




INDIGO - DataCloud

RIA-653549

# First implementation of our ideas by INDIGO-DataCloud



Patrick Fuhrmann



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

# From the Paper “Advances in Cloud”

## EC Expert Group Report on Cloud Computing,

- <http://cordis.europa.eu/fp7/ict/ssai/docs/future-cc-2may-finalreport-experts.pdf>

The Expert Group identified the following unsolved issues in the currently deployed ecosystem :

- Open Interoperation across (proprietary) Cloud solutions at IaaS, PaaS and SaaS level has not yet been developed.
- No solutions are available to manage multitenancy at large scale and heterogeneous environments.
- No dynamic and seamless elasticity from in-house Cloud to public Clouds ...
- Data management: Problems with bandwidth, security and privacy between public and private clouds.

# INDIGO-DataCloud



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 **Exp**loitation
- **In response** to the technological **needs of scientists** seeking to easily exploit distributed Cloud/Grid compute and data resources.



# User (Scientist) first



# Users first: from here...



INDIGO - DataCloud



## Use-Cases from

LifeWatch
EuroBioImaging
INSTRUCT
LBT
CTA
WeNMR
ENES
eCulture
ELIXIR
EMSO
Dariah
WLCG

100 distinct requirements

Converted to concrete activities in the Project DoW

## Computational

- Software as a Service
- Execution of Workflows
- Cloud Bursting
- X-Site Execution
- Improved Scheduling
- Access to GP-GPU's

\*\*\*

## Storage

- Distributed Storage, accessible via POSIX
- Persistent Data Storage

\*\*\*

## Infrastructure

- Global Level AAI
- Software Defined Networks

\*\*\*

# ... to here ... Community Case Studies



INDIGO - DataCloud

**Report on how several scientific communities are implementing their own requirements into concrete applications using INDIGO-DataCloud components.**

- Monitoring and Modelling Algae Bloom in a Water Reservoir
- TRUFA (Transcriptomes UserFriendly Analysis)
- Medial Imaging Biobanks
- Molecular Dynamics Simulations
- Astronomical Data Archives
- Archive System for the Cherenkov Telescope Array (CTA)
- HADDOCK Portal
- DisVis
- PowerFit
- Climate models inter comparison data analysis
- eCulture Science Gateway
- EGI FedCloud Community Requirements
- ELIXIR-ITA: Galaxy as a Cloud Service
- MOIST – Multidisciplinary Oceanic Information System
- Data Repository platform for DARIAH

<https://www.indigo-datacloud.eu/documents-deliverables>

... t



Report  
scient  
are im  
own re  
concr  
us  
DataCl



### Using “Champion” approach :

Communities have to provide a scientist, becoming an expert in computing and INDIGO terminology.



INDIGO - DataCloud

# Now about QoS in Storage



# Where it came from ...

## Amazon

- S3 : online

- Glacier : nearline

## Google

- Standard

- Durable Reduces Availability (DRA)

- Nearline

## IBM (HPSS, GPFS)

- Storage classes (user defined)

## dCache

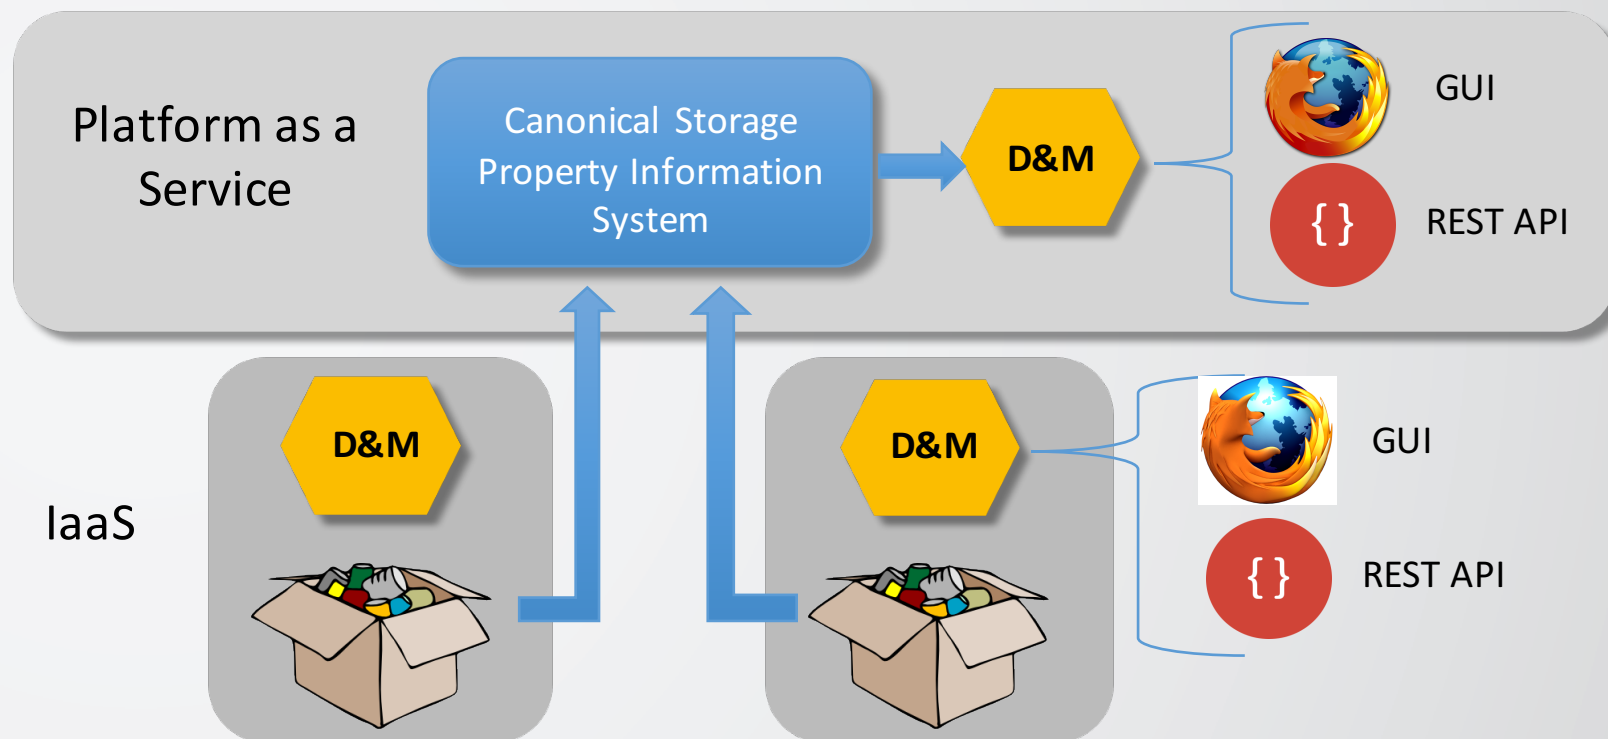
- Storage groups (user defined)

- Tape

- Disk (spinning or SSD)

- Resilient Management ('n' copies)

# The first ideas for a solution



# Theoretical approach .....



INDIGO - DataCloud



Built a common (agreed) vocabulary e.g. within RDA

Map agreed vocabulary to protocol spec, e.g. with SNIA

Provide a reference Implementation



# Since we only had 30 months and Augusto would like to see some results ...

Built a common (agreed) vocabulary e.g. within RDA

Map agreed vocabulary to protocol spec, e.g. with SNIA

Provide a reference Implementation

Time →

# Protocol decision

- 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 it usable.
- Although we started with the SNIA reference implementation, we had to rewrite a large part of it.

# 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



INDIGO - DataCloud

## 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 clouds – 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



# CDMI Considerations


- CDMI is an industry standard.
- 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. 😊

# The Architecture




HPSS

High Performance Storage System



# Evaluation Deployment

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























INDIGO - DataCloud

# Federated View (Real Screenshot)

## Available Qualities of Storage

	Name	Access Latency [ms]	Number of Copies	Storage Lifetime	Location	Storage type	Available Transitions
	disk	100	1		DE	 Processing	
	tape		2		DE	 Processing	
	DiskAndTape	10	3	20 years	DE	 Processing	TapeOnly
	DiskOnly	10	3	20 years	DE	 Processing	
	profile1	10	3	20 years	DE	 Processing	profile2
	profile1	100	2		DE	 Processing	profile2
	profile2	10000	2		DE	 Archival	profile1
	profile2	100	2		DE	 Processing	profile1
	tape	600000	1		DE	 Archival	
	TapeOnly	50000	2	20 years	DE	 Archival	DiskAndTape

# Generic QoS (Details, Real Screenshot)

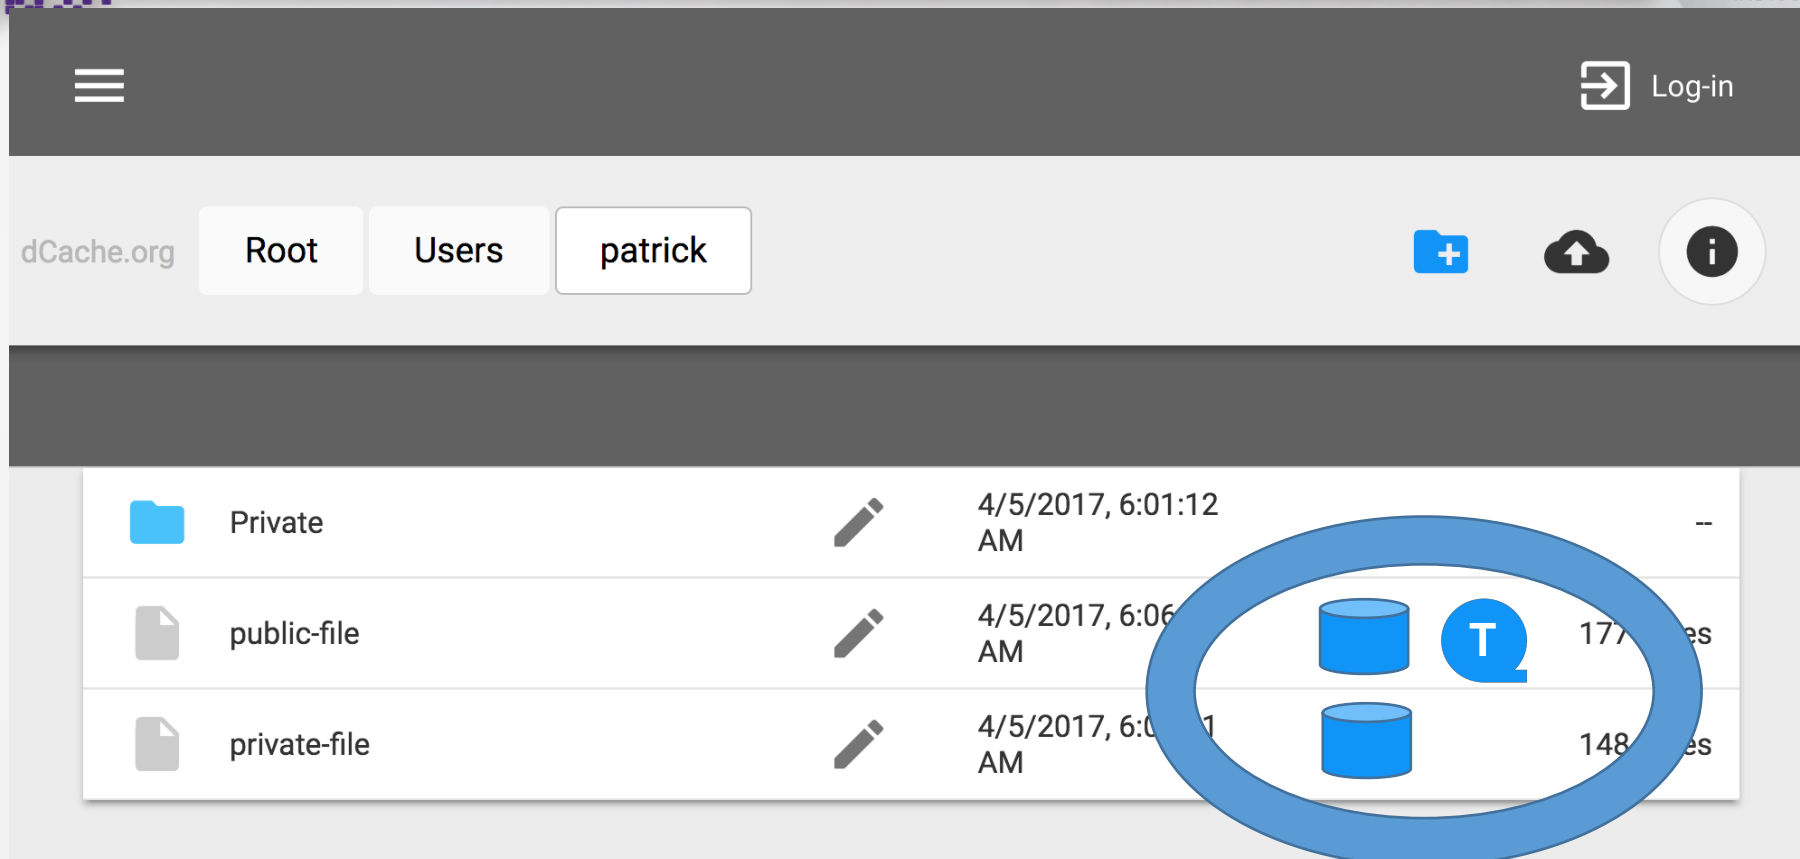


INDIGO - DataCloud

KIT / DiskOnly

Data storage lifetime	20 years
Latency	10
Throughput	4194304
Capability lifetime action	migrate-to:/cdmi_capabilities/dataobject/DiskAndTape
Capabilities allowed	
Capability lifetime	0:30:00
Geographic placement	<ul style="list-style-type: none"><li>• DE</li></ul>
Data redundancy	3

# QoS dCache View (Real Screenshot)

A screenshot of the dCache.org web interface. The top navigation bar includes a hamburger menu, a "Log-in" button, and the site name "dCache.org". Below this is a breadcrumb trail with buttons for "Root", "Users", and "patrick". To the right of the breadcrumb are icons for adding a folder, uploading a file, and a user profile. The main content area displays a table of files and folders. A blue circle highlights the "public-file" and "private-file" entries, which show database icons and a "TQ" label. The "Private" folder entry is also visible.

Icon	Name	Icon	Timestamp	Size	Details
Folder	Private	Pencil	4/5/2017, 6:01:12 AM	--	
File	public-file	Pencil	4/5/2017, 6:06:11 AM	177 MB	177 MB
File	private-file	Pencil	4/5/2017, 6:06:11 AM	148 MB	148 MB

# Conclusion



INDIGO - DataCloud

- Apologies for not doing it the “right way”
- But we had to provide an implementation within 30 months.
- However, we can prove that we are serious.
- Process with SNIA is painful but helps to understand the difficulties, to map our ideas to a real protocol.
- Implementing the protocol helps to understand the issues with the different storage systems.
- We even now support limited transitions.
  - Dangerous !!! (Tape is a tricky beast 😊 )



INDIGO - DataCloud

# The End

<https://www.indigo-datacloud.eu>

**Better Software for Better Science.**





INDIGO - DataCloud

- Brokering IG and Brokering WG
- Vocabulary groups
- NEW : data preservation :
- Overlap with data management plans.