INDIGO - DataCloud

RIA-653549

INDIGO-DataCloud Quality of Service in Storage

Patrick Fuhrmann

For the INDIGO-DC Collaboration

Cremlin Connecting Russian and European Measures for Large-scale Research Infrastructures

INDIGO-DataCloud is co-founded by the Horizon 2020Framework Programme

Disclaimer

This initiative is funded until Sep 30, by the INDIGO-DataCloud project,
in order to standardize the interaction between the platform and the infrastructure layer in the storage area.
Funding will continue with the “eXtreme DataCloud” project, starting Nov 1, 2017.

Reminder: INDIGO-DataCloud

- **An H2020 project** approved in January 2015 in the EINFRA-1-2014 call
  - 11.1M€, 30 months (**from April 2015 to September 2017**)
- **26 European partners** in 11 European countries
  - Coordination by the Italian National Institute for Nuclear Physics (INFN)
  - Including developers of distributed software, industrial partners, research institutes, universities, e-infrastructures
- **Develop an open source Cloud platform** for computing and data ("DataCloud") tailored to science.
- **Targeting Multi-disciplinary scientific communities**
  - E.g. structural biology, earth science, physics, bioinformatics, cultural heritage, astrophysics, life science, climatology
- Deployable on **hybrid (public or private) Cloud infrastructures**
  - INDIGO = **IN**tegrating **D**istributed data **I**nfrastructures for **G**lobal Expl**O**itation
- **In response** to the technological **needs of scientists** seeking to easily exploit distributed Cloud/Grid compute and data resources.

INDIGO - DataCloud

Now about Quality of Service in Storage

The most obvious QoS in storage

- WLCG
  - Disk and Tape or Custodial, Online, Nearline, Replia
  - Plus transitions.
- Amazon
  - S3
  - Glacier
- Google
  - Standard
  - Durable Reduces Availability (DRA)
  - Nearline
- IBM and dCache
  - Various user defined storage classes including transitions

Mapping QoS to simple media

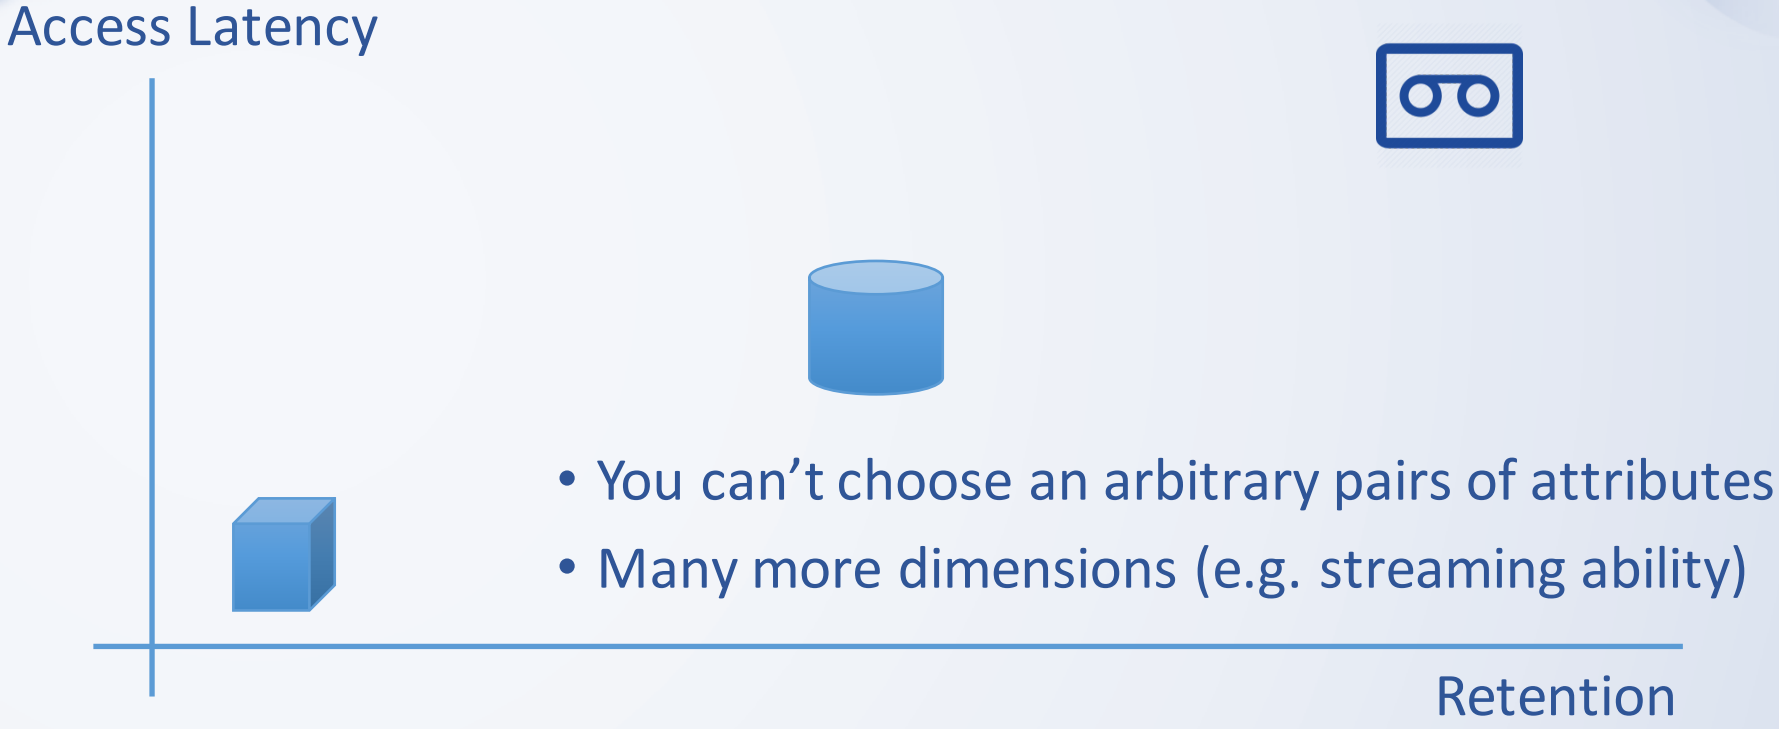

- You can’t choose an arbitrary pairs of attributes
- Many more dimensions (e.g. streaming ability)

Straw man's attempt to classify QoSiS

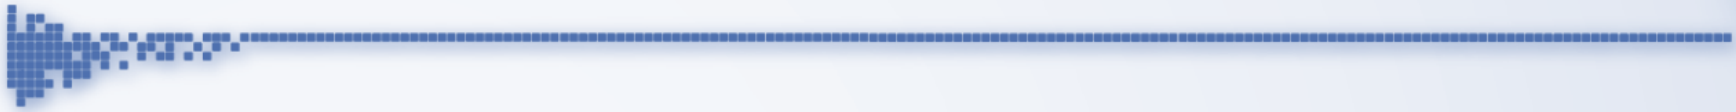

- Storage Media Quality (Durability , Retention)
- Storage Access Quality (Online, Nearline, Offline)
- Immutable, non-Immutable
- Which QoS transition are possible (disk->tape, Tape->disk, etc)
- Time depended QoS policy (Data Life Cycle)
- Access protocols, Authentication Protocols
- And many more (Reagan Moore was suggesting more than 200 of those)

Problems

- Definitions and access mechanisms are mostly proprietary, ambiguous and as a result not comparable between systems.
- Clients, if they want to use different systems, must know about all their different QoS levels and the way on how to access them.
- When running a procurement of cloud storage the required qualities are difficult to compare to the offered ones. See HNSciCloud.
- When writing a Data Management Plan e.g. in the context of an EU or national proposal, one needs to be an expert to understand what storage is required during the different phases of the data life cycle of you project.

At which point(s) can a more precise definition being helpful ?

- **Data Management Plans** would be **more precise** and much easier to write if precise definitions would exist which one could just refer to.
- It would **simplify the procurement process** if a storage service quality definitions would be more tuned towards scientific needs, supporting typical scientific data life cycles.
  - **Moreover after a decision of the EC in summer** : The rules on European standardisation allow the European Commission to identify information and communication technology (ICT) technical specifications - that are not national, European or international standards - to be eligible for referencing in public procurement. This allows public authorities to make use of the full range of specifications when buying IT hardware, software and services, allowing for more competition in the field and reducing the risk of lock-in to proprietary systems.

At which point(s) can a more precise definition being helpful ?

- On the technical level, a proper service quality definition plus a storage network control protocol would enable
  - Applications, platforms and frameworks to programmatically select the most appropriate storage service from different endpoints w/o knowing the specifics of the endpoint.
  - Storage broker systems (e.g. EGI or EUDAT) to select the cheapest storage endpoint for a particular use case.

Which could look like ....

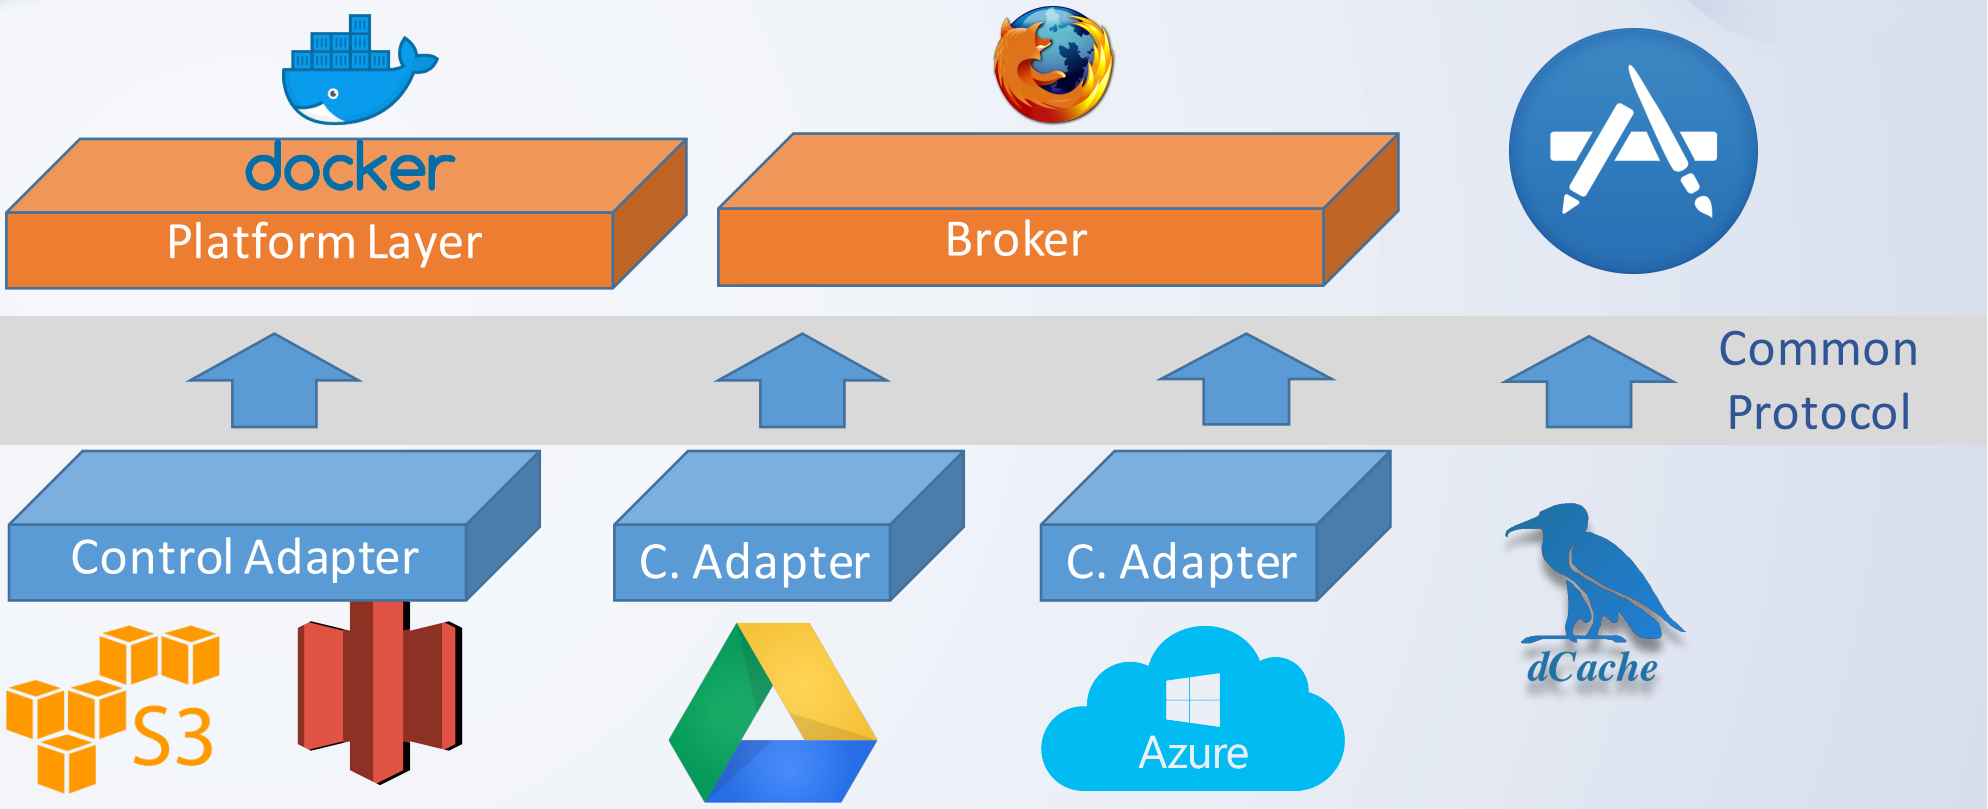

Now, how to tackle this ?

How would one tackle the issue ?

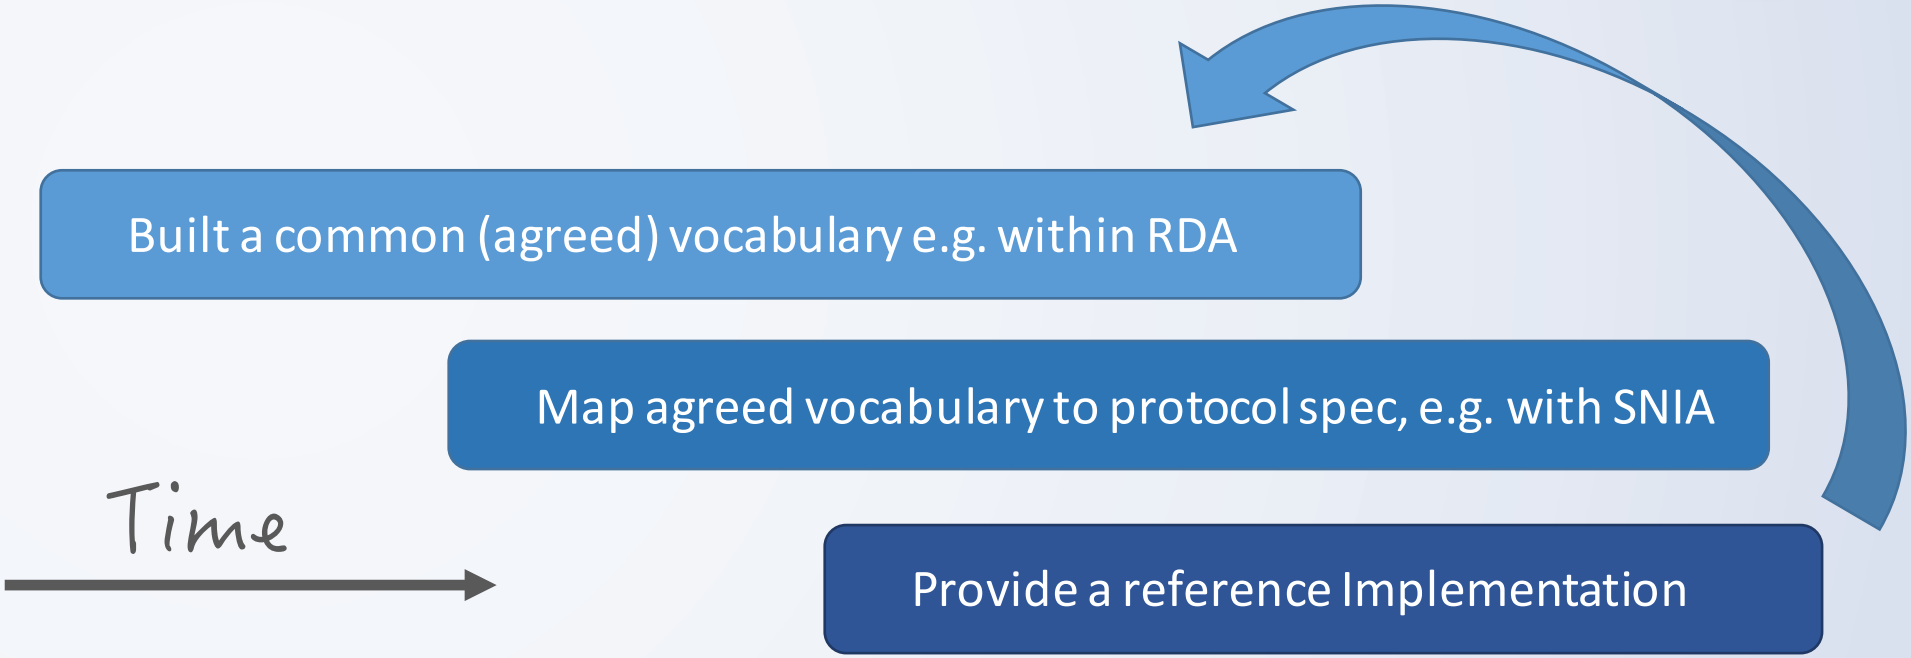

Since we only had 30 months and we would like to see some results before the EoP

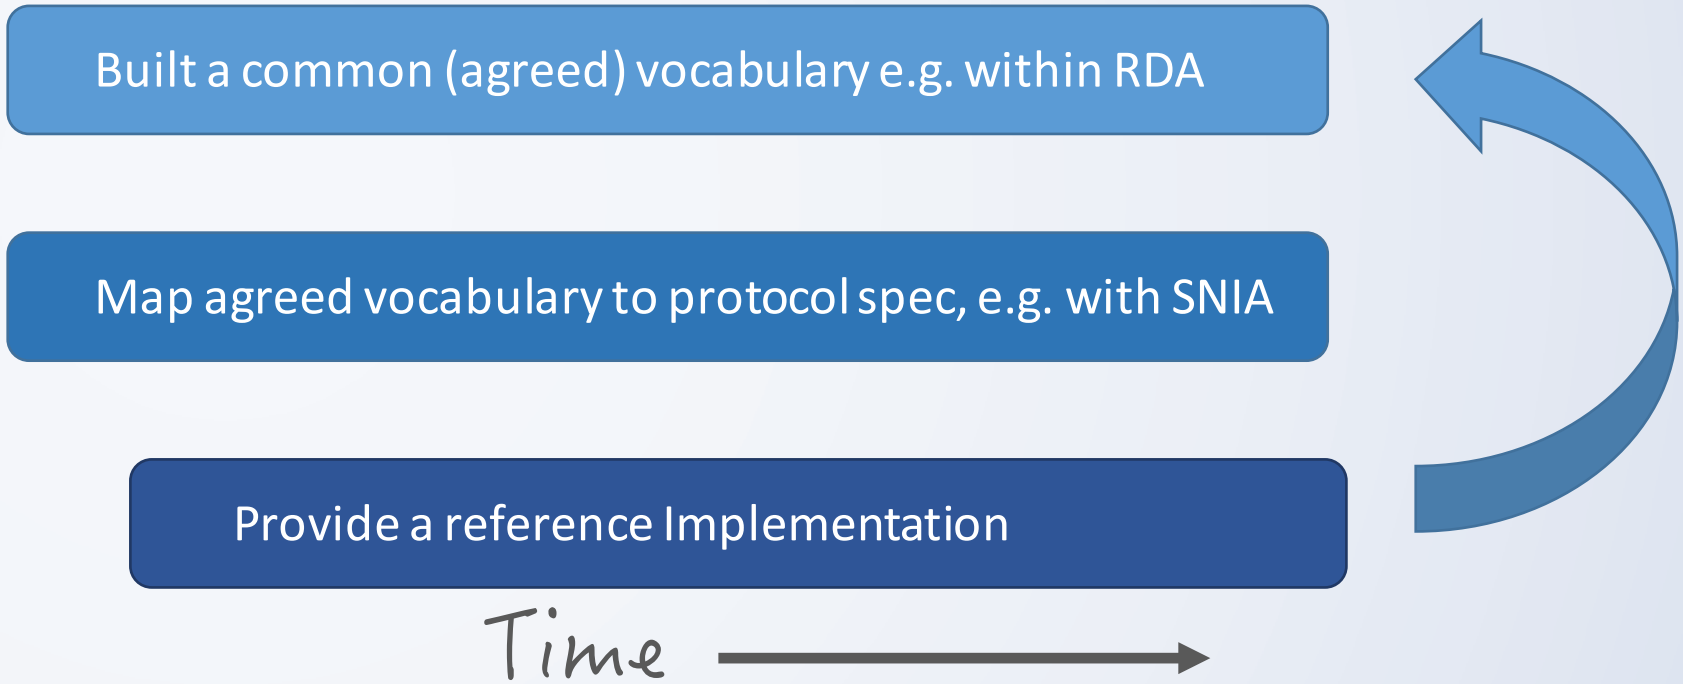

Recap

INDIGO - DataCloud

- Work on the 'Storage Service Definition' ontology with RDA
- Selecting a standard protocol to render the SSD ontology.
- Working on a server reference implementation.
- Building a 'show case'.


RDA
RESEARCH DATA ALLIANCE

RDA Work

- We submitted a proposal for a working group in RDA which was originally called : QoS in Storage and Data Life Cycle
- We provided uses cases and a case statement.
- The request was rejected
  - Name too long : now : “Storage Service Definition”.
  - Not enough participates (none) outside of the EU
  - No connection to other RDA groups, working on similar activities.
  - Bits and pieces
- At RDA 10 (last week) we started a new approach
- Turned out some Australian agencies had the same problem and were already working on a solution. So they joined our group.
- Good chances to get it approved this time. (Not giving up)
- Using tools to render our ontology.


SNIA
India
TM

The CDMI SNIA Part

SNIA : The Storage Networking Industry Association

The Storage Networking Industry Association (SNIA) is a non-profit organization made up of member companies spanning information technology. A globally recognized and trusted authority, SNIA's mission is to lead the storage industry in developing and promoting vendor-neutral architectures, standards and educational services that facilitate the efficient management, movement and security of information.

The CDMI SNIA Part

CDMI: Cloud Data Management Interface

The SNIA Cloud Data Management Interface (CDMI) is an ISO/IEC standard that enables cloud solution vendors to meet the growing need of interoperability for data stored in the cloud. The CDMI standard is applicable to all types of cloud – private, public and hybrid. There are currently more 20 products that meet the CDMI specification.

INDIGO Products on SNIA Web Pages

Shipping Commercial CDMI Servers

Arsys CloudStorage (Powered by Scality)

Indigo Project - Storage Quality of Service and Data Lifecycle

Coho Data

Compuverde Object Store

Critical Path Messaging Platform (Powered by Scality)

DDN WOS

Mezeo MezeoCloud (Zimbra)

NetApp StorageGRID 9

NetApp StorageGRID Webscale

ProphetStor

Scality Ring

SGI OmniStor (Powered by Scality)

Tarmin GridBank

XOR Systems - Cloud Aqua

Open Source CDMI Servers

CDMI-Server

dCache

FI-WARE Project

JClouds

OpenStack Swift

SNIA CDMI Reference Implementation

Stoxy

Venus-C

Protocol decision (CDMI)

- The decision to use CDMI (SNIA) as the QoS control protocol was already made at the time of the proposal.
- Very difficult to change this decision, as CDMI was the only industry standard, somehow working in our direction.
- So we joined SNIA
- And contributed to the CDMI reference implementation.
- We actually moved it into GitHub and made is usable.
- Although we started with the SNIA reference implementation, we had to rewrite a large part of it.

CDMI Considerations

- CDMI is not very widely spread.
- CDMI doesn't cover our use cases.
- But CDMI provides the possibility of 'extensions', which we are using.
- Based on our experience with WLCG (Storage Resource Manager) we have a much better idea on how to define those protocols than SNIA.
  - QoS in CDMI is very much shoehorned.
  - Multi user QoS transitions are not mapped correctly.
  - INDIGO, based on its DoW was bound to CDMI.
- INDIGO is going on SNIA's nerves. ☺
- We agreed on an extension to CDMI, covering our initial thoughts.

News on CDMI and SNIA

- SNIA accepted the INDIGO reference implementation.
- Now available from the SNIA github repository.

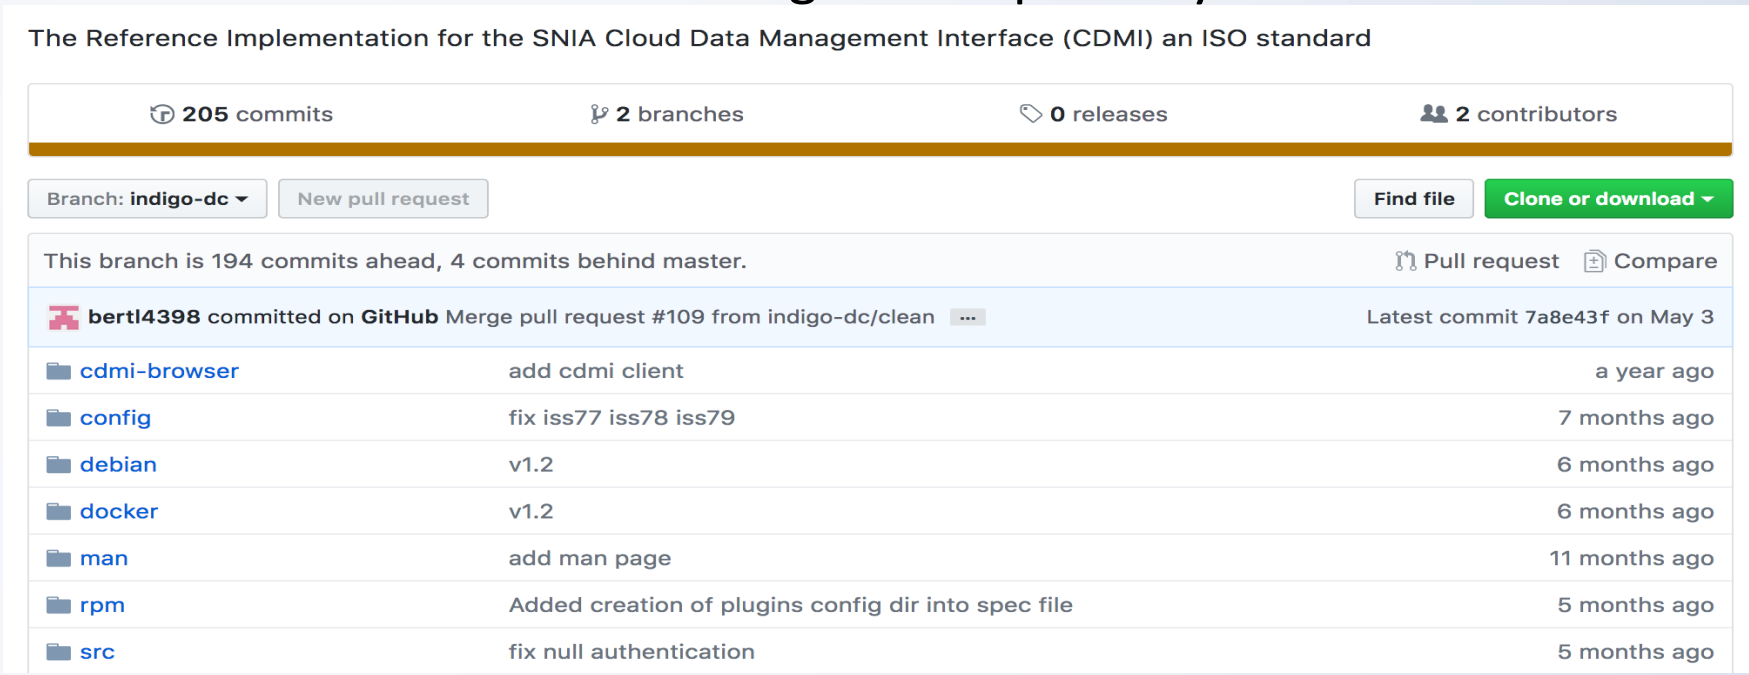

Installation

INDIGO-DataCloud CDMI Server

This project ports the SNIA CDMI-Server reference implementation to a Spring Boot application.

Requirements

- JDK 1.8+
- Maven 3+
- (optional) Redis

Build & Run & Configure

The project uses the Maven build automation tool that will build one fat jar Spring Boot application.

It depends on the cdmi-spi Java SPI cdmi-spi library.

The cdmi-spi library is provided at http://cdmi-qos.data.kit.edu/maven/ and should be included automatically.

Our reference design

The Architecture

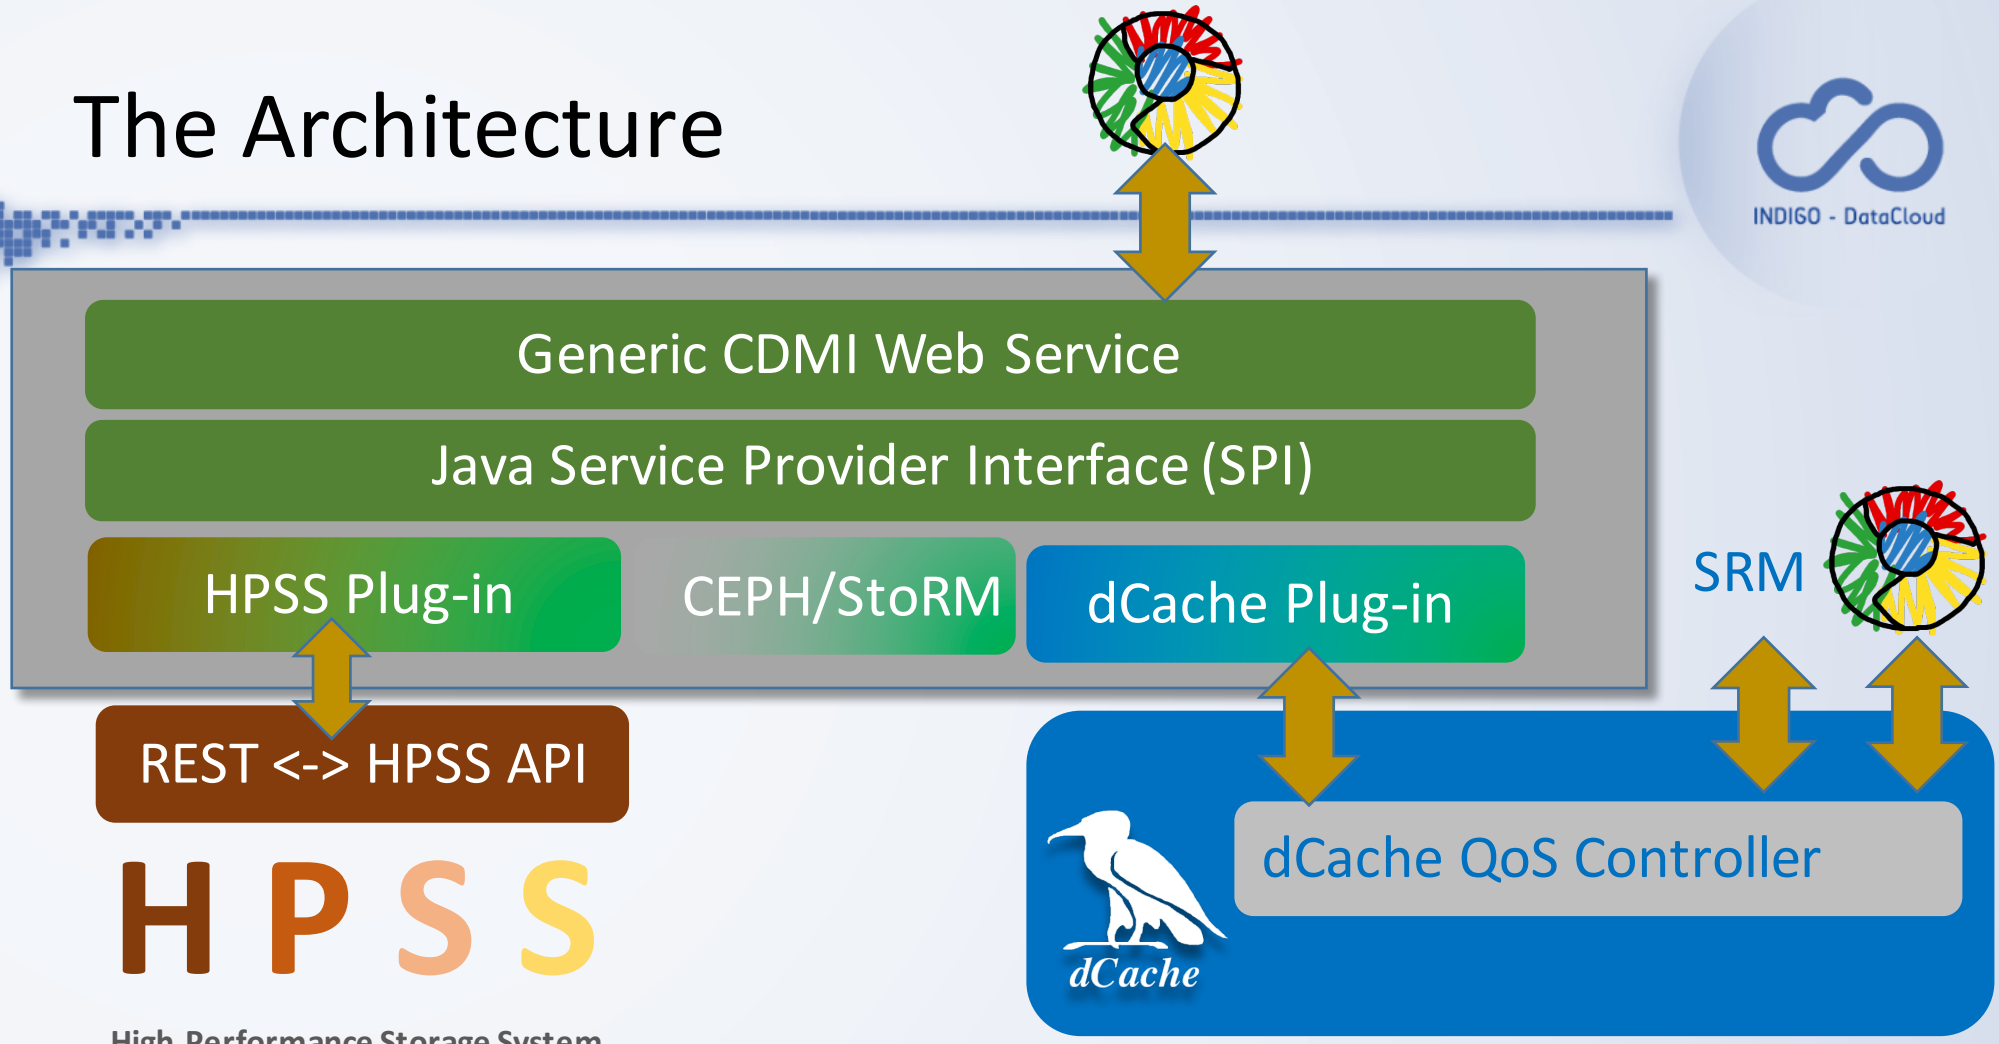

How is this done in dCache ?

dCache Media Control

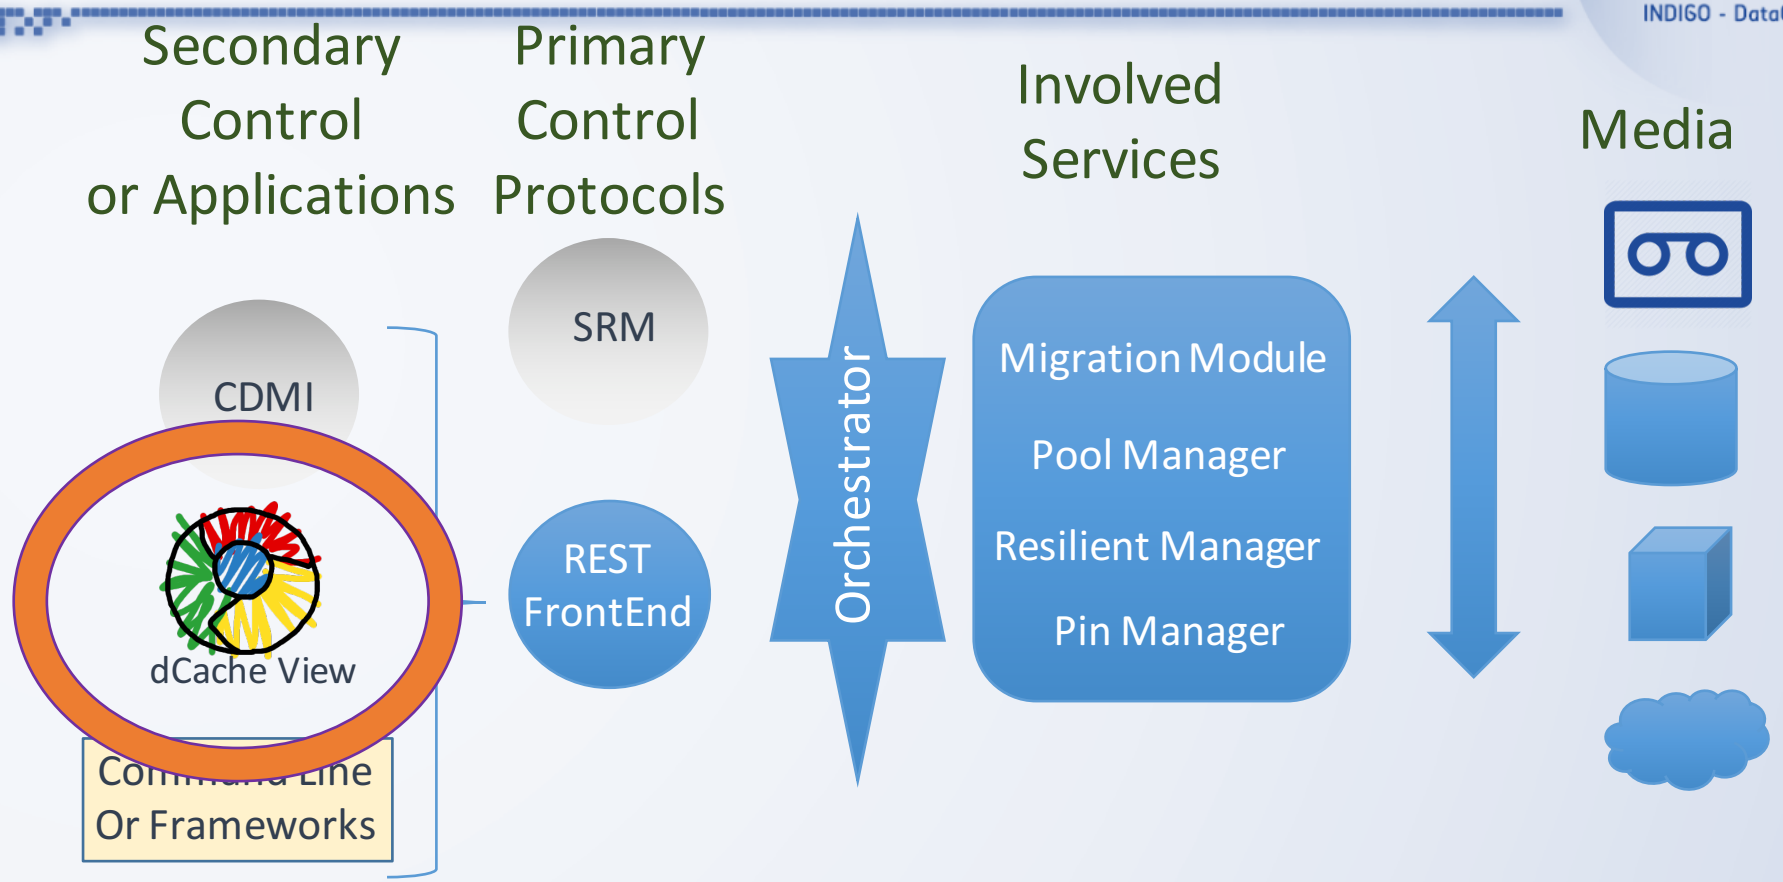

QoS dCache View (Real Screenshot)

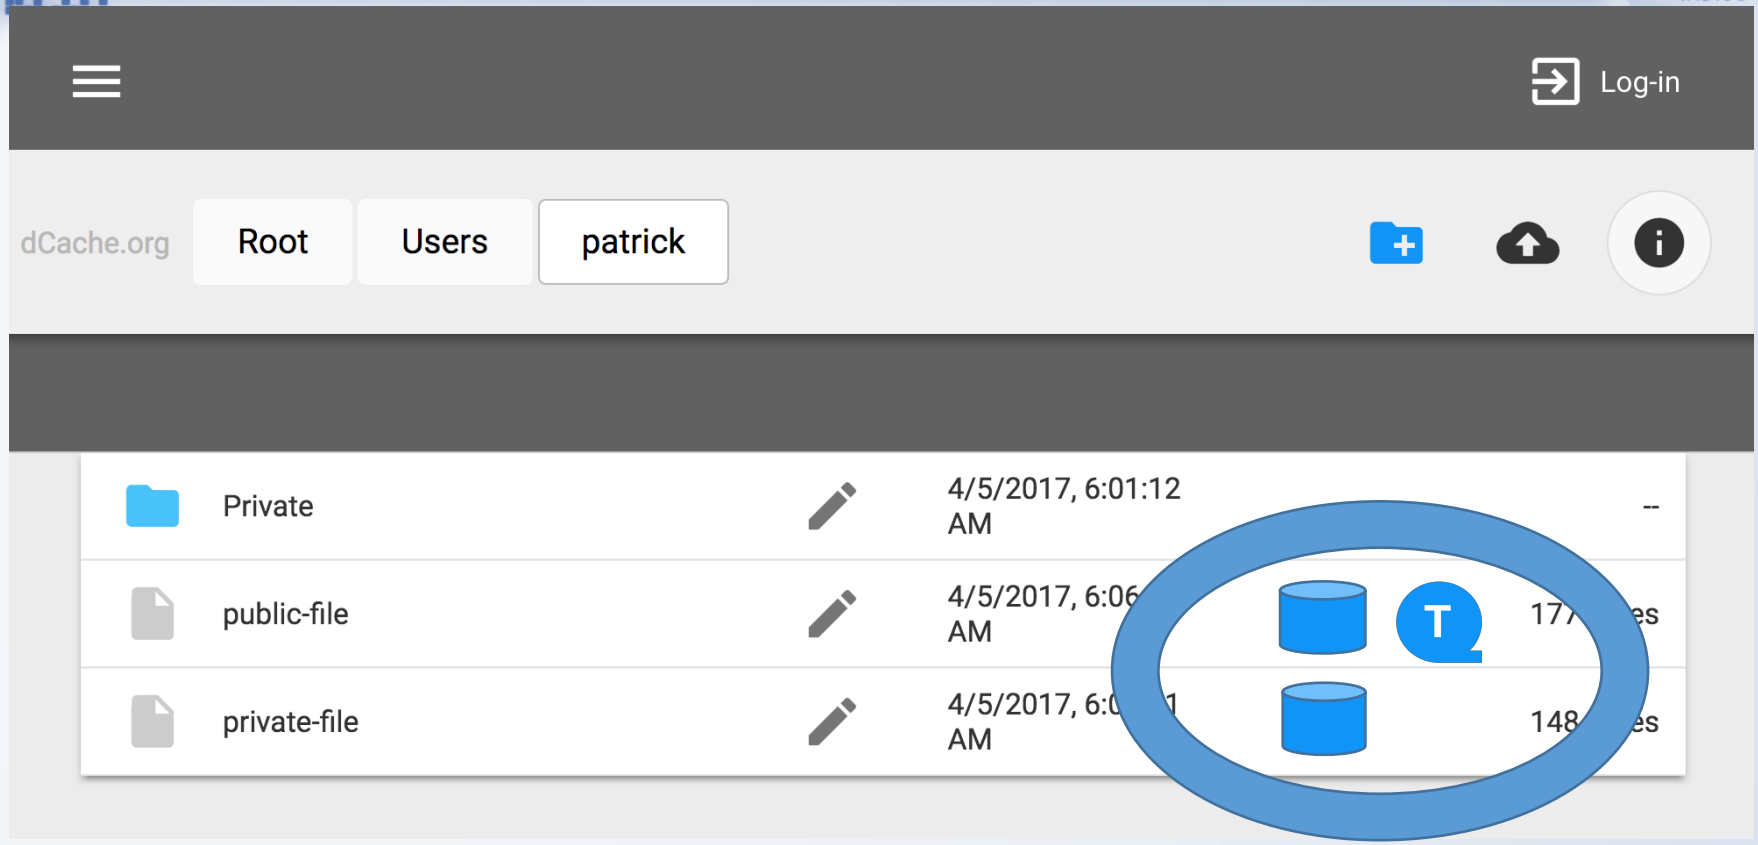

Supported Transitions in dCache

INDIGO Evaluation Deployment

- KIT (master server)
- KIT (GPFS, HPSS, mixed Tape, Disk)
- CNAF (StoRM)
- DESY (dCache, mixed Tape, Disk)
- Poznan (CEPH)

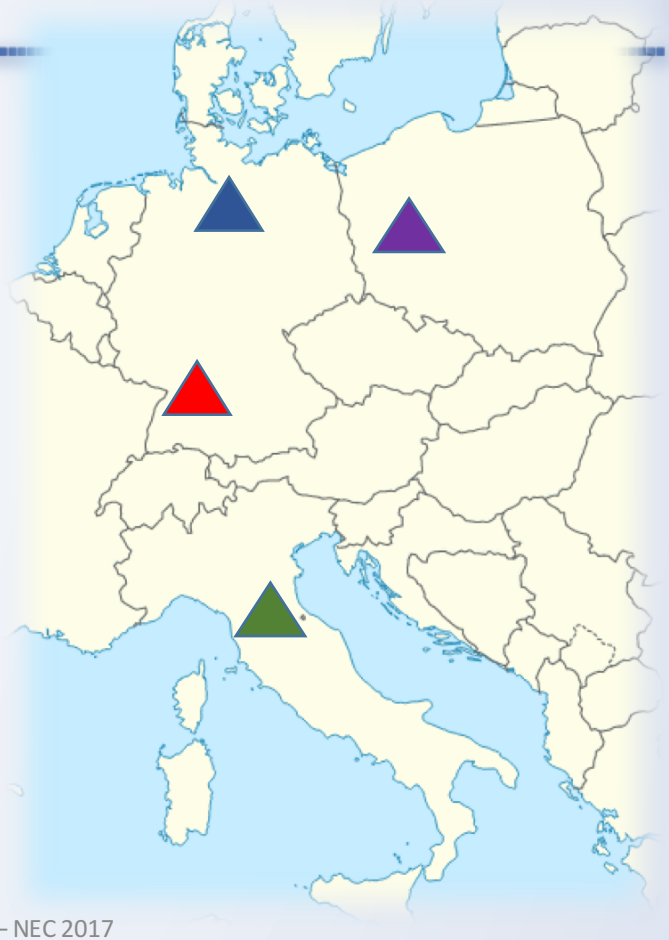

The Setup

- Each site is using a different storage technology, but exposing details through the CDMI extension.
- We created an example landing page, a service, frequently scanning the known endpoints for QoS details and render them ‘human readable’
- We are using proper Open ID Connect authentication for the endpoints and the landing page.
  - You may use your Google ID, or INFN/CNAF IAM service.

The European Storage Landing Page

INDIGO - DataCloud

Available Qualities of Storage

	Name	Access Latency [ms]	Number of Copies	Storage Lifetime	Location	Storage type	Available Transitions
PSNC	DataobjectProfile1	3000	1	20 years	PL	Archival	
PSNC	DataobjectProfile2	2000	2	20 years	PL, UK	Archival	
PSNC	DataobjectProfile3	500	3	20 years	NL, ES, PL	Archival	
DISK	disk	100	1		DE	Processing	tape, disk+tape
DISK	disk+tape	100	2		DE	Processing	tape
KIT	DiskAndTape	50	3	20 years	DE	Processing	TapeOnly
INFN	DiskAndTape	50	2		IT	Processing	
KIT	DiskOnly	50	3	20 years	DE	Processing	
INFN	DiskOnly	50	1		IT	Processing	
	profile1	10	3	20 years	DE	Processing	profile2
	profile2	10000	2		DE	Archival	profile1
DISK	tape	600000	1		DE	Archival	disk+tape

Selecting DESY, DISK

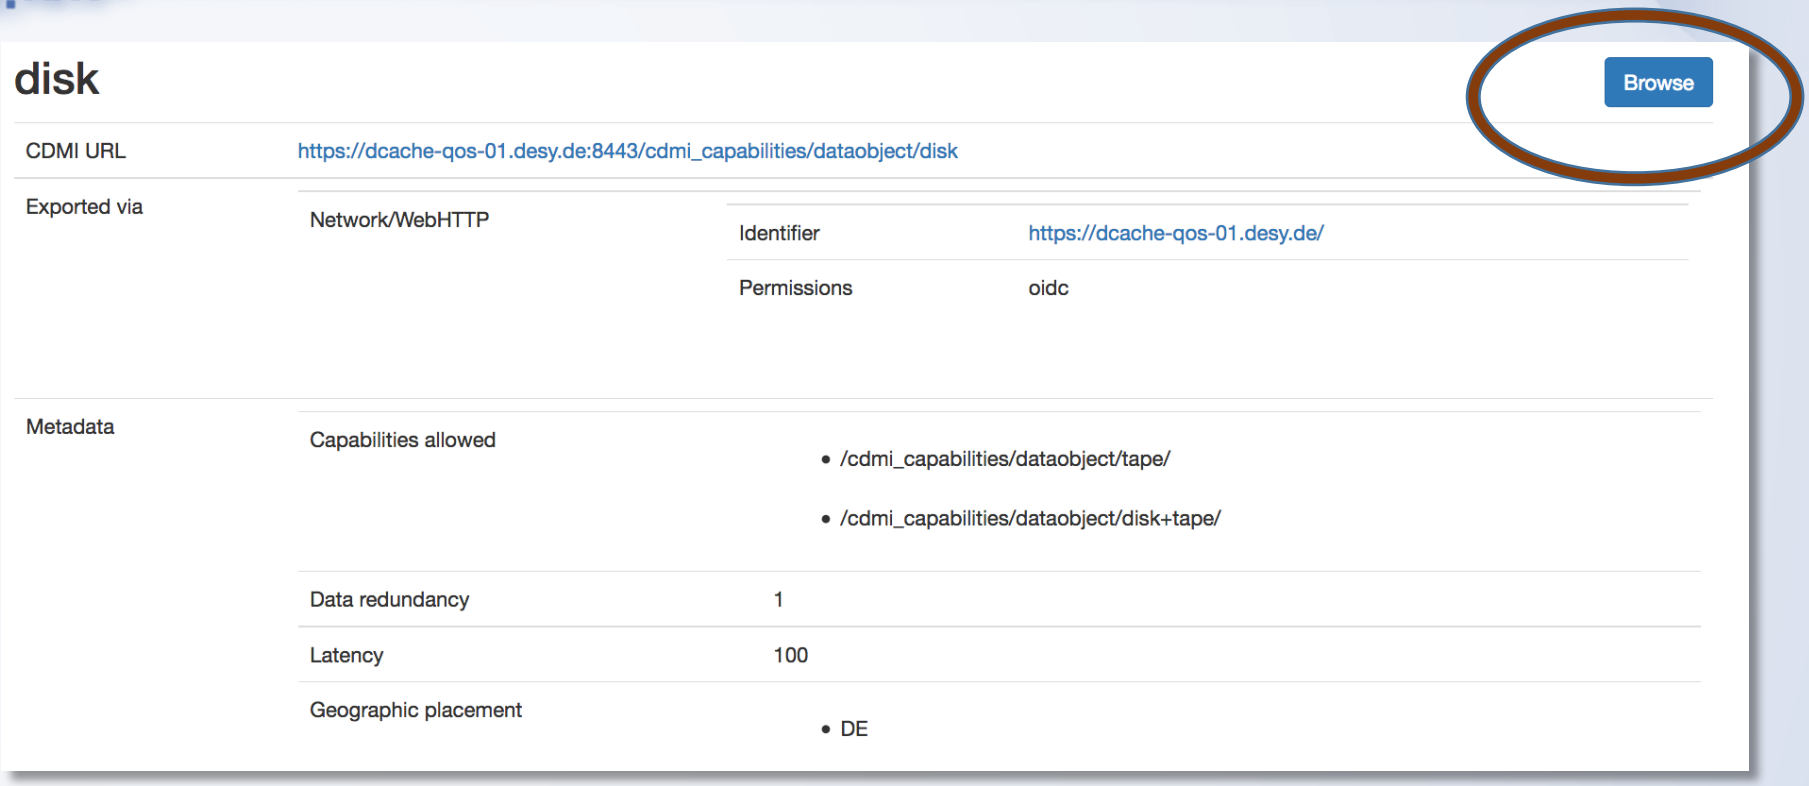

Browsing selection

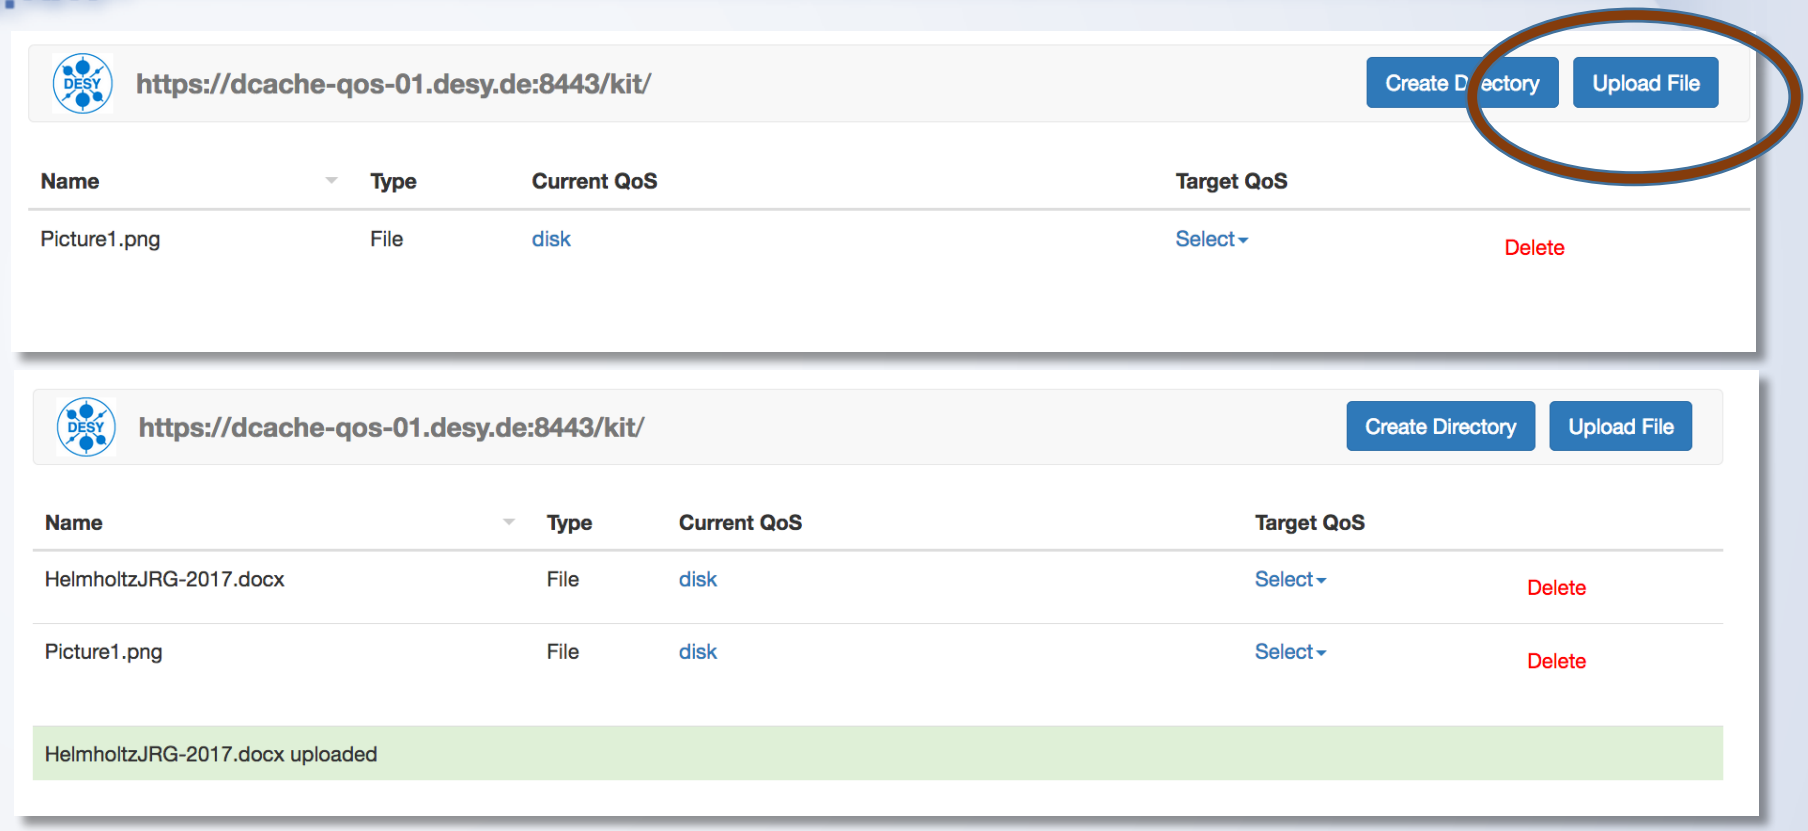

Conclusion

- A ‘Storage Service Definition’ agreement would be of benefit in various areas.
- We try to agree community wide on such a definitions with RDA.
- Process with SNIA is painful but helps to understand the difficulties to map our ideas to a real protocol.
- SNIA now supports our code improvements and accepted our improved reference implementation.
- Implementing the protocol plug-ins helps to understand the issues with the different storage systems.
- We already now support a limited number of transitions.
- Our European Reference System is working. (for DEMO purposes)
- We hope to attract people and slowly get to something which is of benefit for most sciences.


The End
https://www.indigo-datacloud.eu
Better Software for Better Science.