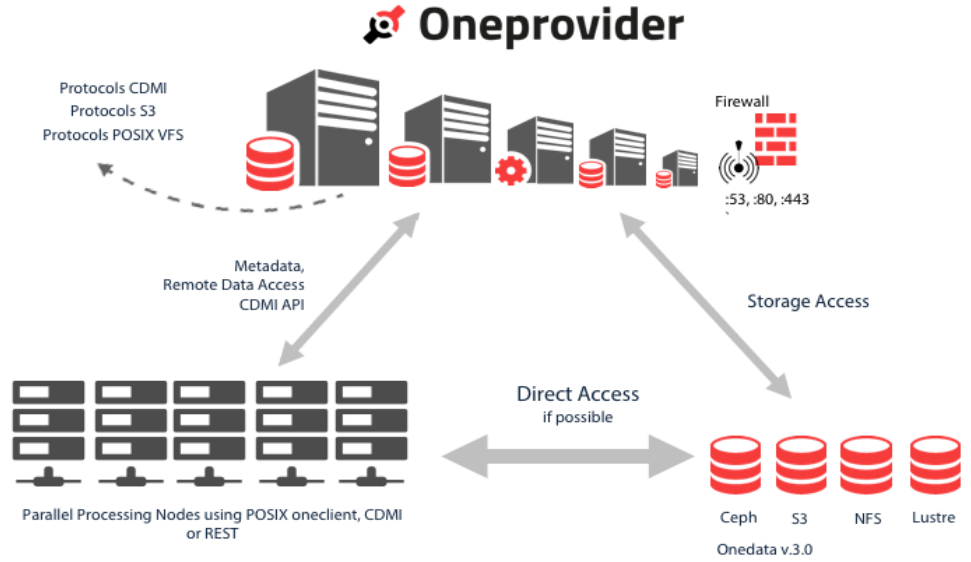

Onedata federations

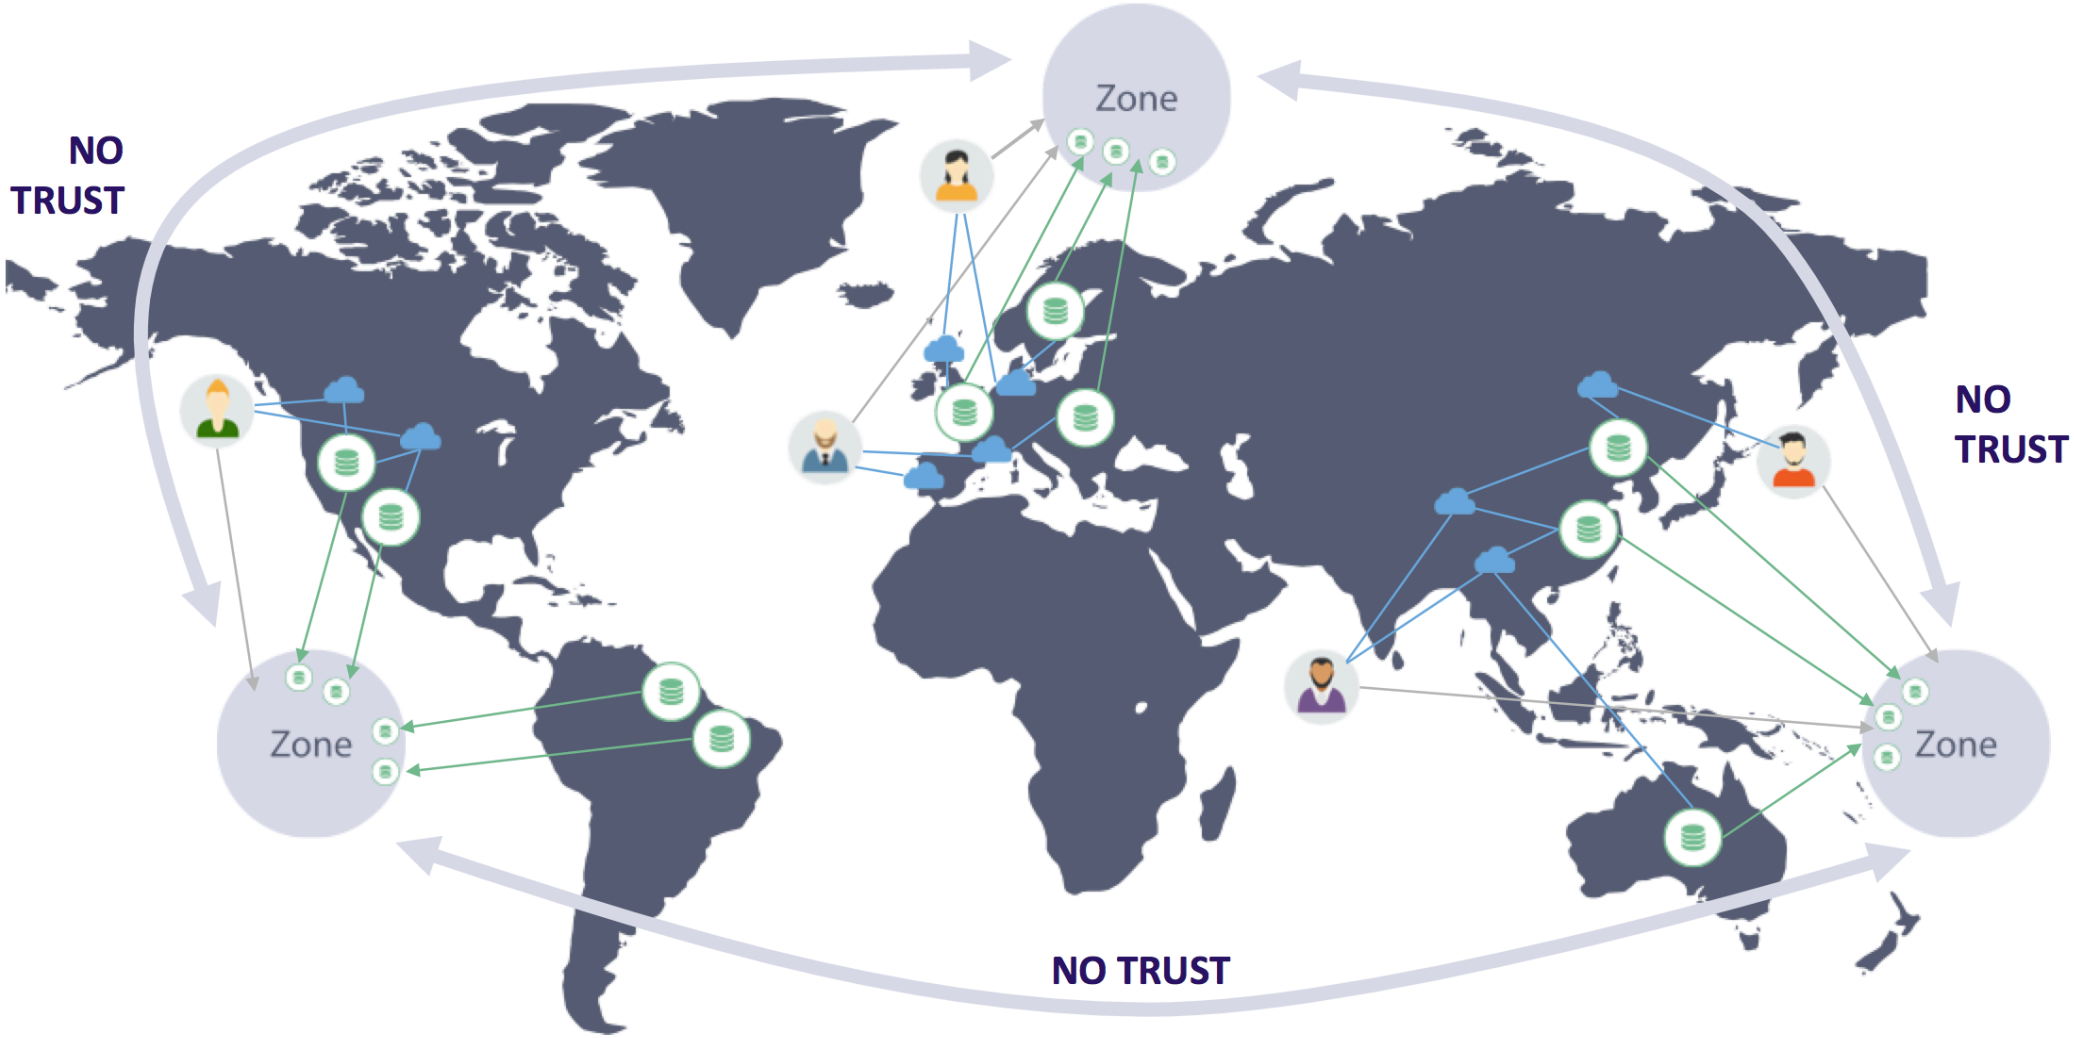

Metadata handling use cases

LIFEWATCH

- Metadata management to handle heterogeneous and large datasets
  - Different data types, formats, source and ways to access
  - e.g. Copernicus data: ~16PB per year
- Used as input for water quality forecasting systems
- Use of standards like EML (Ecological Metadata Language) and adopting best practices like FAIR+R principles

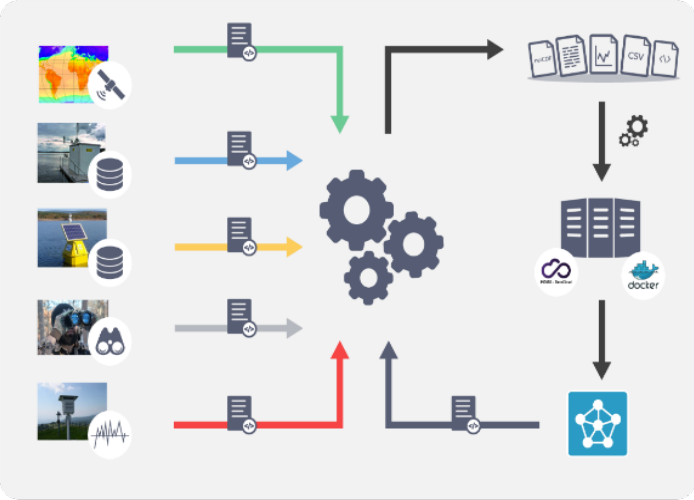

CTA

- The CTA distributed archive lies on the « Open Archival Information System » (OAIS) ISO standard. Event data are in files (FITS format) containing all metadata.
- Metadata are extracted from the ingested files, with an automatic filling of the metadata database.
- Metadata will be used for the further query of archive.
- The system should be able to **manage replicas**, tapes, disks, etc, with data from low-level to high-level.

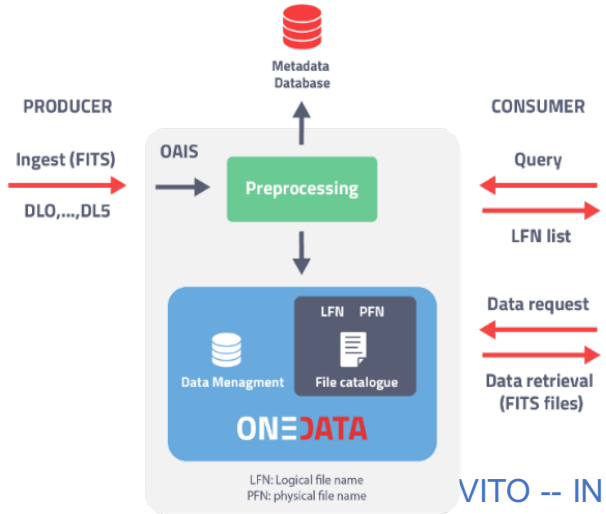

ECRIN

- Clinical trial data objects available for sharing with others
  - a variety of access mechanisms
  - wide variety of different locations
    - growing number of general and specialised data repositories
    - trial registries
    - Publications
    - the original researchers' institutions
  - 'discoverability' will become much worse in the future as more and more materials is made available for sharing

Status of the technical activities

- The single WP arch is mostly known
- The arch documents for both the WP is generally quite in a final state
- We are finalizing the release schedule for the first major release
- Most of the components are in a advanced state of testing and integration among them
- The integration between the INDIGO PaaS Orchestrator and Rucio, is in an advanced state
- We already have some demo on distributed caches (WP4.2) with specific protocols and underneath storage solutions

Status of the technical activities

- Now we have a more clear idea of the Onedata evolutions needed from the use cases
  - We still need few interactions: in order to better shape the fine details.
  - And to help the uses to correctly exploits the possibilities
- The interaction between Onedata and the rest of µServices platform, still require more deep and technical discussion.
  - Message bus, QoS, transport/sync protocols

Status of the technical activities: timing

- ***We are very late with the General architecture deliverable:***
  - D1.6 - Month 6 (R,PU) - General Architecture [INFN]
  - This is still work-in-progress
  - We need to clean up a bit few grey point in order to be able to complete it
- The services in WP4 are already working hard in order to follow the foreseen integration path, but the most interesting features, may be ready for a future upgrade after the fist major release
- Onedata is already working with the main use cases to facilitate the integration.

Overall architecture

- Unified architecture:

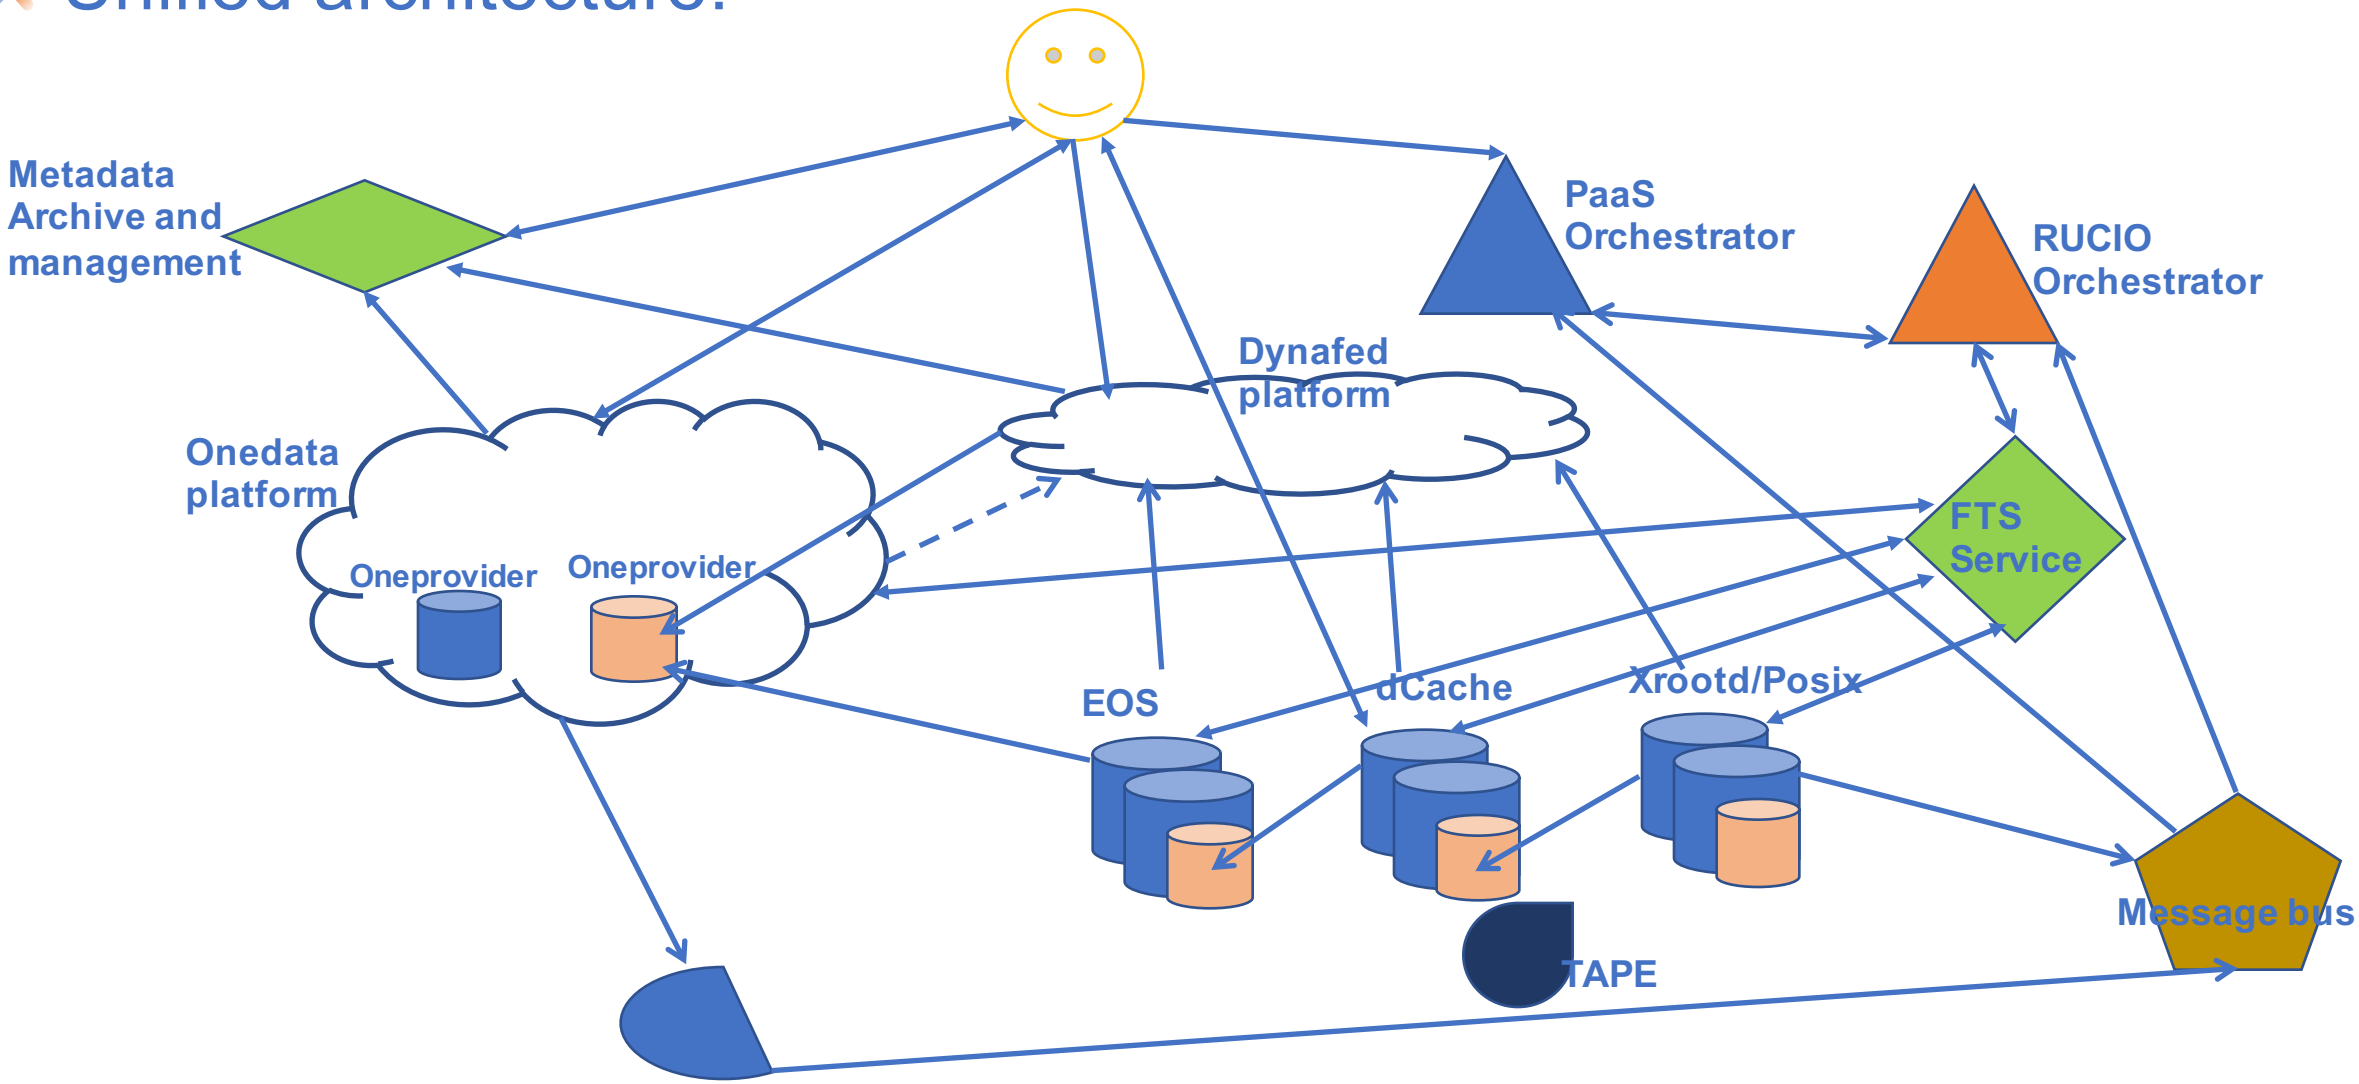

All hands objective about the technical activities

- Better understanding of the ingestion workflows needed by the end-users
- Final decision on the features and capabilities to be released in the XDC first major release
- Final view on the joint overall architecture so that we are able to finish the deliverable D1.6
- Final decision on specific item of integration between WP4 and WP5:
  - AAI relations between Onedata and the WP4 components
  - third party transfers between Onedata and others storage solutions (dCache, EOS, etc) using http
  - using Onedata as cache also for external storage (dCache, EOS, etc) using http
  - dealing with QoS
  - Storage notifications and how to deal with messages gathering from the storage and others components
- A more deep understanding of the needs of the user communities, and the possibility to work with them on concrete issues.
- **The goal is to have a real "working meeting" and not having tons of slides!!**

Conclusions

- I see all the technical people present and well involved
  - This is very good!!! ☺
- We have a small but very important session WP4+WP5
- Please exploit the JRA + user community session (ECRIN, Lifewatch, CTA, etc) in order to work on practical issues and possibly solve them!! ☺
- Remember that we can host non-foreseen side meeting, to work with colleagues on specific items,
  - You may ask Patrick at any time to know where you can stay
  - Please exploit this possibility as you need it