




INDIGO - DataCloud

RIA-653549

Quality of Service in Storage by INDIGO-DataCloud



Anupam Ashish



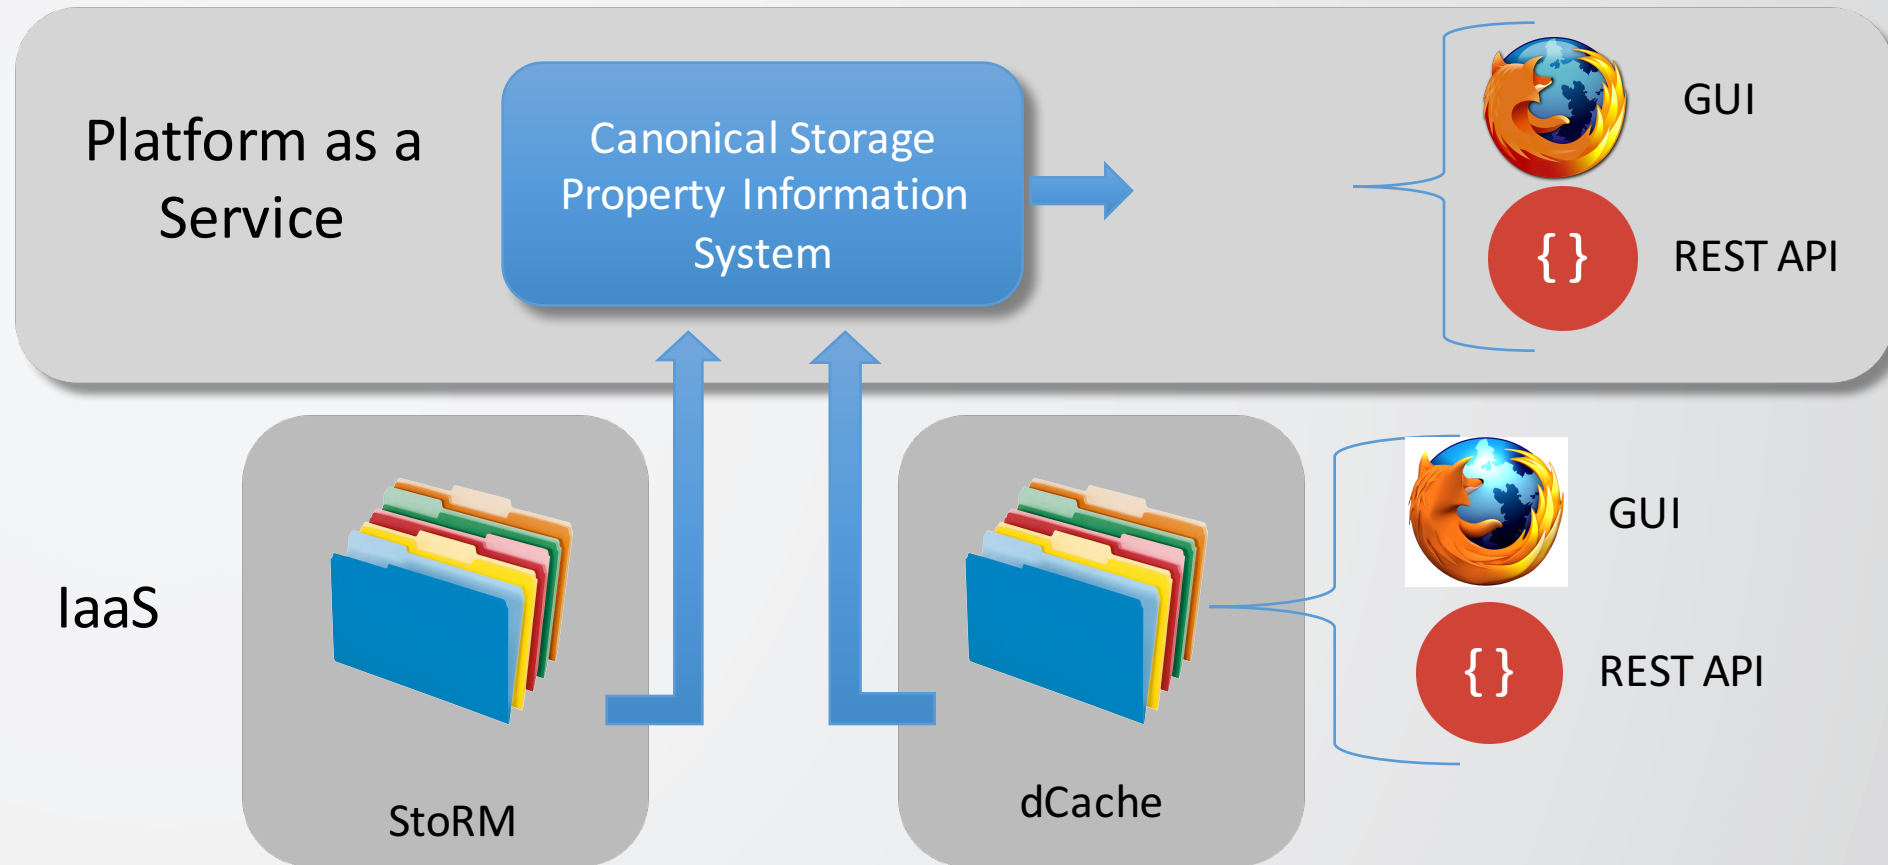
INDIGO-DataCloud is co-funded by the
Horizon 2020 Framework Programme

Where did it come from?



- Amazon
 - S3 : online
 - Glacier : nearline
- Google
 - Standard
 - Durable Reduces Availability (DRA)
 - Nearline
- IBM (HPSS, GPFS)
 - Storage classes (user defined)
- dCache
 - Tape
 - Disk (spinning or SSD)
 - Resilient Management ('n' copies)

First Ideas



The INDIGO – DataCloud Approach



Built a common (agreed) vocabulary for Data Management

Map agreed vocabulary to a QoS protocol spec

Provide a reference Implementation

CDMI Cloud Data Management Interface



- SNIA Cloud Data Management Interface (CDMI)
 - ISO/IEC standard
 - **Interoperability for Data stored in the cloud**
 - Cloud solution vendors
- More than 20 products that meet the CDMI

Storage Networking Industry Association



- **Non-Profit Organization**
 - Industrial and Scientific members from information technology
- **Mission**
 - Promoting vendor-neutral architectures
 - Standards and Educational Services
- **Facilitate**
 - Efficient Management of Data
 - Movement of Data
 - Security of Data

CDMI v1.1.1



- Discover capabilities available with storage provider
- Manage containers and the data that is placed in them
- RESTful principles in the interface design
- Providers can support a subset of CDMI
 - Must expose the limitations in the capabilities reported

CDMI v1.1.1



- Object Model

- Data Objects ↔ Files

- Store Values

- Containers ↔ Directories

- Capability Object

- Functionality provided by the storage system

- Each data object or container has one or more capability objects

- `cdmi_retention_period`

- QoS Class

- Collection of capability objects

CDMI v1.1.1



- Some Quality of Storage Attributes available
 - Data Redundancy **cdmi_data_redundancy**
 - Geo-Location **cdmi_geographic_placement**
 - Latency **cdmi_latency**
 - Retention **cdmi_retention_period**
- Data System can express an actual value to the exposed attributes
 - Geo-Location **DE**
 - Latency **3000**
- User can request a specific value during data creation or update
 - Geo-Location **FR, DE, JP**

CDMI Considerations

- CDMI is an industry standard.
 - Allows dCache to express its QoS for other systems
- Not widely Used
- Doesn't cover our use cases
 - Transition from one QoS to other
 - No clear way to determine if
 - the transition is allowed or
 - the system is capable of providing time
 - Lifetime of QoS
 - Pinning in dCache
 - Lock QoS on a data object

CDMI Considerations



- Pros

- CDMI is an industry standard.
- Allows dCache to express its QoS for other systems

- Cons

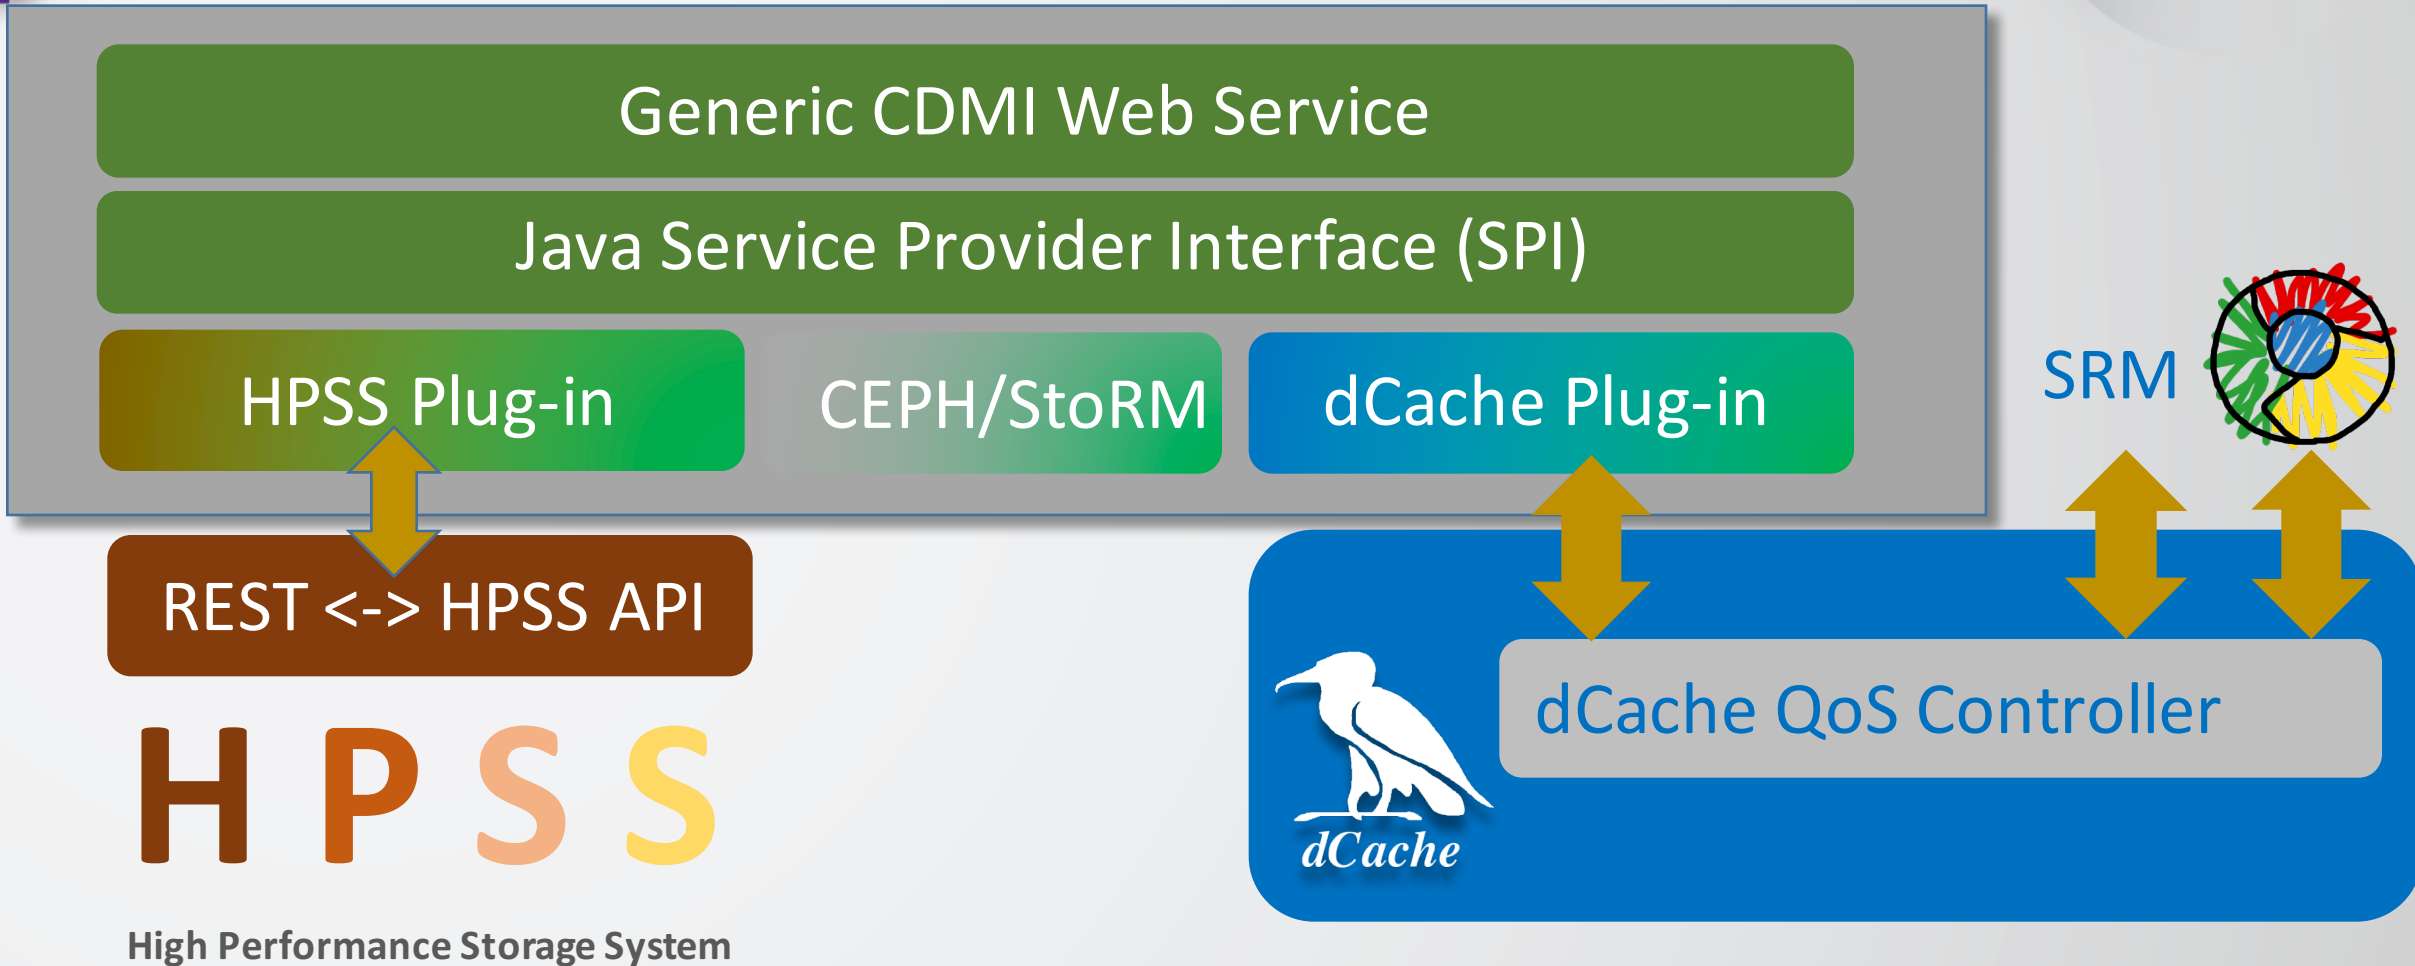
- Not widely Used
- Transition from one QoS to other
 - No clear way to determine if
 - the transition is allowed or
 - the system is capable of providing time
- Can't express Lifetime of QoS
 - Pinning in dCache
 - Lock QoS on a data object
- No way to manipulate the set of capabilities

INDIGO CDMI Extension




- Default attributes based on storage system configuration
- Assign QoS Class on creation and update of objects
- Changes to QoS restricted
 - Discover allowed transitions from one QoS Class to another
 - To the ones permitted by the storage system
 - Changes to individual capability of a QoS class not allowed
- Attribute for Capability Lifetime
- Additional attribute for monitoring QoS transitions

INDIGO QoS Architecture



Current Deployment

- KIT (master server)
- KIT (GPFS, HPSS: Tape, Disk)
- CNAF (StoRM: disk for now)
- DESY (dCache: Tape, Disk)
- Poznan (CEPH: disk only)



Infrastructure View



Log-in










dCache.org

Root

Users



















patrick



	Private		4/5/2017, 6:01:12 AM	--
	public-file		4/5/2017, 6:06:00 AM	  177 Bytes
	private-file		4/5/2017, 6:06:01 AM	 148 Bytes

Federated View

Available Qualities of Storage

Name	Access Latency [ms]	Number of Copies	Storage Lifetime	Location	Storage type	Available Transitions
 DataobjectProfile1	3000	1	20 years	PL	 Archival	
 DataobjectProfile2	2000	2	20 years	PL, UK	 Archival	
 DataobjectProfile3	500	3	20 years	NL, ES, PL	 Archival	
 disk	100	1		DE	 Process	
 disk+tape	100	2		DE	 Process	
 DiskAndTape	50	3	20 years	DE	 Process	
 DiskAndTape	50	2		IT	 Process	
 DiskOnly	50	3	20 years	DE	 Process	
 DiskOnly	50	1		IT	 Process	

Collected Capabilities

- Access Latency
- Number of Copies
- Storage Lifetime
- Location
- Available Transition

Back to Infrastructure View

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">• DE
Data redundancy	3



The End

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

Better Software for Better Science.





The End


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

Better Software for Better Science.

Conclusion



- 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 😊)

- 
- A decorative horizontal line composed of purple pixels, starting with a larger, more complex pattern on the left and becoming a simple dotted line towards the right.
- Brokering IG and Brokering WG
 - Vocabulary groups
 - NEW : data preservation :
 - Overlap with data management plans.