

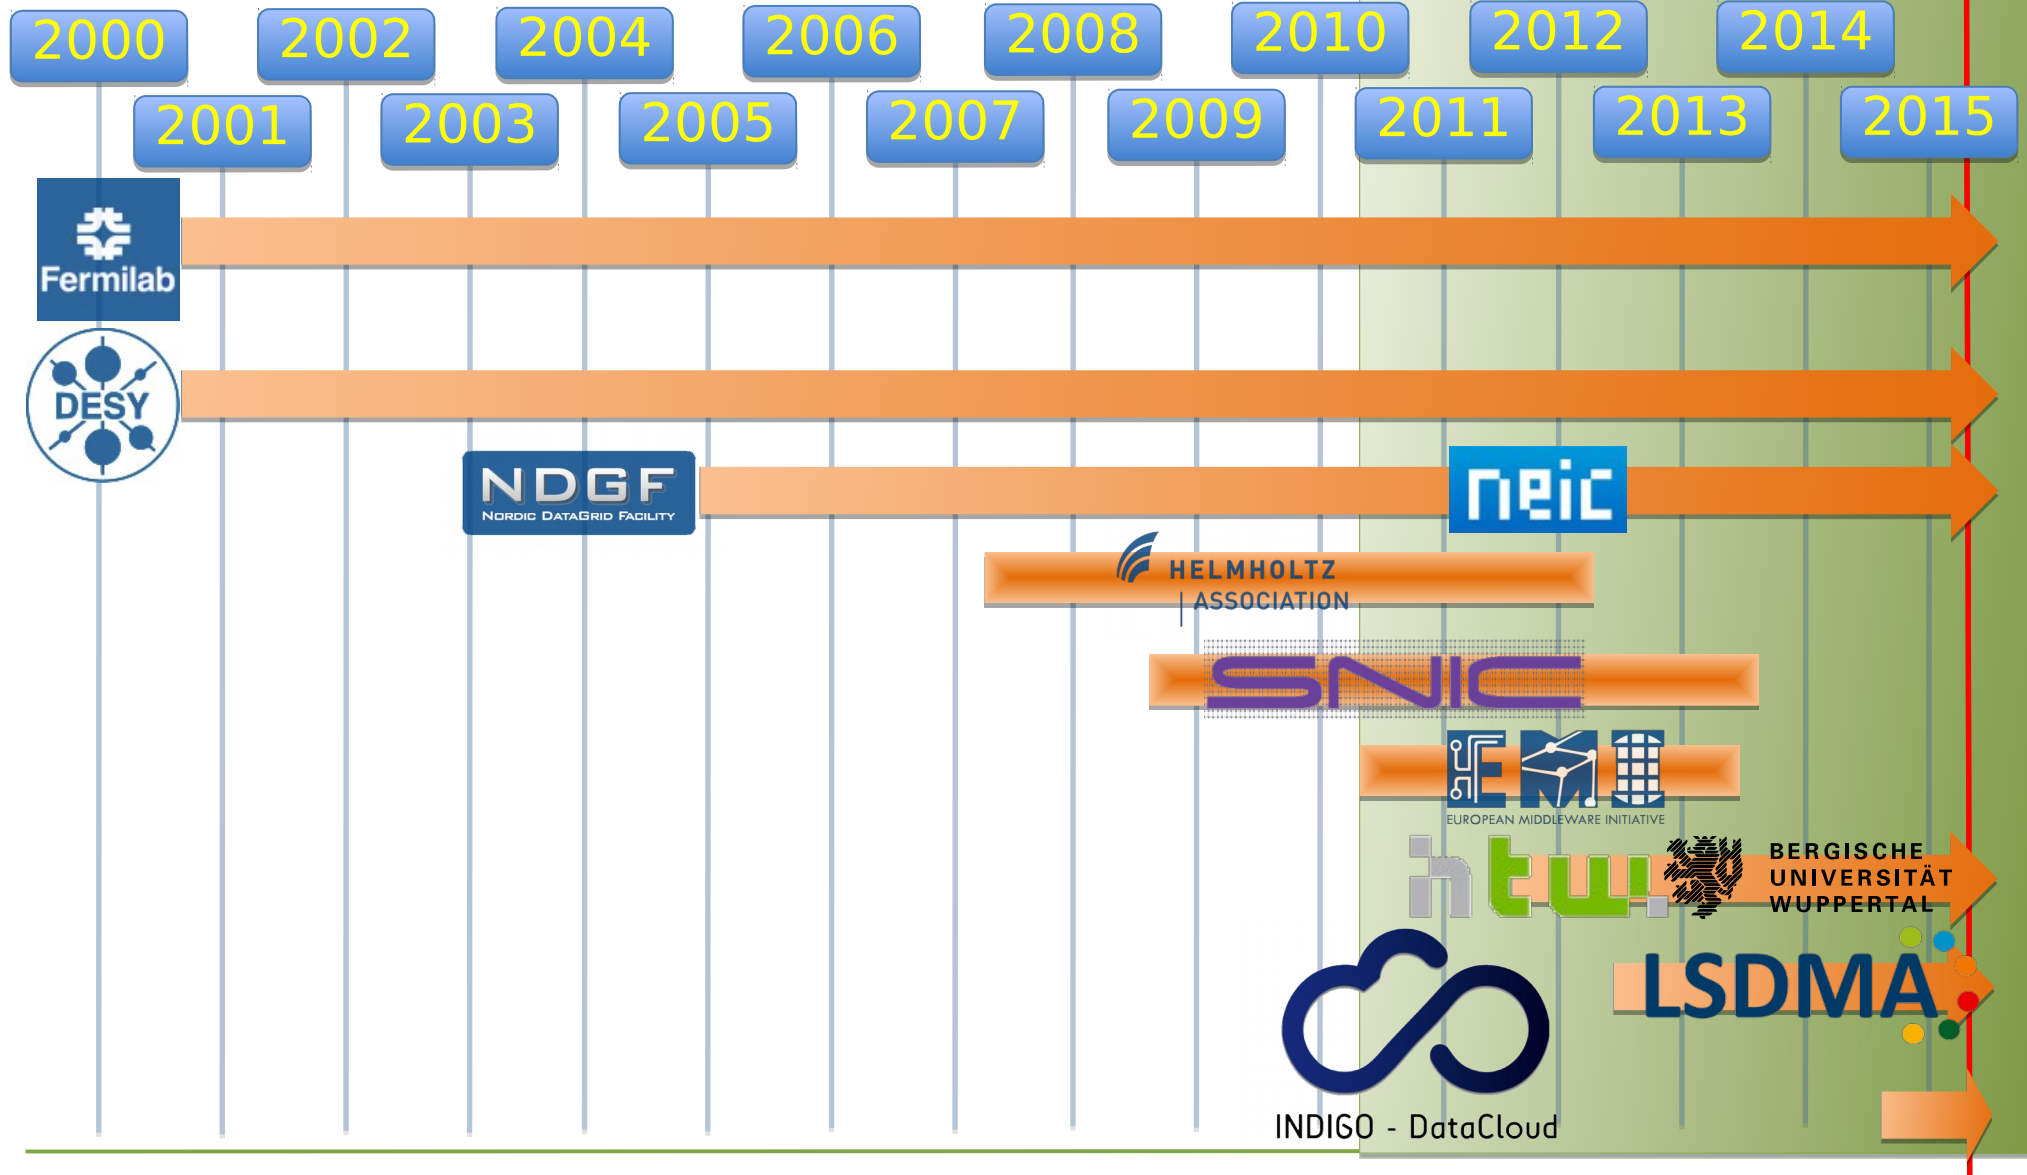
dCache developments

Paul Millar

Physics at the Terascale (2015-11-18)



Association and funding



INDIGO-DataCloud: cheat sheet



INDIGO - DataCloud

- A Horizon-2020 project
 - **Approved:** January 2015; **Started:** April 2015; **Ends:** September 2017
 - 26 partners (inc. DESY and KIT) from 11 European countries,
 - **Budget:** over €11 million
 - **Objective:** develop an Open-Source platform for computing and data, deployable on public and private cloud infrastructures.
 - Requirements from 11 INDIGO communities.
-


What is INDIGO-DataCloud



INDIGO - DataCloud

- **Biological and medical science,**

Biological, molecular and medical imaging, life science research applied to medicine, agriculture, bio-industries and social, structural biology.



- **Social science, arts and humanities,**

Georeferencing (e.g., of current and historical maps), cultural heritage, smart sensors.



- **Environment and earth science,**

Biodiversity and ecosystem research, interactions between geosphere, biosphere and hydrosphere, earth system modelling.



- **Physical science,**



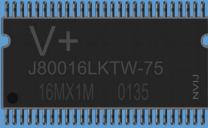
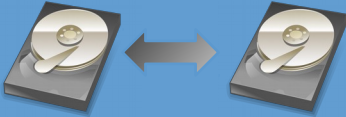
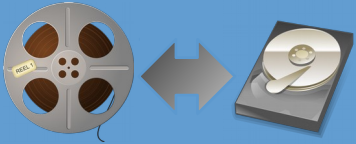
Astrophysics, theoretical and experimental research in physics.




The dCache core team

- The **permanent** effort:
 - 1 FTEs at NeIC,
 - 3 FTEs at DESY,
 - 1.5 FTEs at Fermilab.
 - The **project money** effort:
 - 2 FTE LSDMA → 1 FTE in 2016,
 - 4 FTE INDIGO-DataCloud.
- Total:** 11.5 FTE (10.5 from 2016).
-

Future: improved media handling

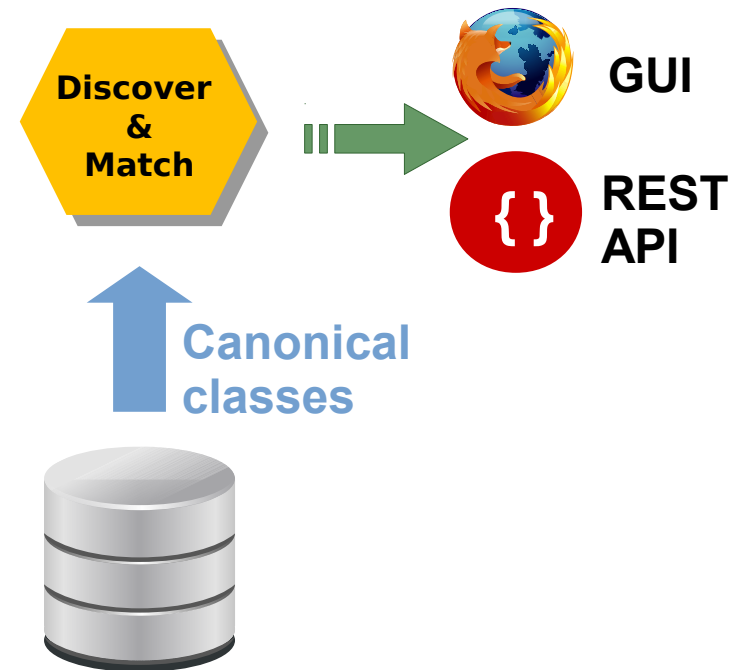
Media Quality					
Access Latency	HIGH	MEDIUM	LOW	MEDIUM	MEDIUM
Durability	OK	MEDIUM	Not so clear	Quite OK	OK
Data rate	OK	OK	MEDIUM	OK	OK
Cost	Very low	Reasonable	Very high	MEDIUM	MEDIUM

Making QoS choices meaningful

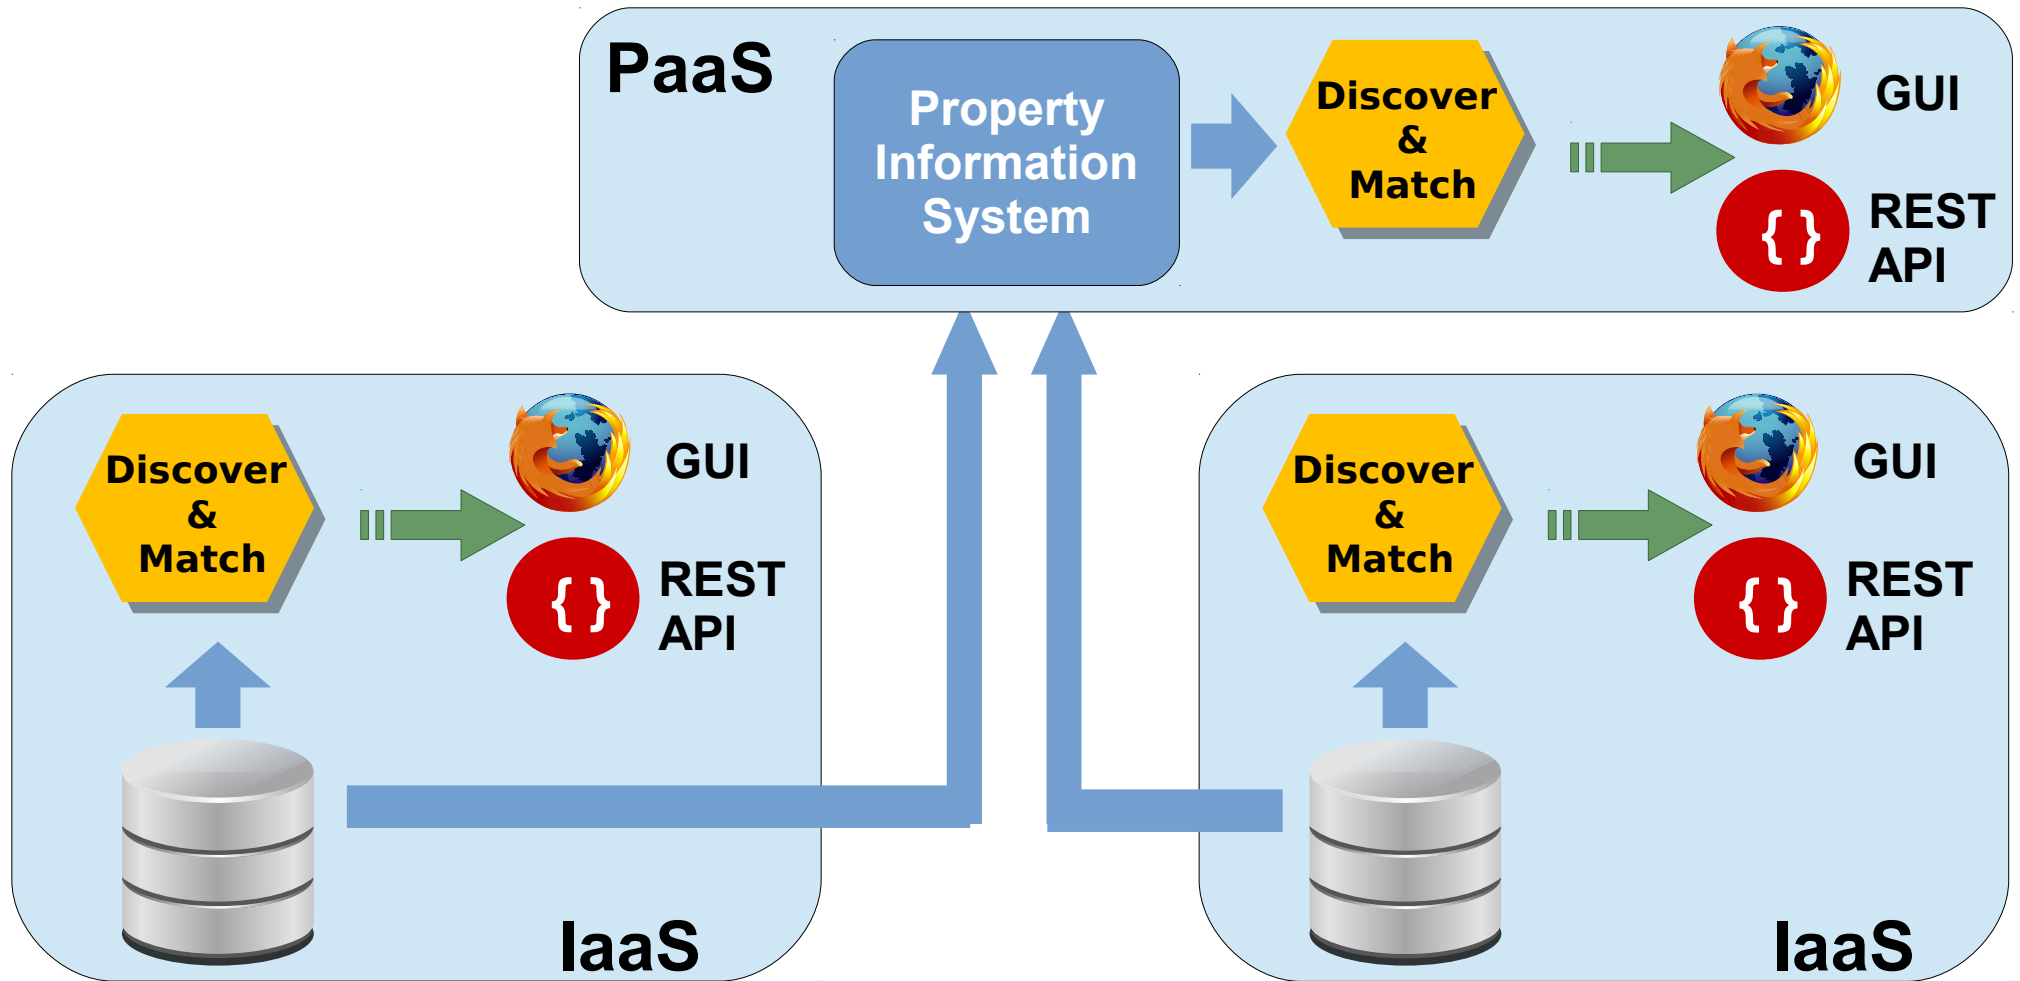


- Low latency & lowest price → Class #1
- High throughput & super durable → Class #2
- Large volume & cheap & archive → Class #3

VS




Federating QoS Choices



Improved media handling: Data-LC

Data Lifecycle is just time dependent changes of

- Storage Quality of Service
- Ownership and Access Control: PI Owned, limited access → Site Owned, Public access
- Payment model: pay-as-you-go → pay-in-advance for rest of lifetime
- Maybe other things



{Replica → Resilience} Manager

- Complete redesign
 - New features include:
 - Better handling of **temporary offline pools**,
 - Useful **diagnostic information**,
 - Supports **multiple replication strategy**:
 - per storage-class, configurable replication & partitioning,
 - Integrated **alarm support**,
 - **Reduced load** on dCache:
 - Information from PoolManager & namespace, not directly querying pools,
 - Fairness in choosing between **foreground or background replication**,
 - **Configurable policy** on internal replication failures.
-

dCache on Ceph

- The start of support for **cluster filesystems**: GPFS, Ceph, ... – Ceph is our initial focus.
 - Two approaches:
 - **Single pool** per cluster (easier, but less useful),
 - **Multiple pool** per cluster (allows load-balancing, harder to achieve)
 - Benefits to dCache:
 - **Data distribution**: delegated to underlying cluster storage,
 - **Integration**: (re-)use existing site storage infrastructure.
 - Benefits to cluster storage:
 - **dCache features**: protocols, authentication, tape integration, ...
 - **Future proof**: migrate from storage technology without down-time.
 - Plan to demo prototype at ISGC 2016 (March next year)
-

Summary

- dCache team is **strong** after recently expanding,
New project money means **more features**,
 - **Resilience manager** as replacement for replica manager,
 - **QoS and Data-Lifecycle** places more control in users hands,
 - dCache on **Ceph** coming soon.
-

Backup slides

