RIA-653549

INDIGO-DataCloud plans for storage

Giacinto Donvito – INFN-Bari

On behalf of INDIGO Collaboration

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

INDIGO: The consortium

- 26 partners from 11 countries
- The exagon contains the partners involved in development activities
- 11M€ of total founding by EC
- 4 Big private companies

From the **EC Evaluation** Summary Report:
"The consortium is of exceptional quality, and complementary, and with good relevant experience and skills"

Scientific Computational Portal “as a Service”

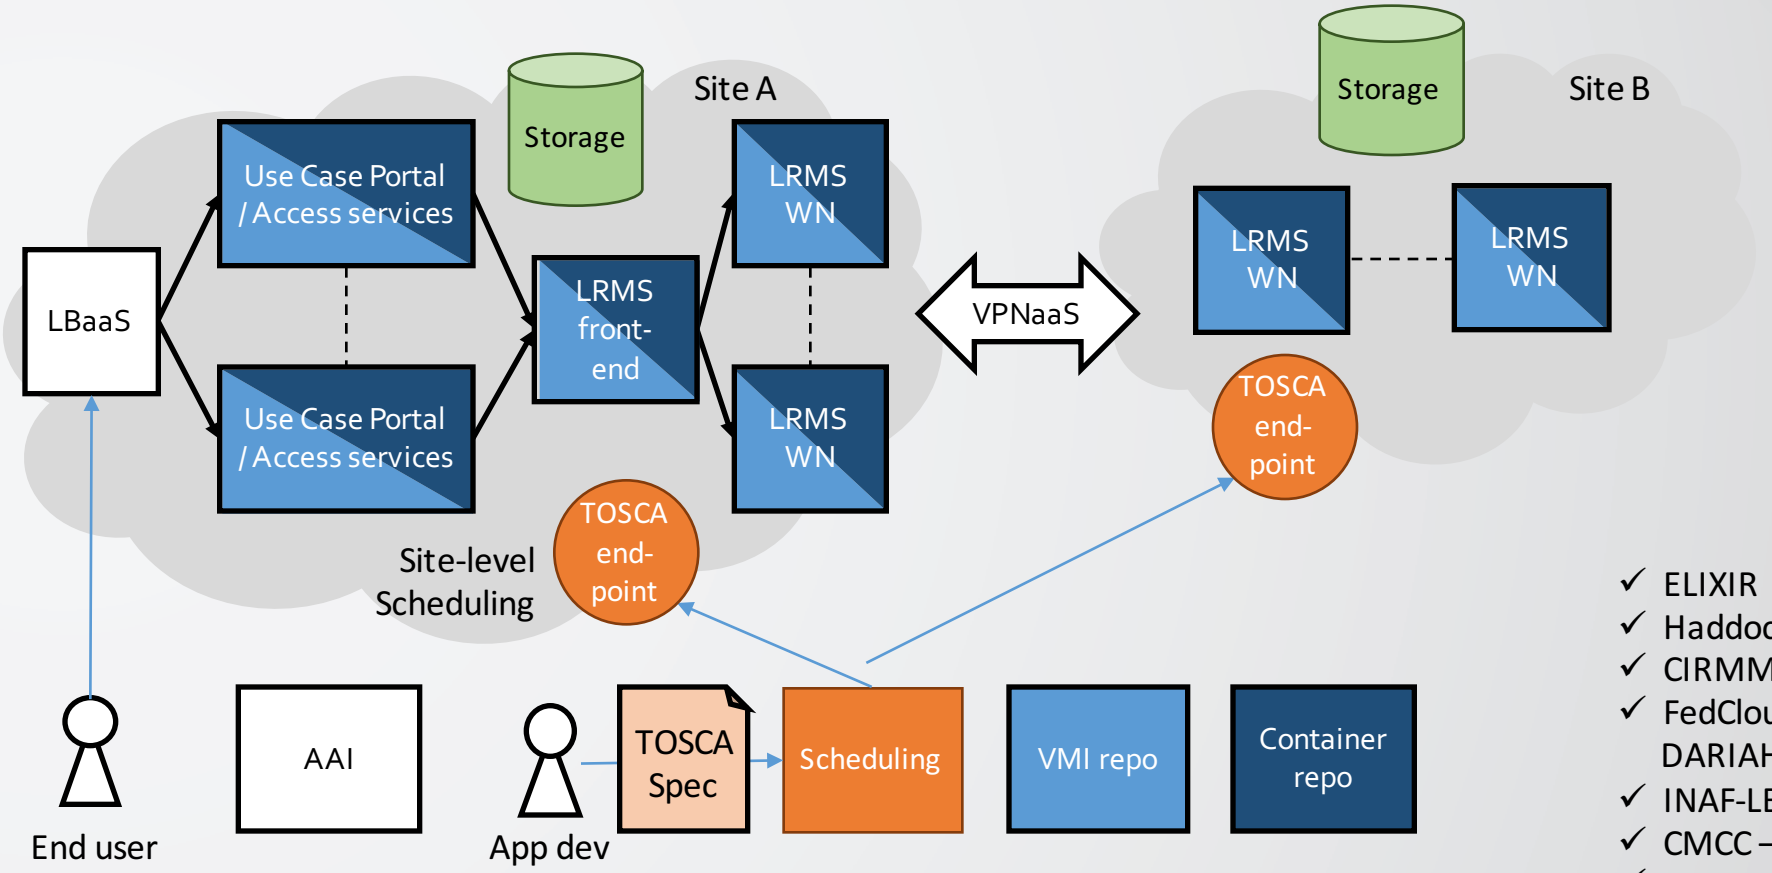

- ✓ ELIXIR
- ✓ Haddock
- ✓ CIRMMP
- ✓ FedCloud
  DARIAH
- ✓ INAF-LBT
- ✓ CMCC – ENES
- ✓ CTA
- ✓ ALGAE-BLOSSOM
  INGV - MOIST

Software as a Service Use Case

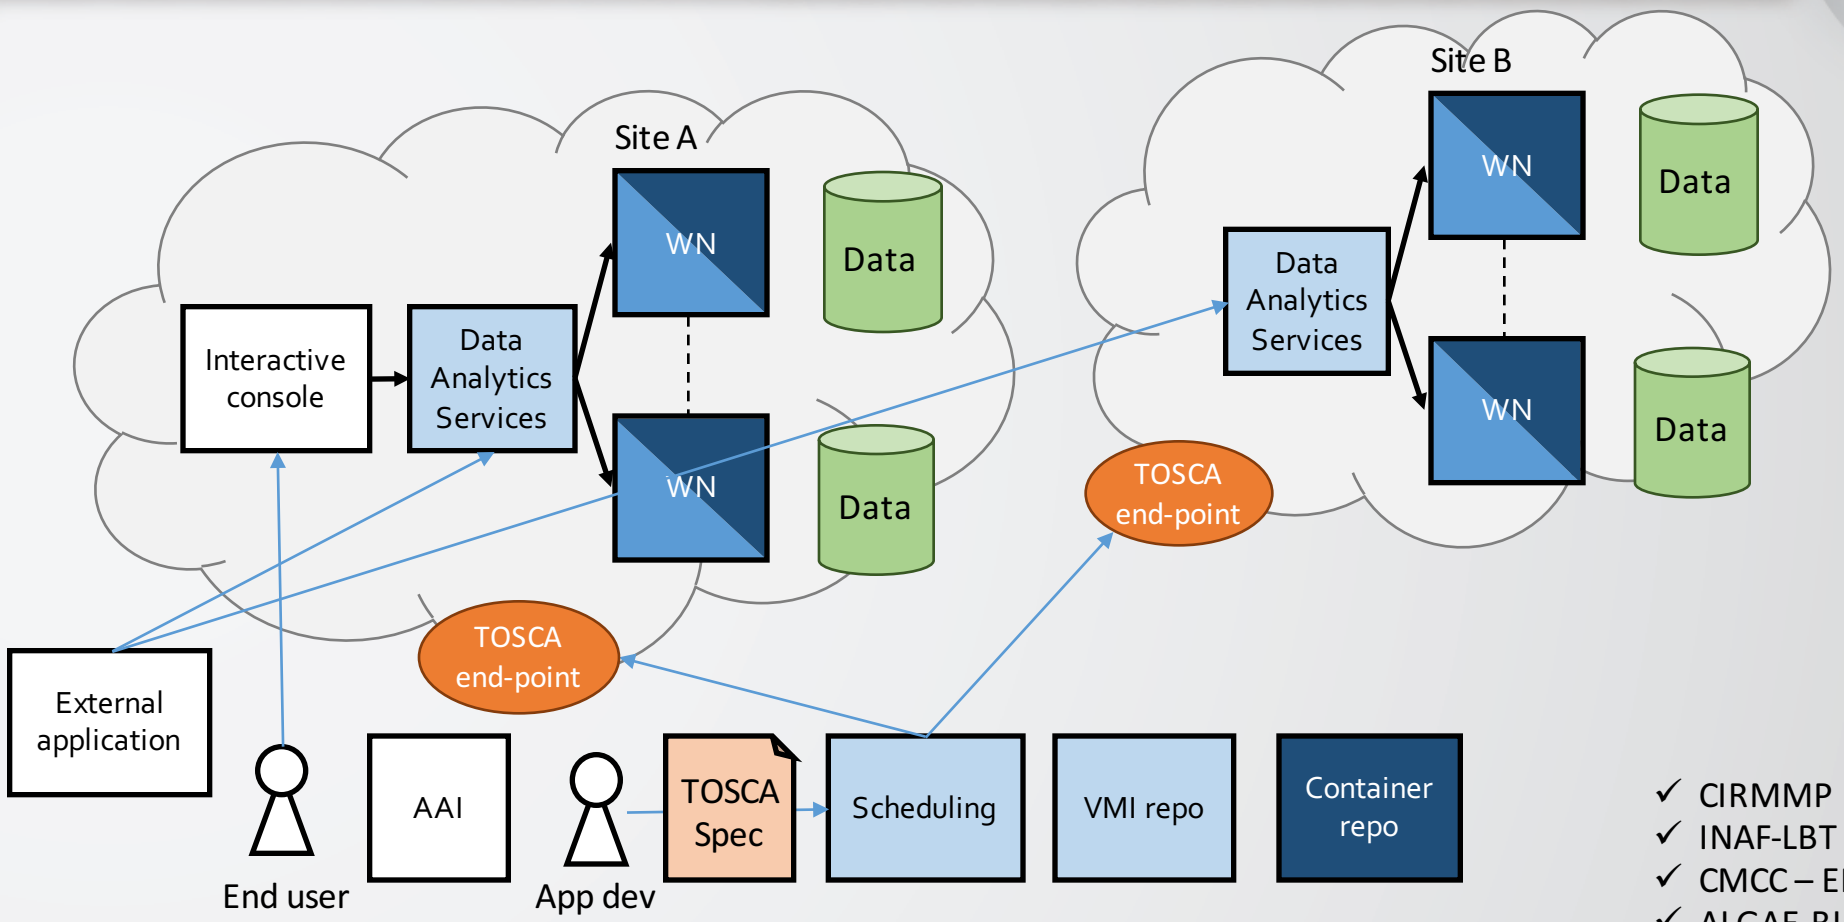

- ✓ CIRMMP
- ✓ INAF-LBT
- ✓ CMCC – ENES
- ✓ ALGAE-BLOSSOM
- ✓ INGV - MOIST

Example: R-Studio

- R-Studio Server provides a browser-based interface (RStudio IDE) to a version of R running on a remote Linux server.
  - How to deploy it using INDIGO?

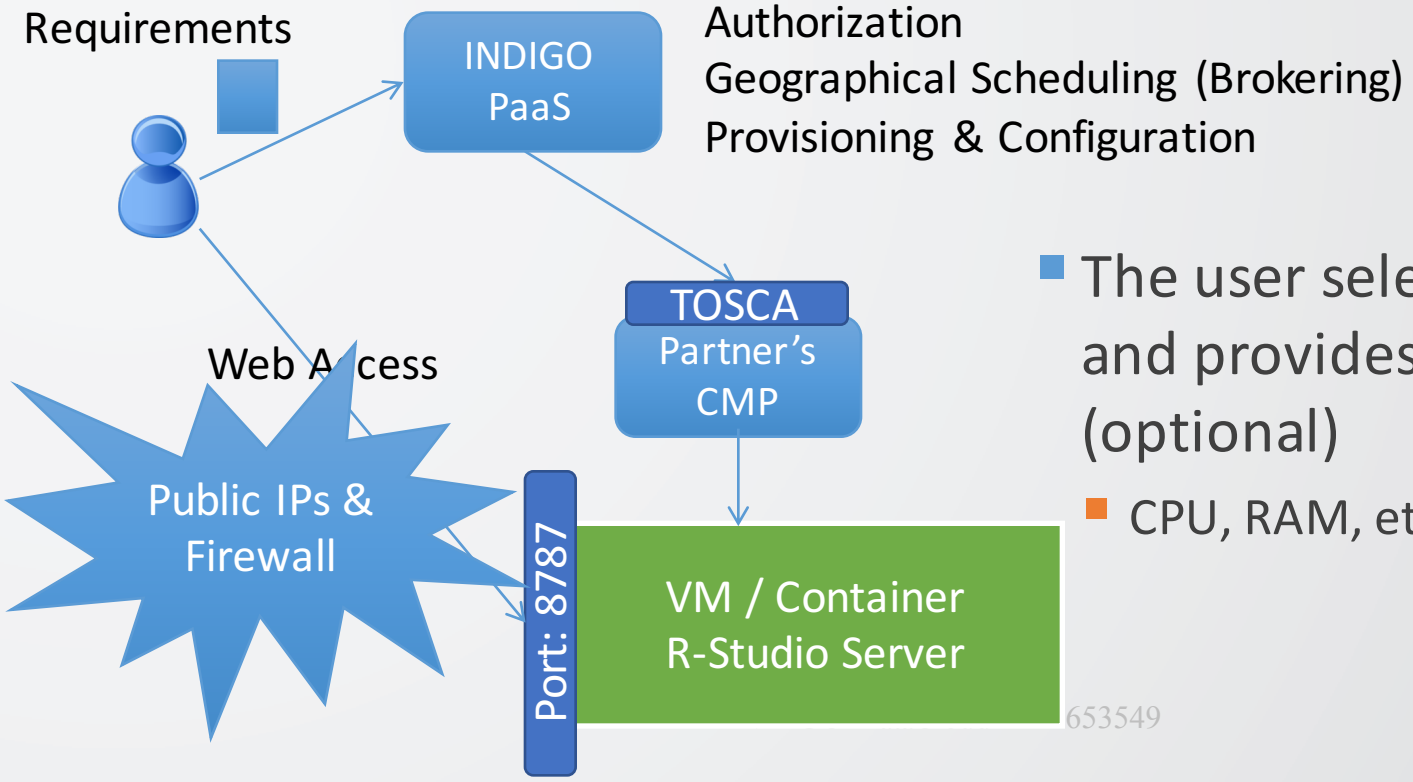

- The user selects the template for R-studio and provides further requirements (optional)
  - CPU, RAM, etc.

INDIGO overall architecture

INDIGO - DataCloud

Color codes:
Yellow: implementation based on already available solution to be improved/changed;

Orange:
Completely new services to be implemented

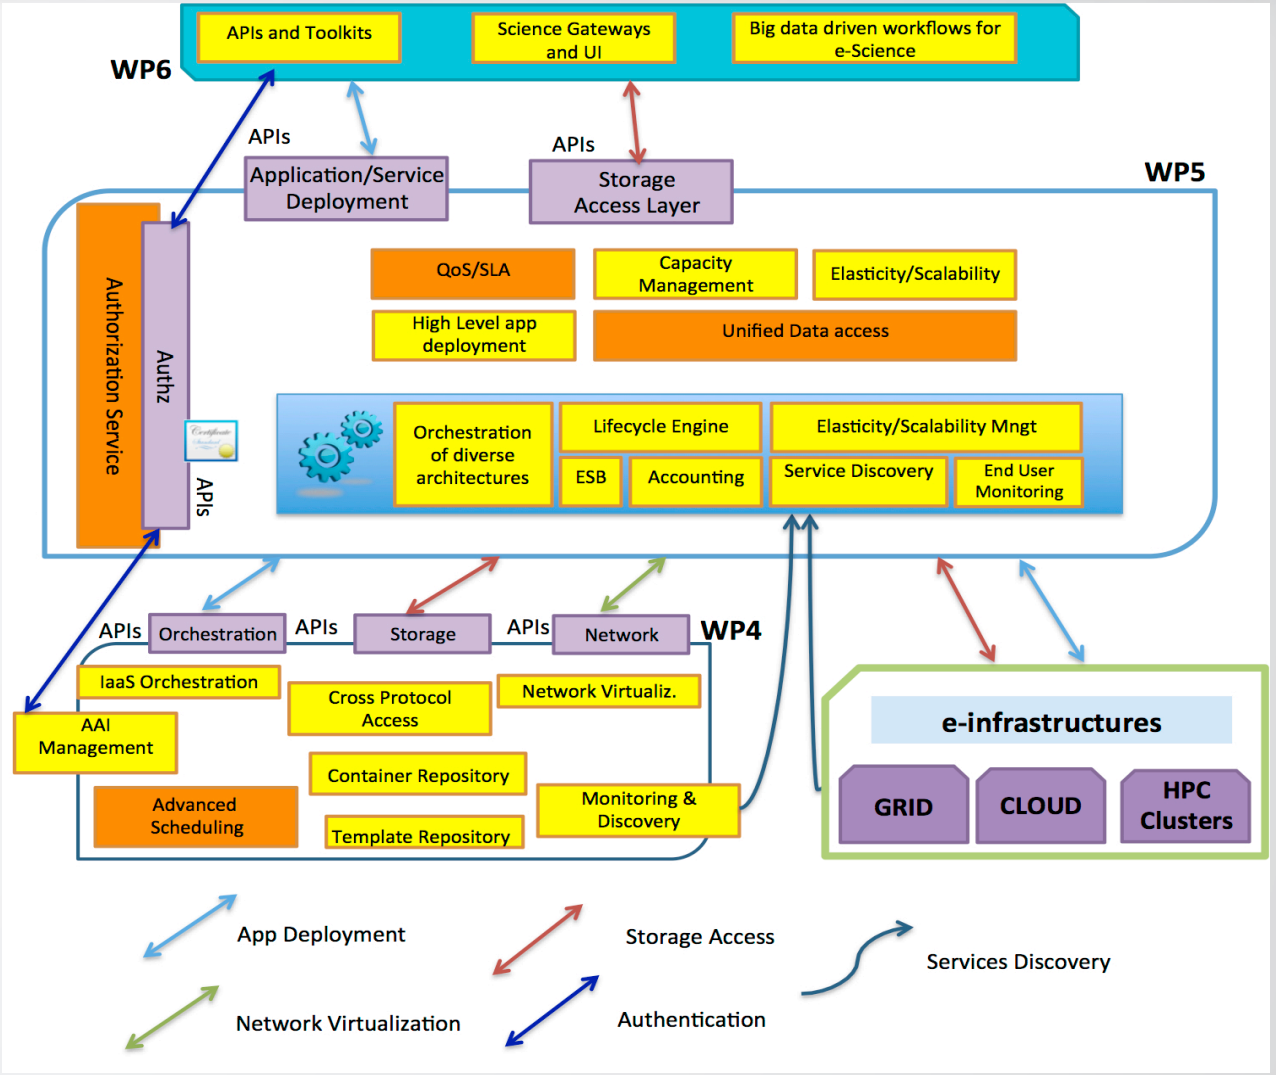

Simplified architecture

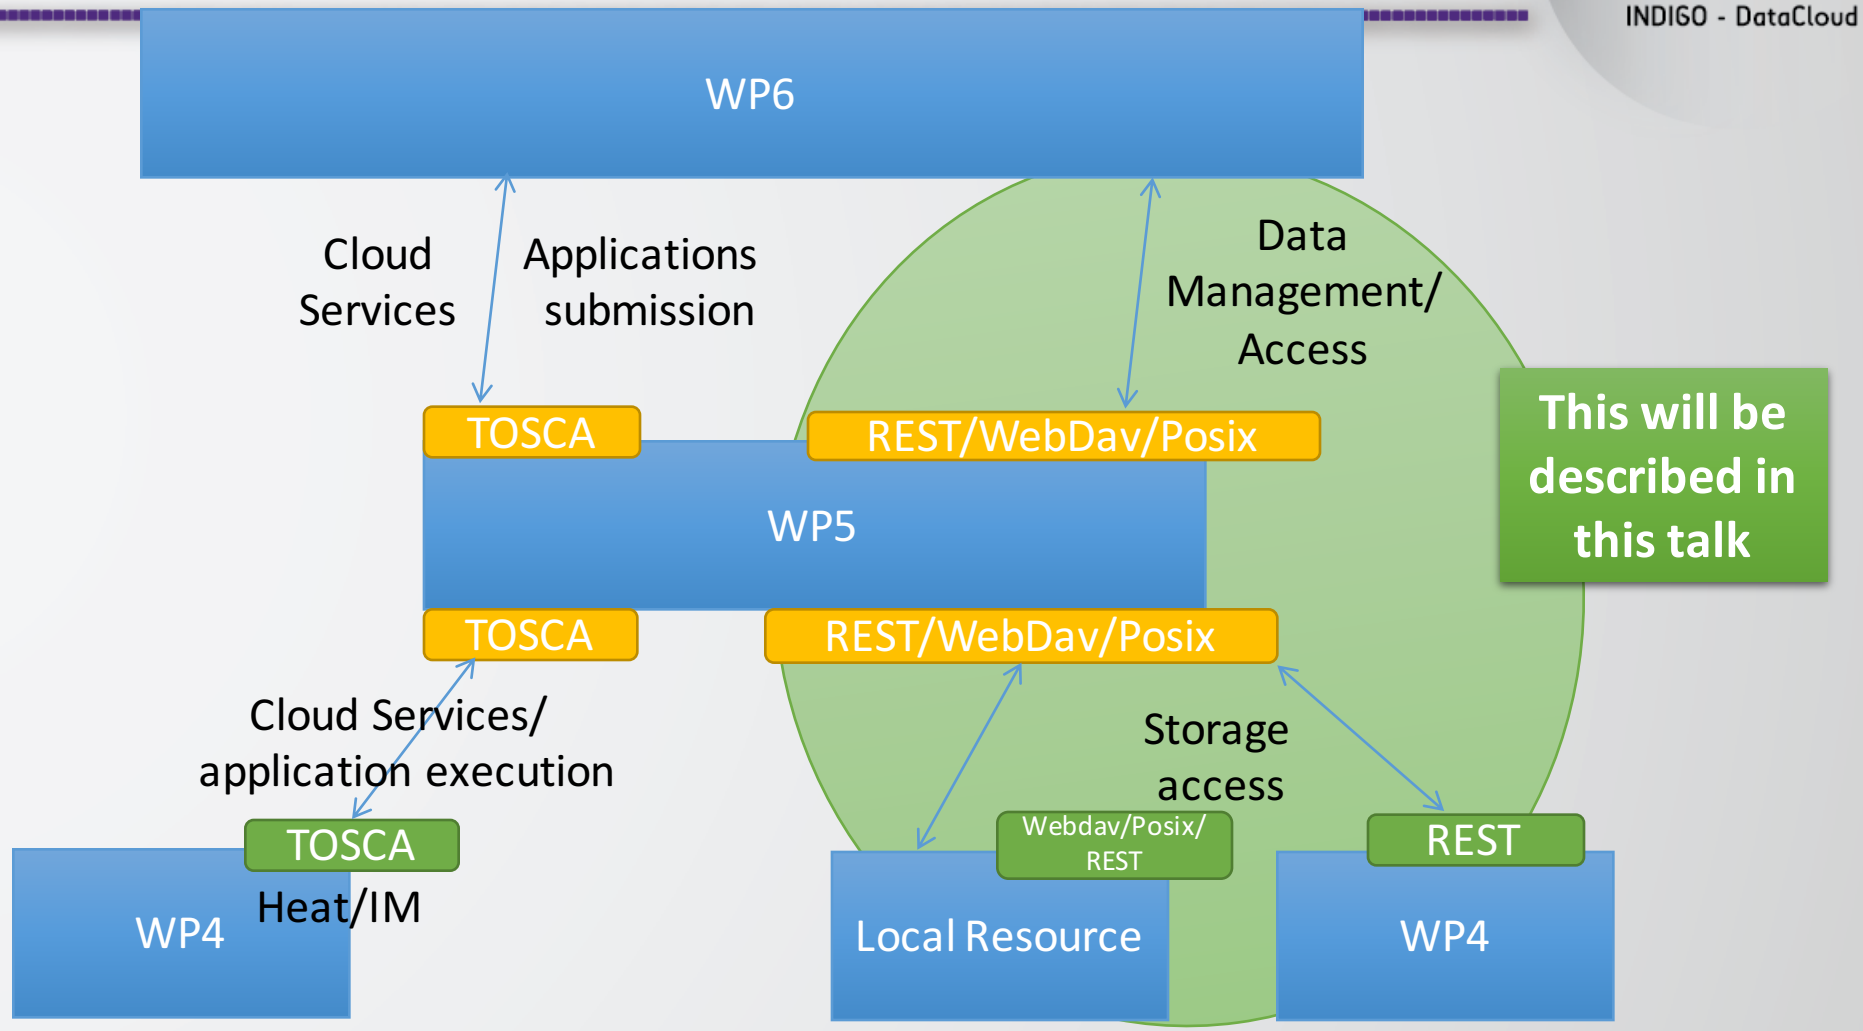

WP4: IaaS Platform Development

INDIGO - DataCloud

- Virtualized **Computing Resources**
  - Full Container support for Cloud Management Infrastructures and Batch
  - Container support for special hardware (Infiniband, GP-GPU's)
  - Spot Instances
  - Fair Share Scheduling
- Virtualized **Storage Resources**
  - QoS and Data Life Cycle for storage (storage management)
  - Access to data by meta data instead of name space
  - Dual access to data (Object Store versus POSIX file name space)
  - Identity Harmonization for storage
- Virtualized **Network Resources**
  - Orchestrating local and federated network resources
  - "Software Defined Network" evaluation Services and Appliances for for virtual networks

The QoS properties

- Is there a sufficiently complete set of properties ?
- In WCLG we only had two properties:
  - Access Latency
  - Retention policy
- That was already too much for most people
- Talking to Reagan Moore (IRODS) at the Paris RDA meeting:
  - He is suggesting about 200 properties
  - That might be a bit over the top for a start

QoS Property Value Ambiguity

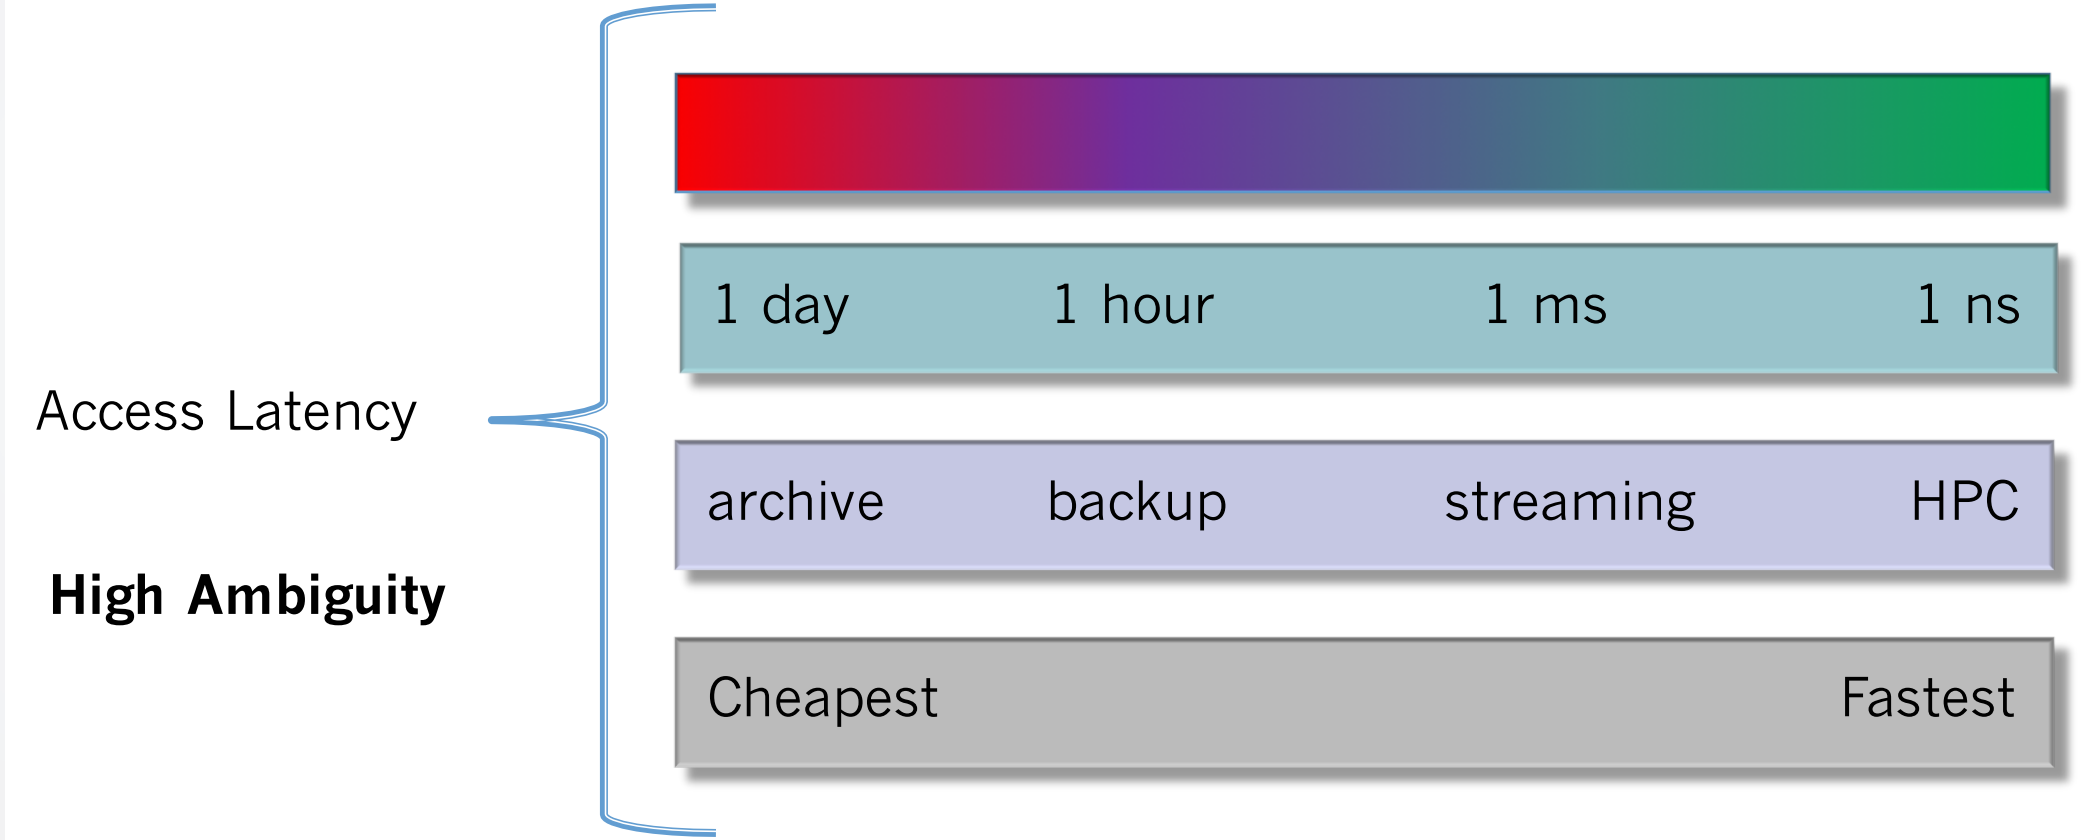

Property dependencies

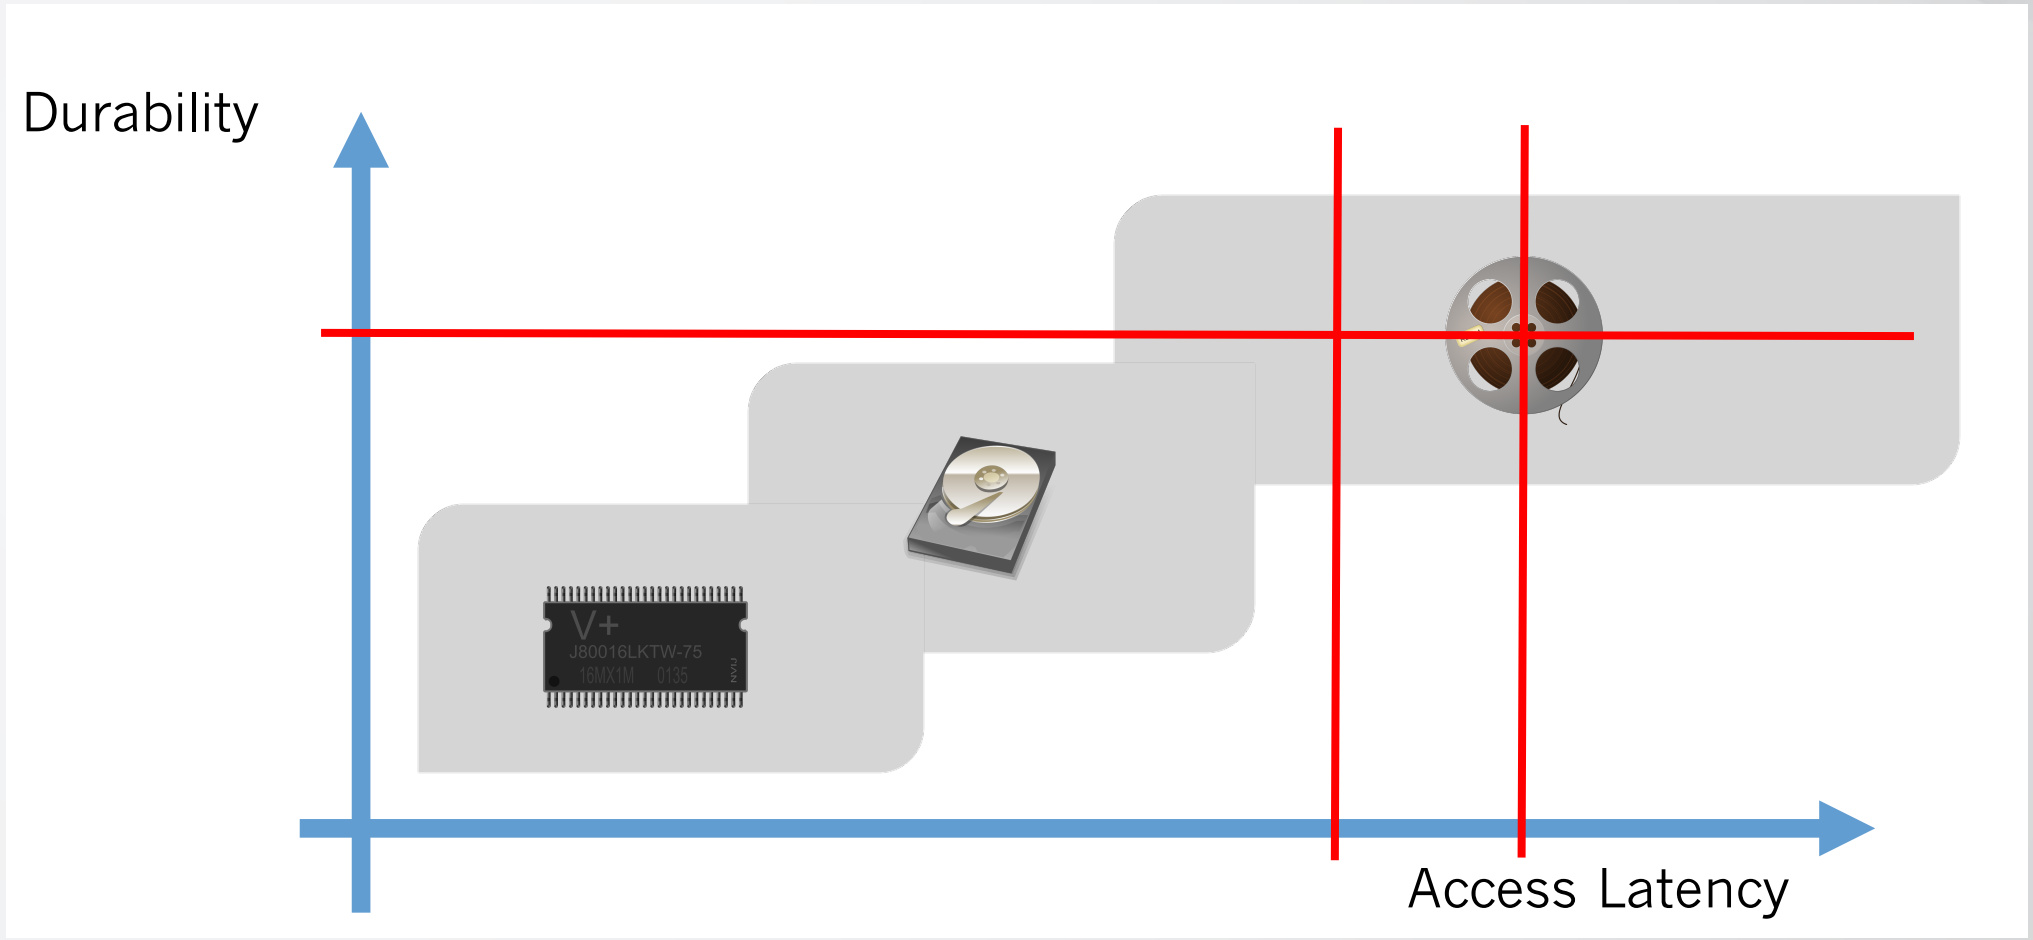

Property dependencies

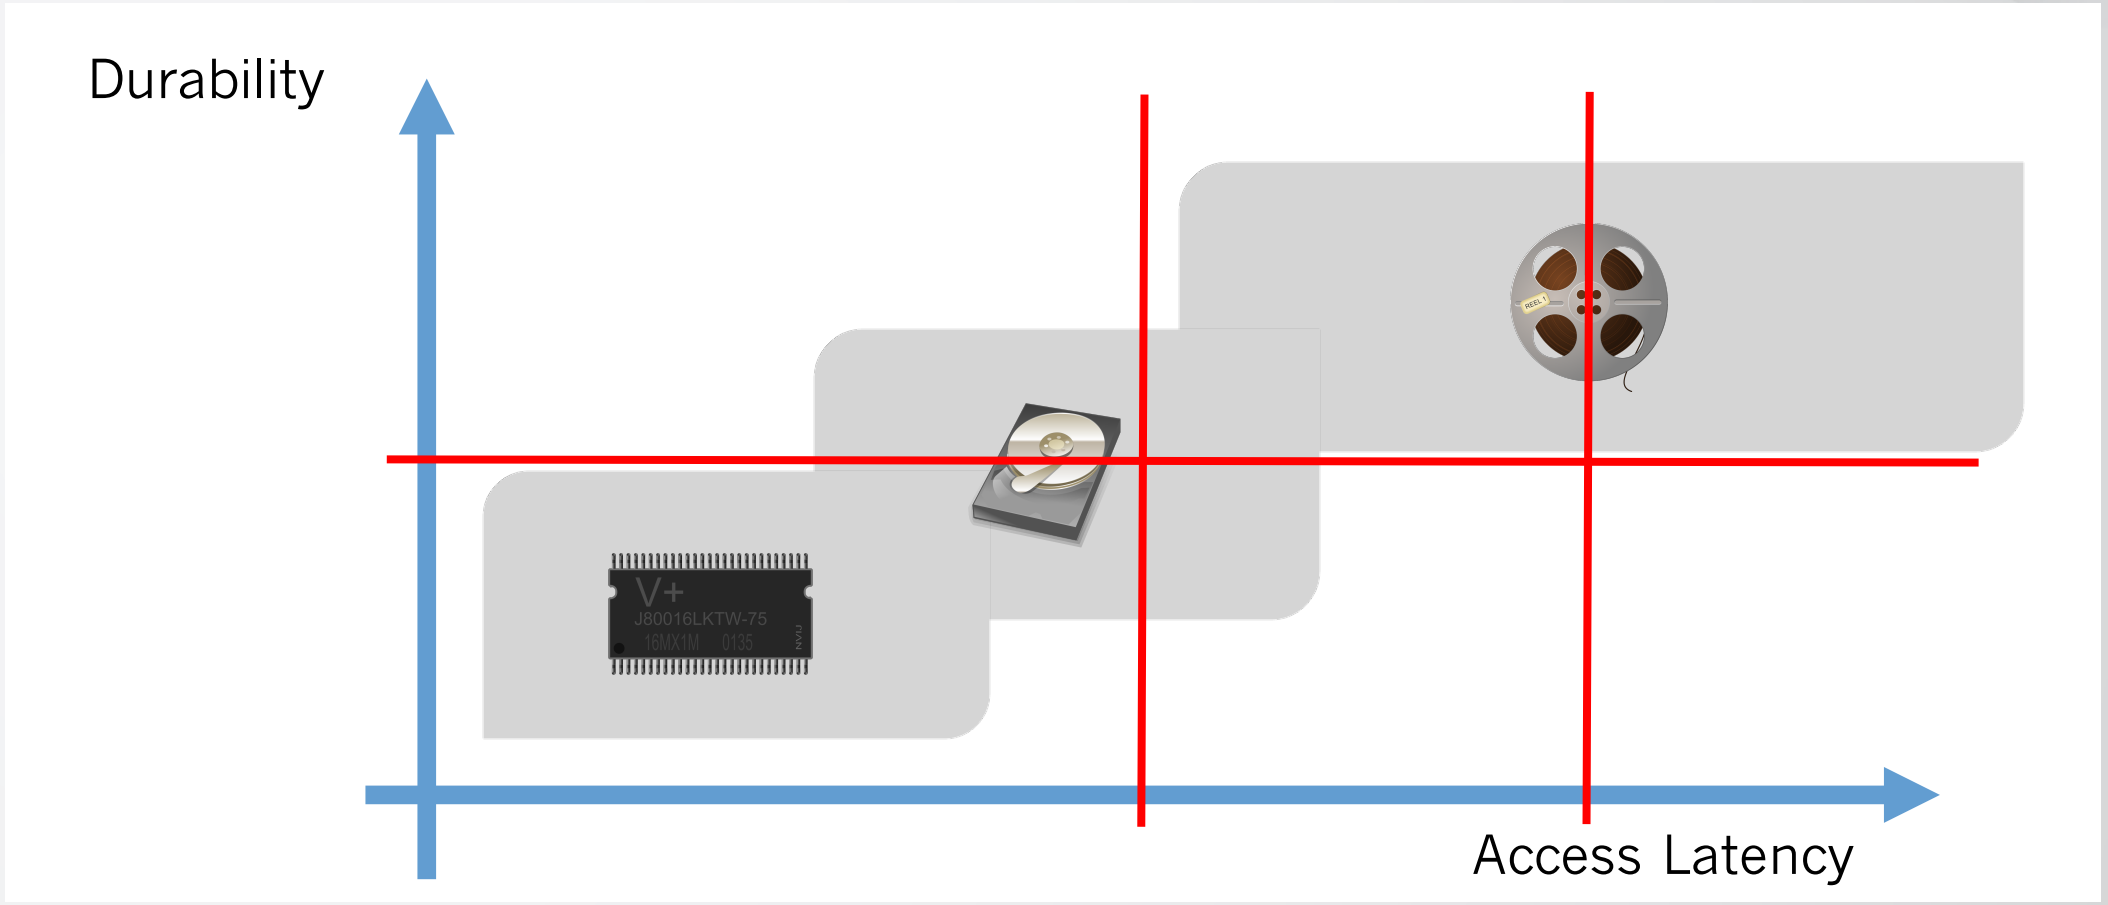

Property dependencies

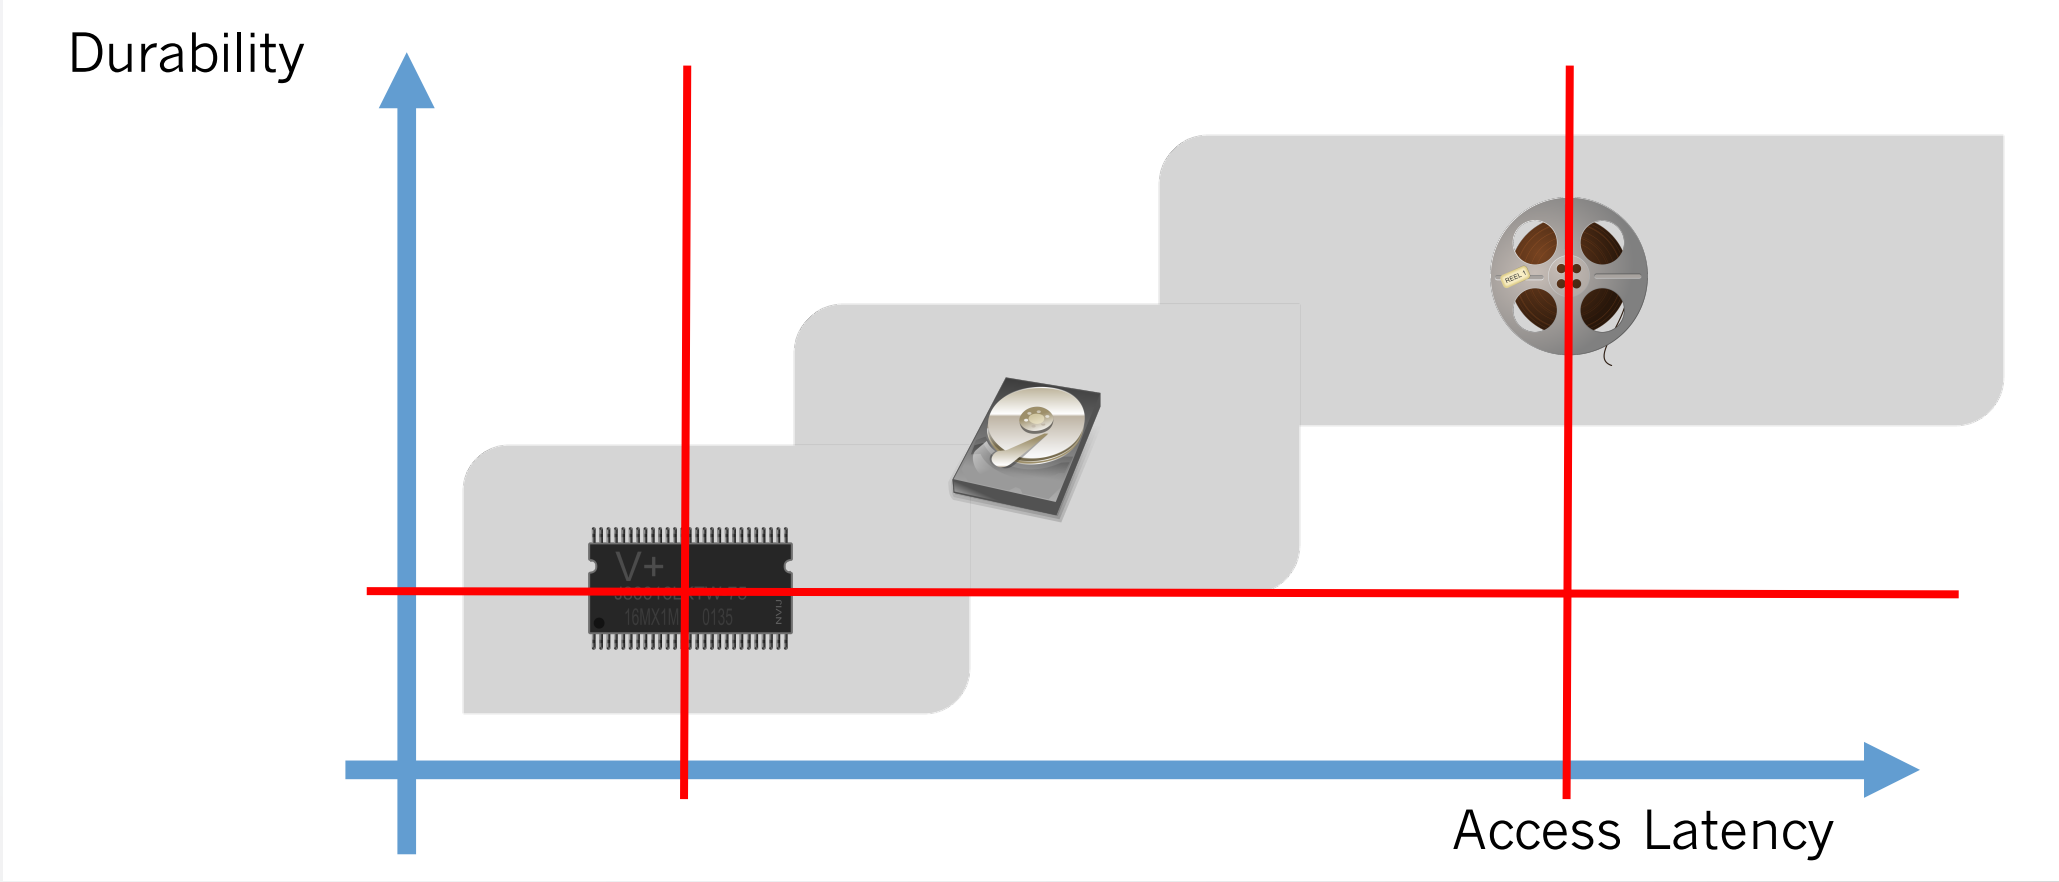

Property Quantization

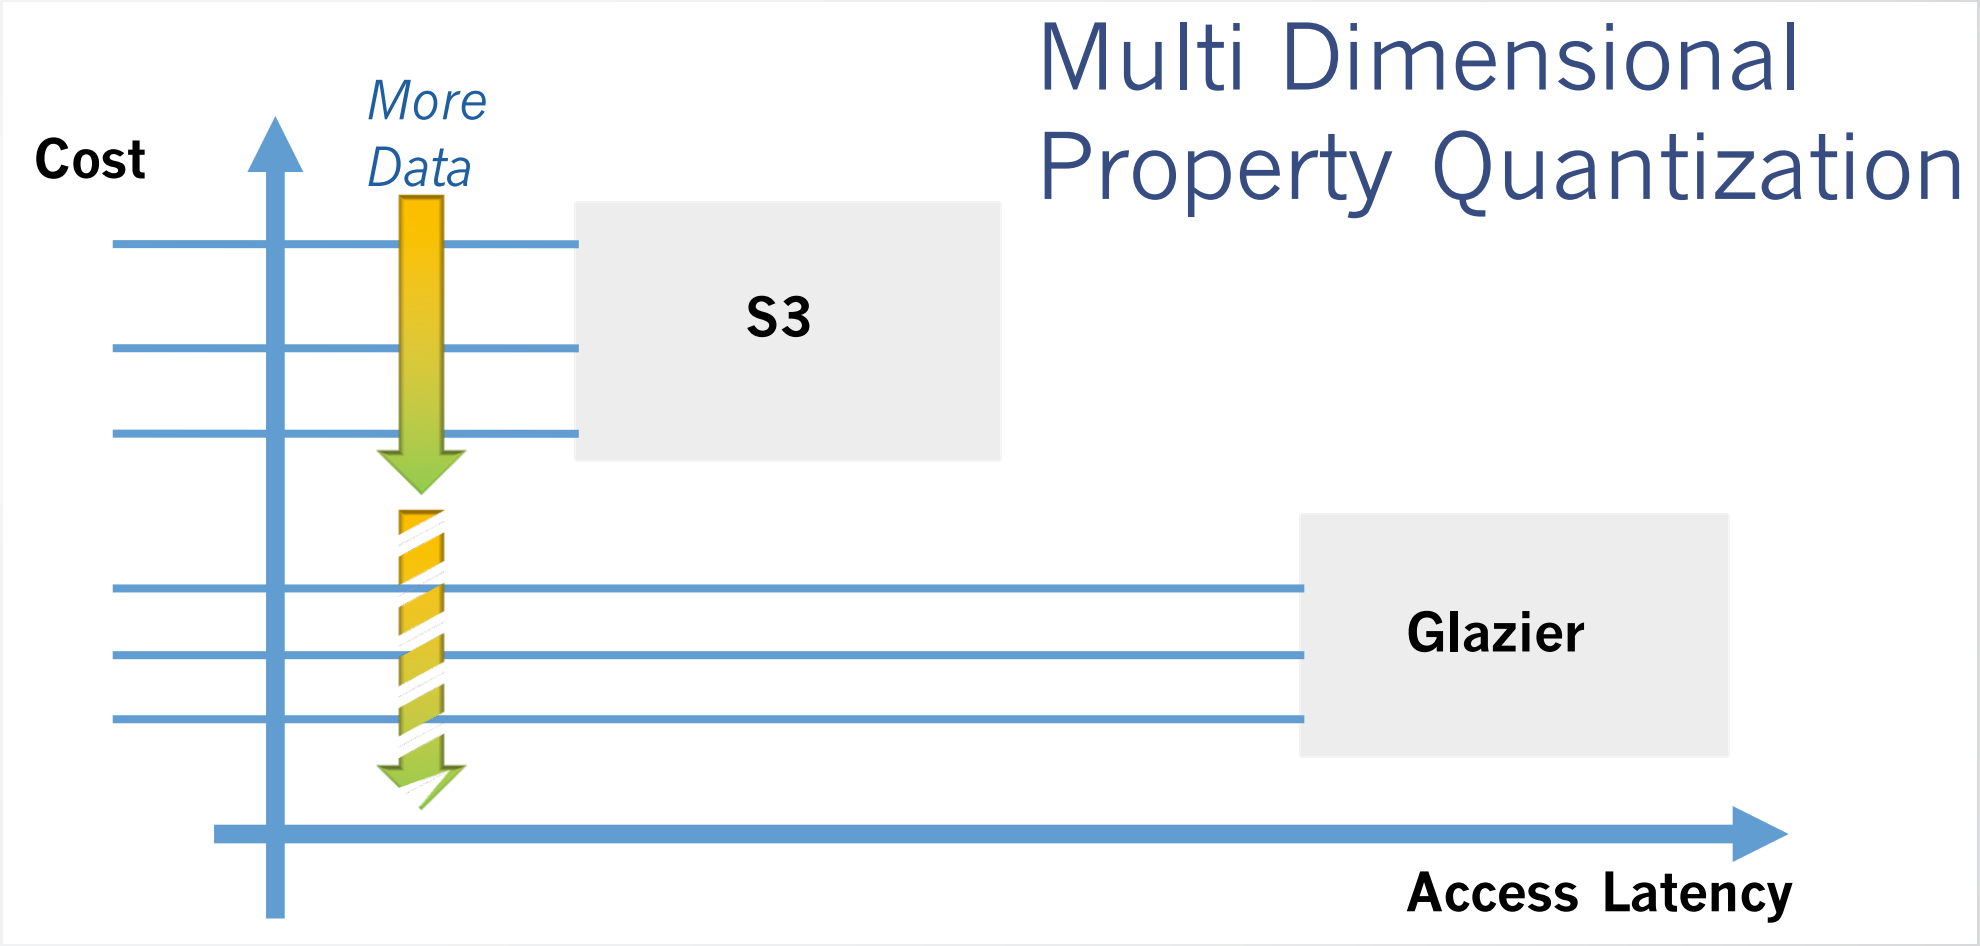

Properties zoo of existing systems

Amazon	S3		Glacier
Google	Standard	Durable Reduces Availability	Nearline
HPSS/GPSS	Corresponds	to the HPSS Classes	(customizable)
dCache	Resilient	disk+tape	TAPE

What are canonical properties ?

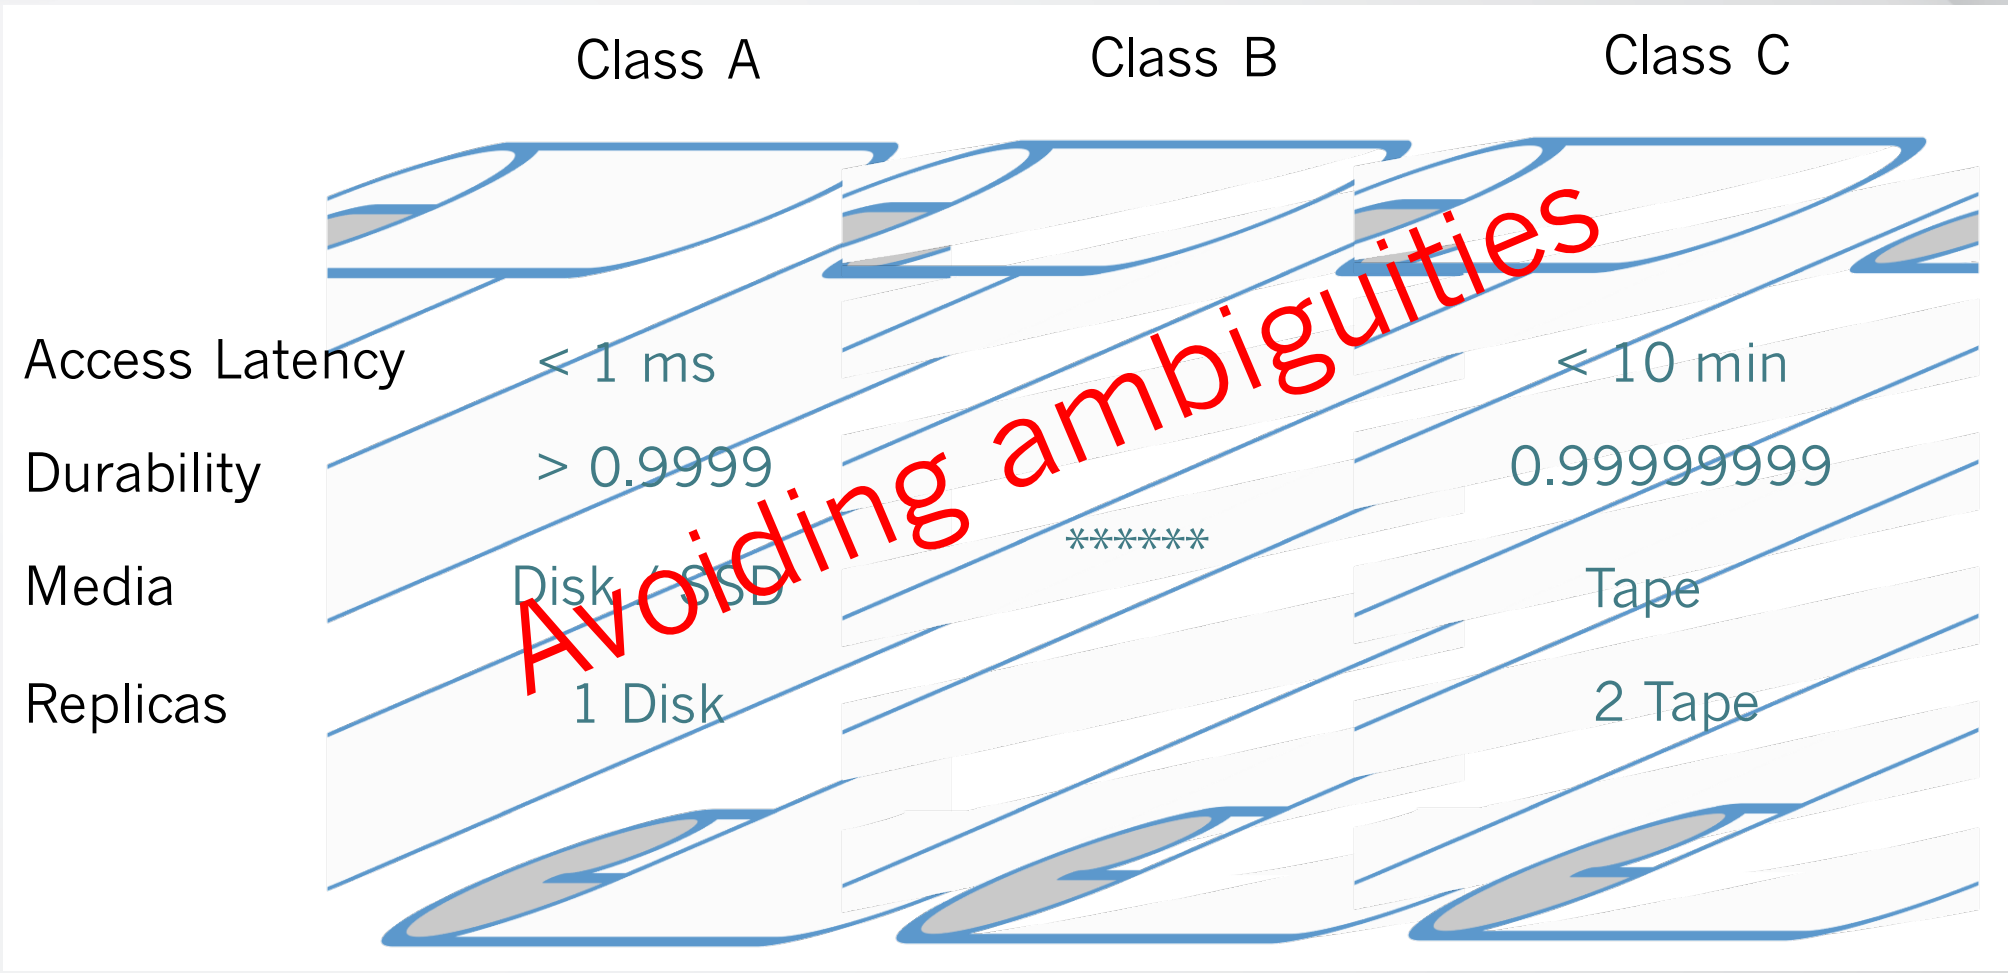

Canonical Storage Properties

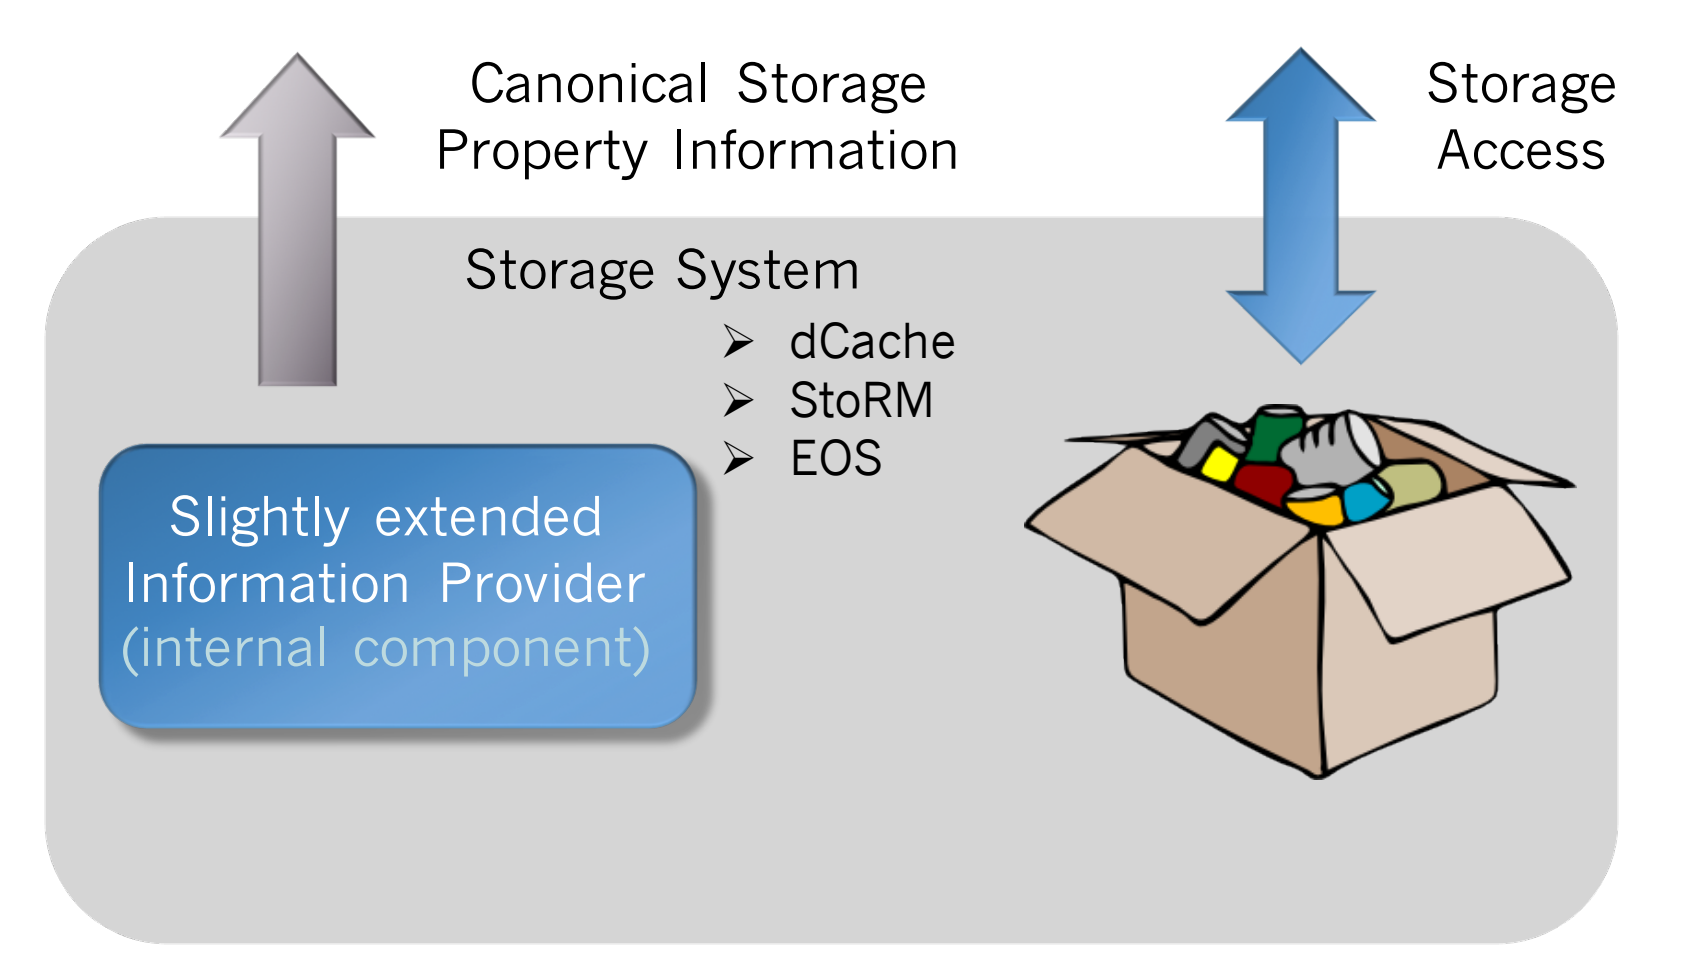

Canonical Storage Properties

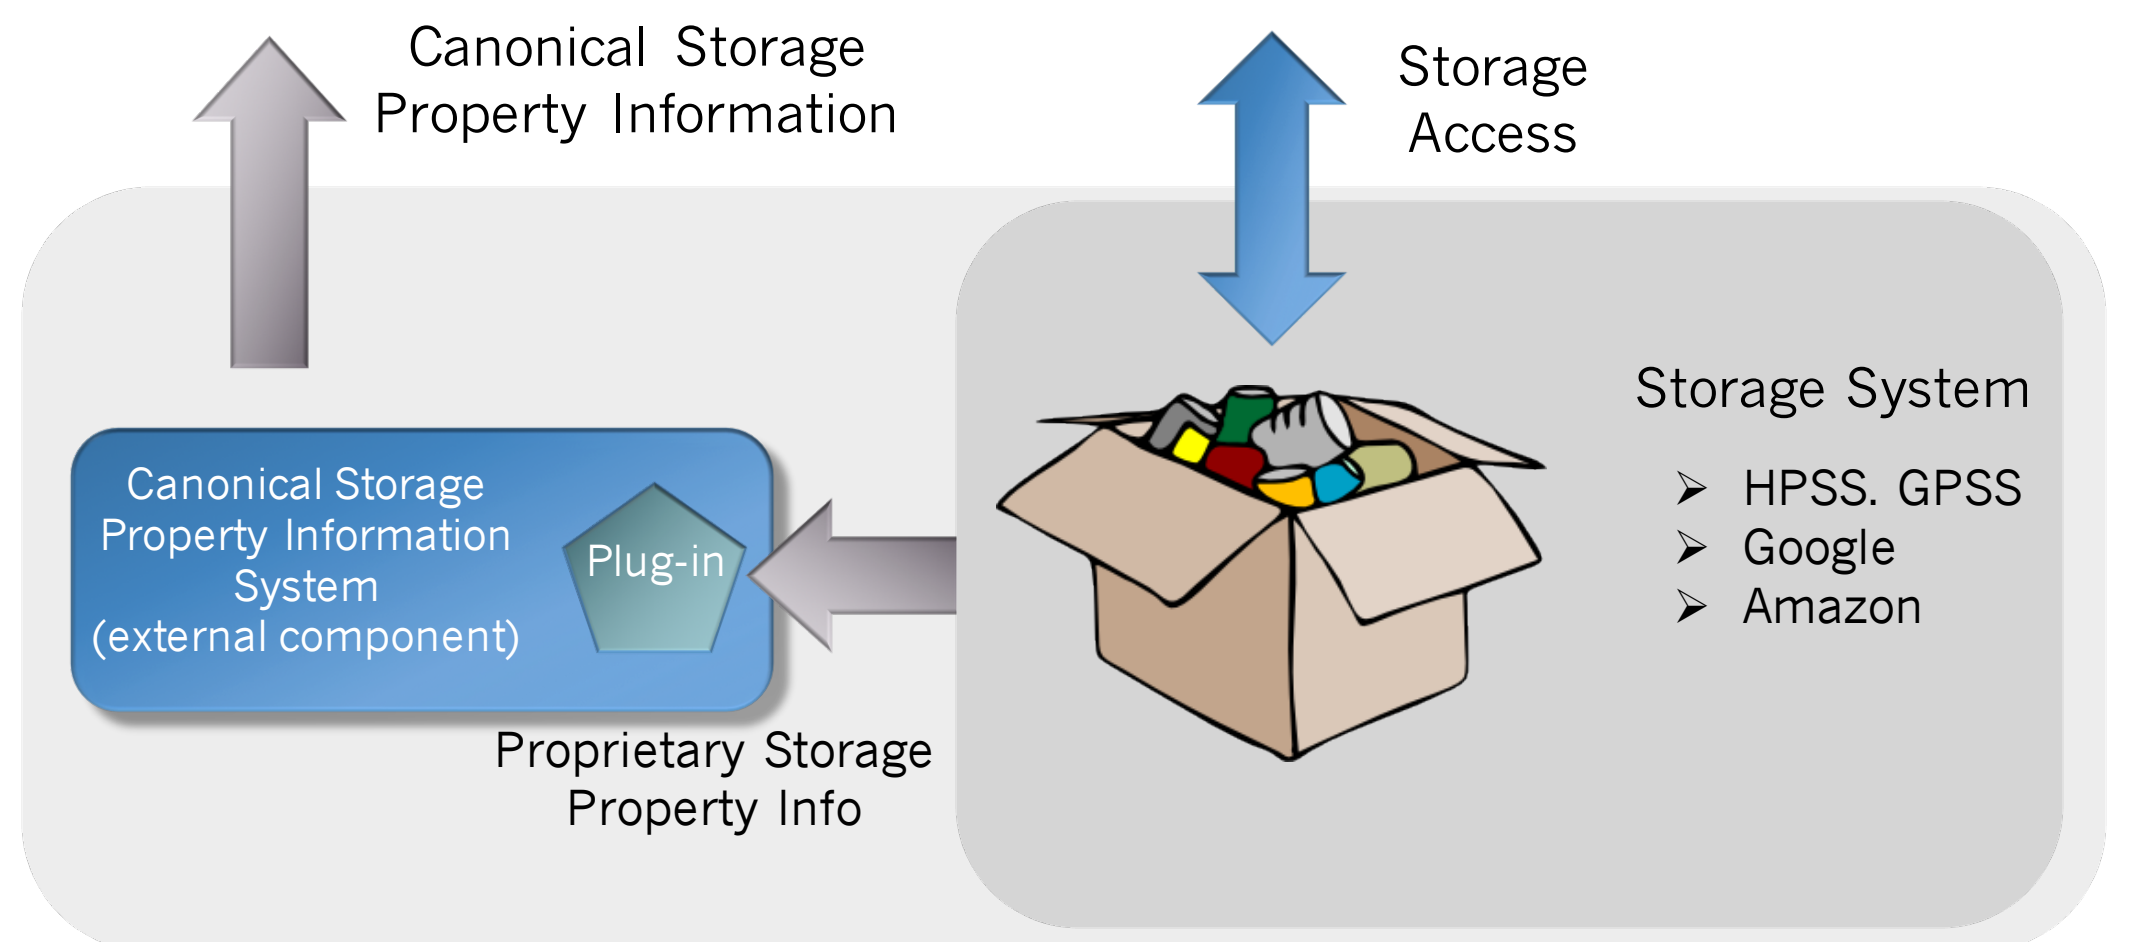

Customer View

- The **canonical view** only helps to describe the system on the **technical level.**
- It’s not very helpful for the storage user.
- We need to introduce more convenient **QoS views**.

QoS views

Ambiguous, non canonical, dependent, combined properties.

Examples :

Low latency & lowest price

High throughput & super durable

Large volume & cheap & archive

Translation and discovery

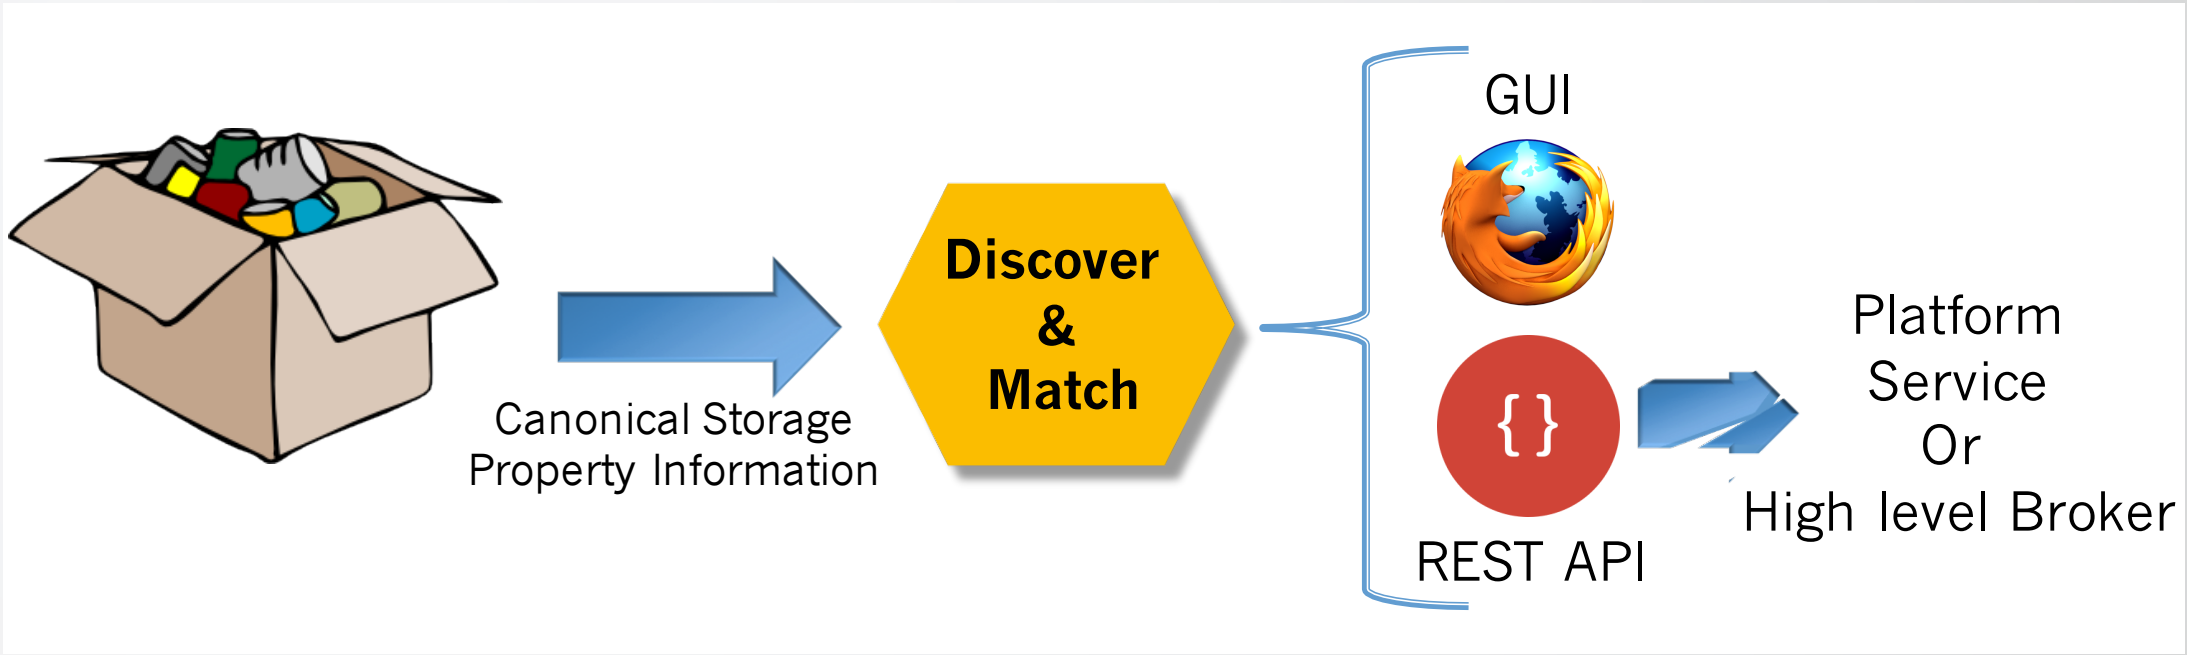

Canonical property federation

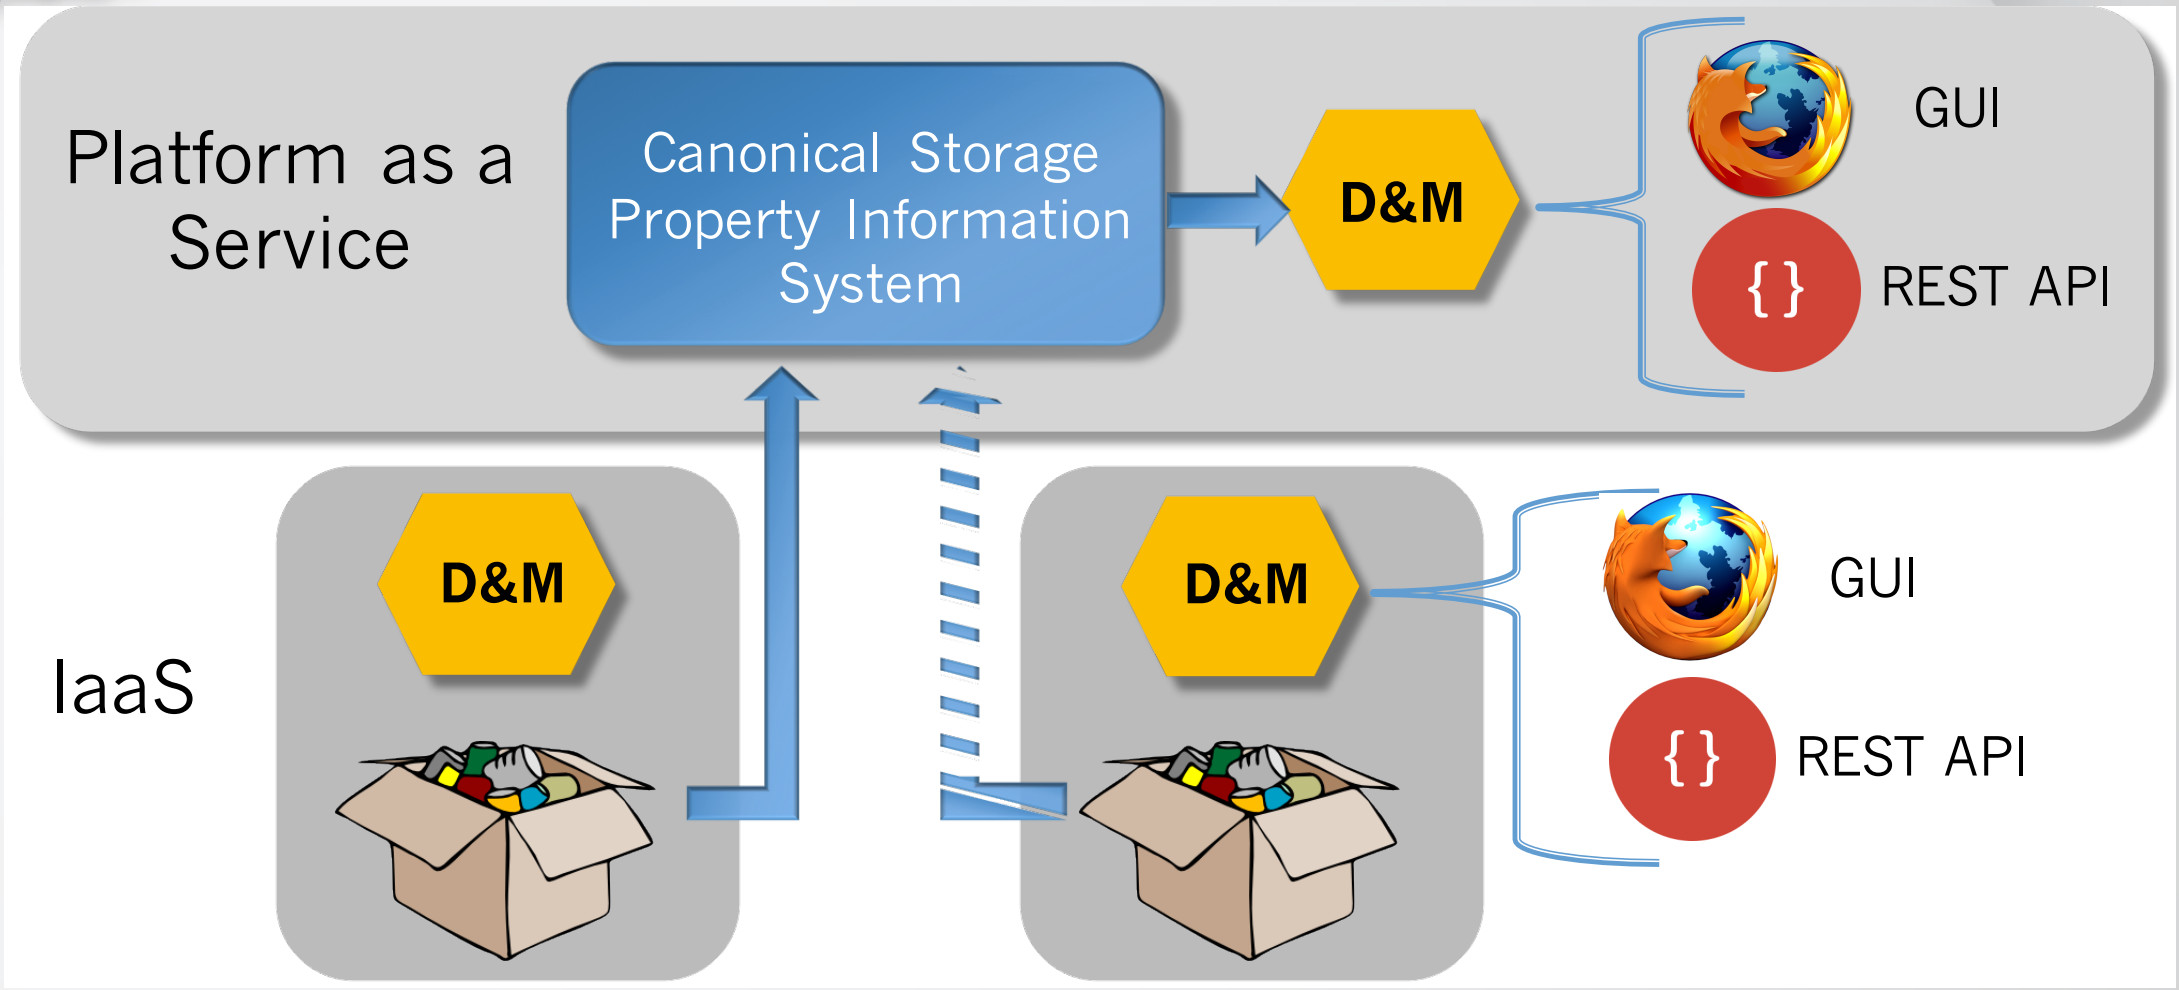

Federated Systems

- The federated system provides additional QoS properties.
  - Number of copies, not in the same location
  - Minimum geographic distance for disaster cases. (fire, earthquakes)
  - Legal implications: Privacy laws
- Federated system might need more higher level services attached:
  - FTS or Globus Online to create replicas
  - DynaFed to federate distributed resources.

More problems to solve

- How does the client provide the storage class to the storage system ?
  - Bucket
  - Directory
  - Additional argument in WebDAV, FTP etc
- The system only provides the class, it doesn't 'promise' the space.
  - Do we need a space reservation protocol ?
  - Is reservation required in systems with unlimited space (Clouds) ?
- Do we allow to change the storage class, assuming the system will do the necessary data movements ?
  - This is of course just a storage system property. Amazon and Goolge don't
  - dCache and HPSS do.

Next step: Data Life Cycle

- Data Life Cycle is just the time dependent change of
  - Storage Quality of Service
  - Ownership and Access Control (PI Owned, no access, Site Owned, Public access)
  - Payment model : Pay as you go ; Pay in advance for rest of lifetime.
  - Maybe other things

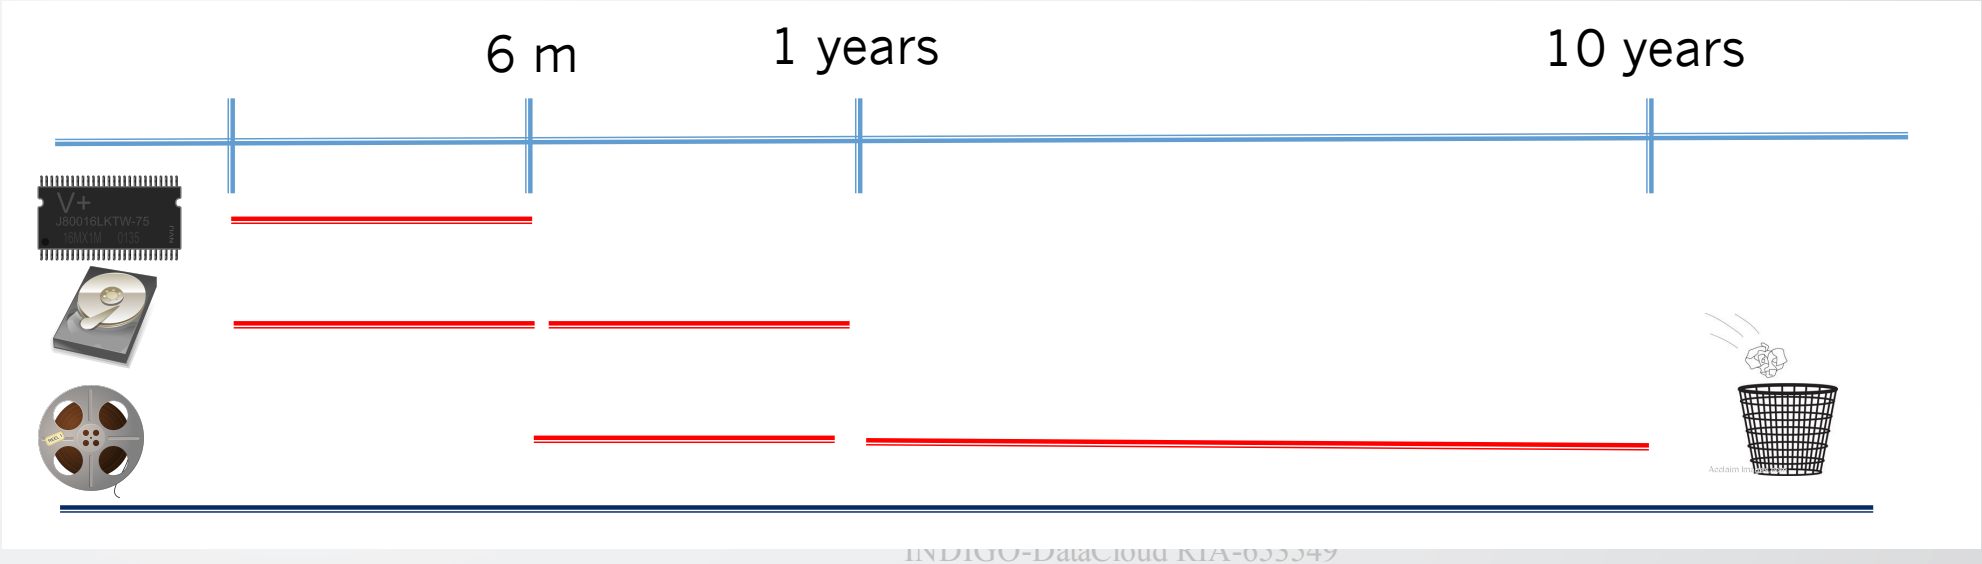

Current status (definitions)

- Introduced at the research data alliance (RDA) in Paris
- Lots of interested communities and sites.
- Creating of interest group in progress.
  - Name still in heavy discussion
- 10 Committed members
- Will be followed up on in Tokyo end of Feb 2016

Current status (technically)

- Canonical Information providers are being build
  - dCache (internal)
  - Common external system for GPFS/HPSS
    - CEPH
    - StoRM/GEMSS
  - Cloud (Amazon and Google)
- Information Provider Protocol in discussion (candidate : CDMI)

IaaS: Metadata-driven storage

INDIGO - DataCloud

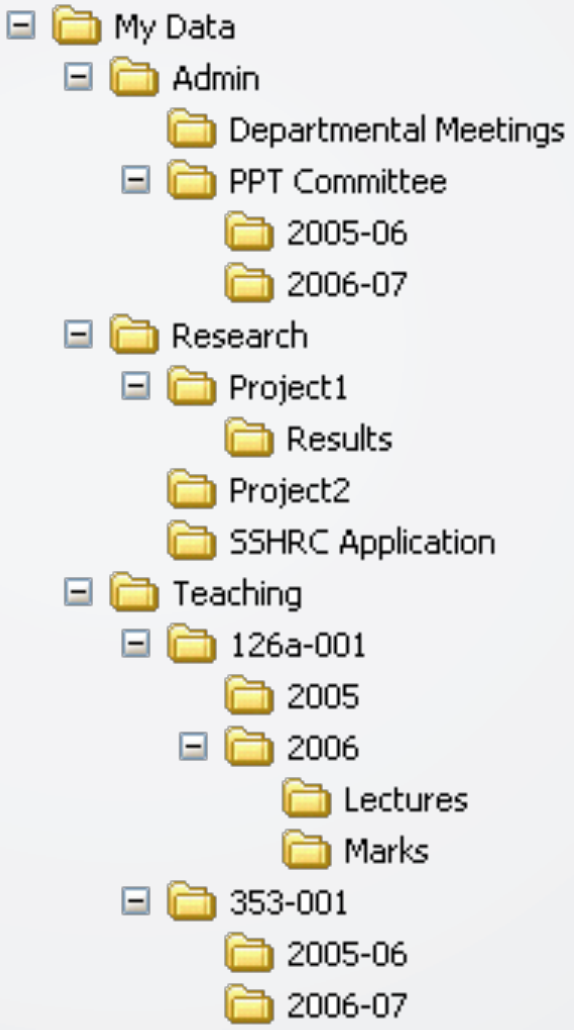

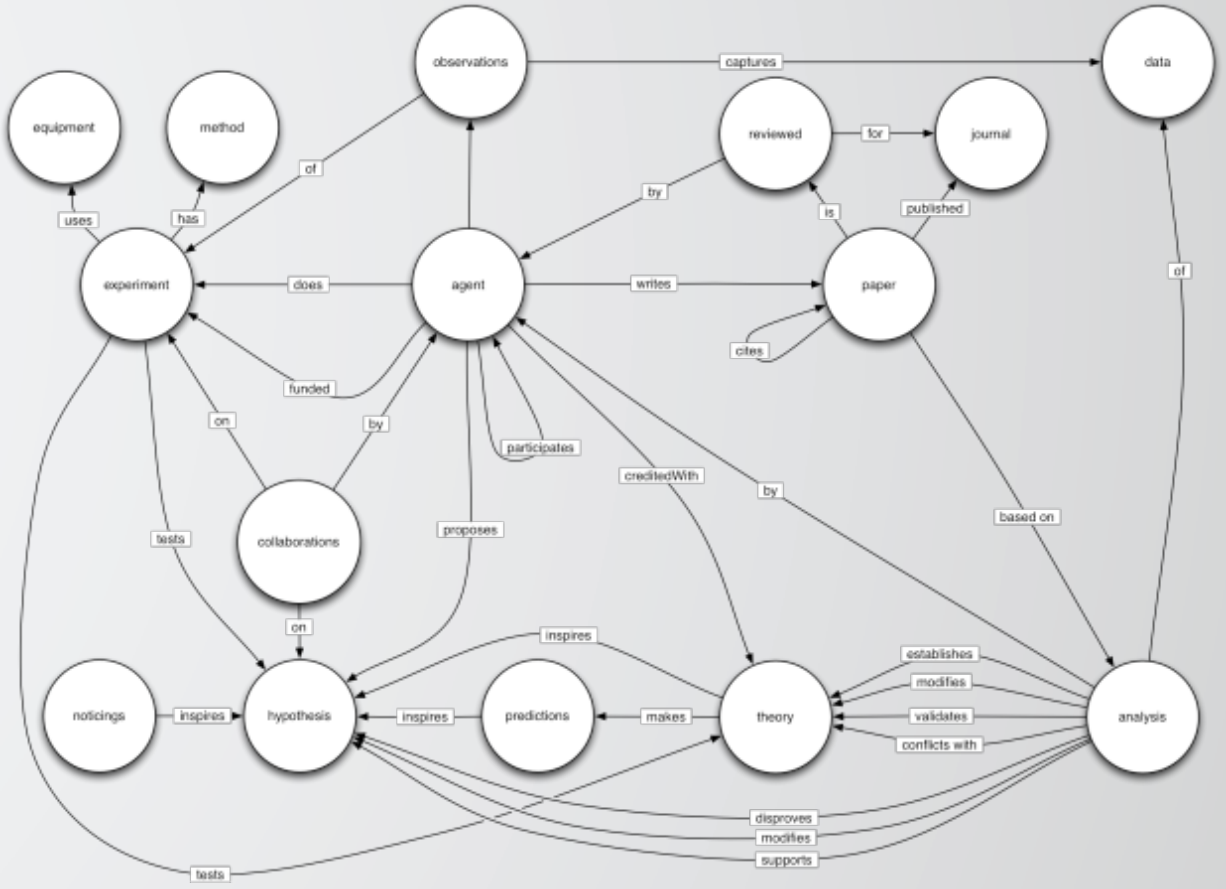

IaaS: laying hierarchical storage

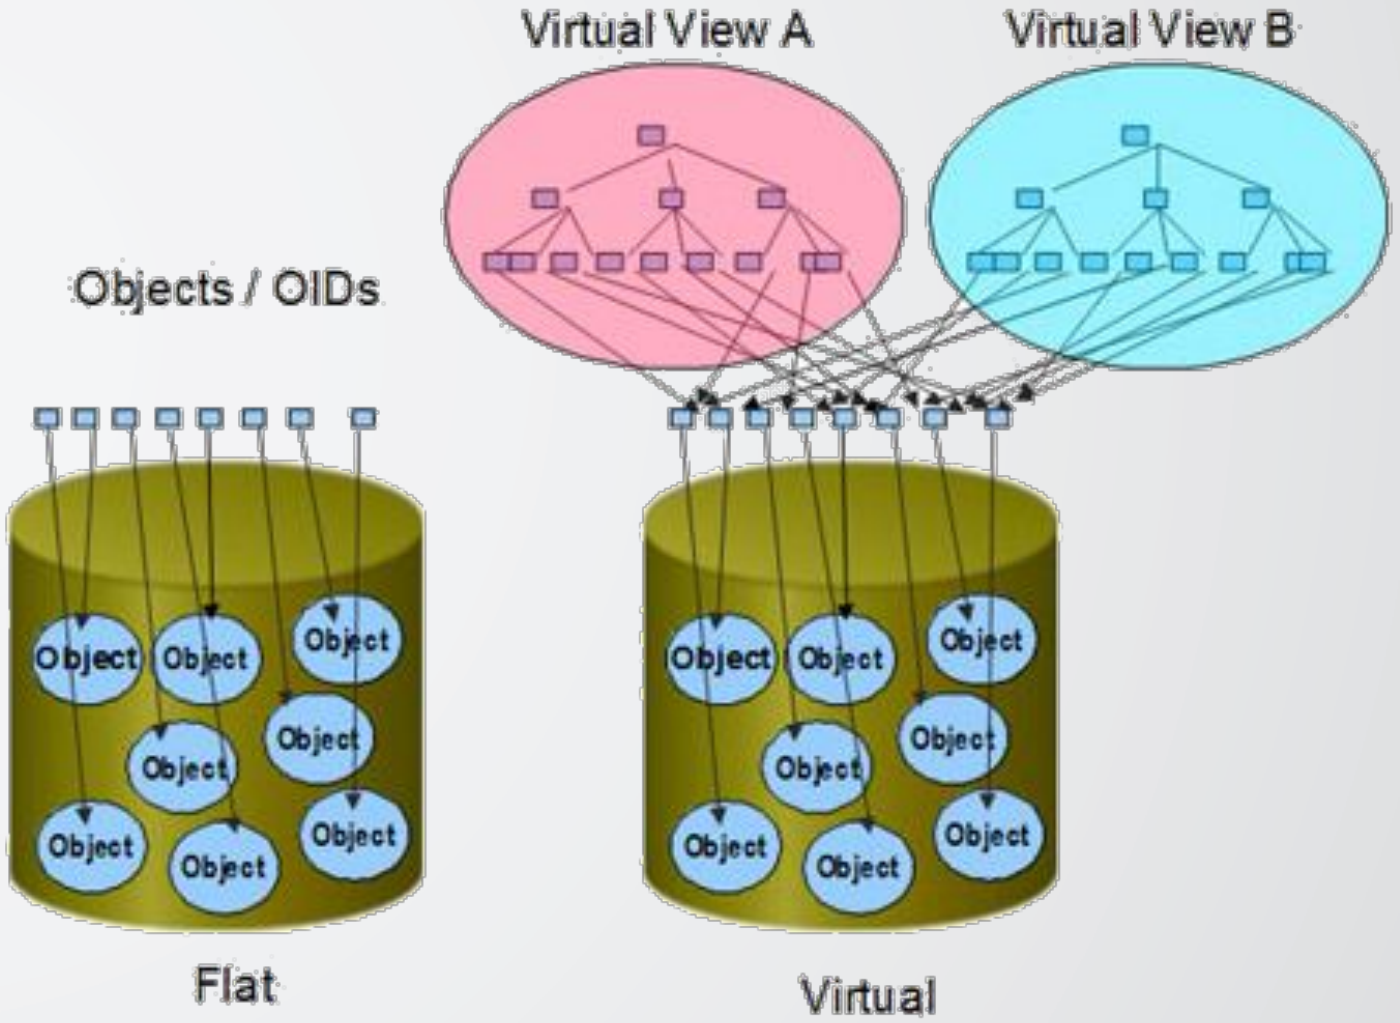

Identity and group-membership

- Allow **different** authentication mechanisms
  - SAML, OpenID-Connect, X.509, ...
- **Harmonise** user identities:
  - User is the same person, irrespective of how they authenticate
- Support **group-membership**:
  - Membership can be used for authorisation decisions.
- Support **third-party** group membership:
  - VOMS-style: where membership *not* asserted by authentication service.

WP5: PaaS Platform Development

INDIGO - DataCloud

- **PaaS** layer is provided by WP5
- It is realized by means of a high-available and scalable µServices cluster
  - Hosted by Kubernetes
- It is able to provide two main approach:
  - High-Level automatized IaaS
  - Real PaaS services and application execution

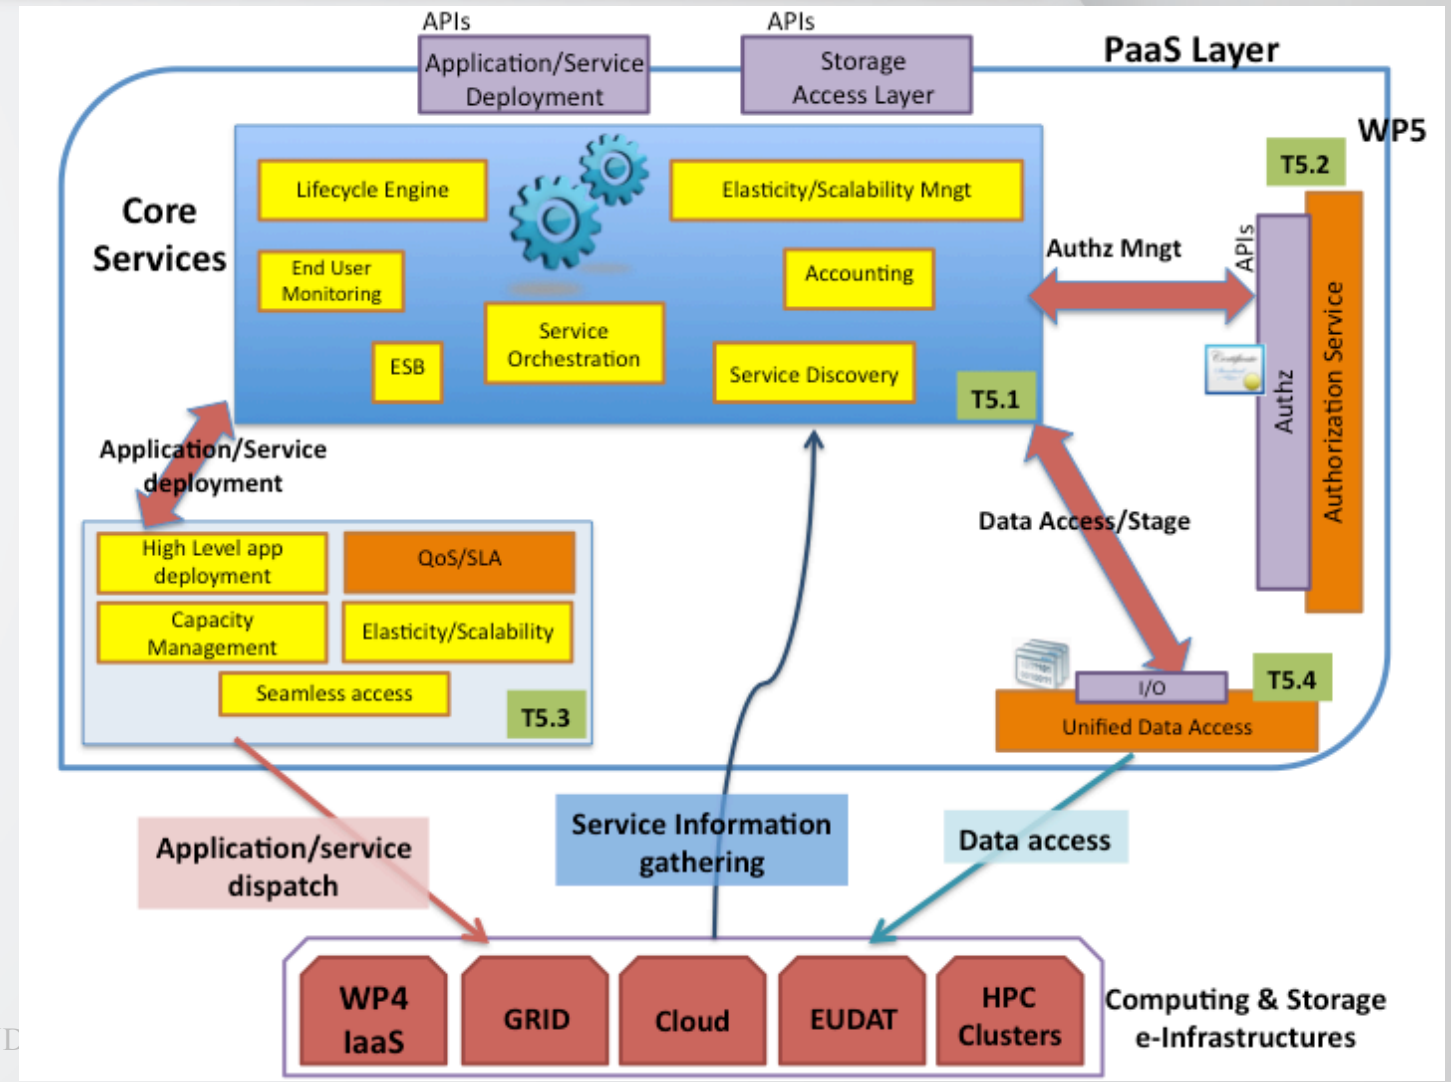

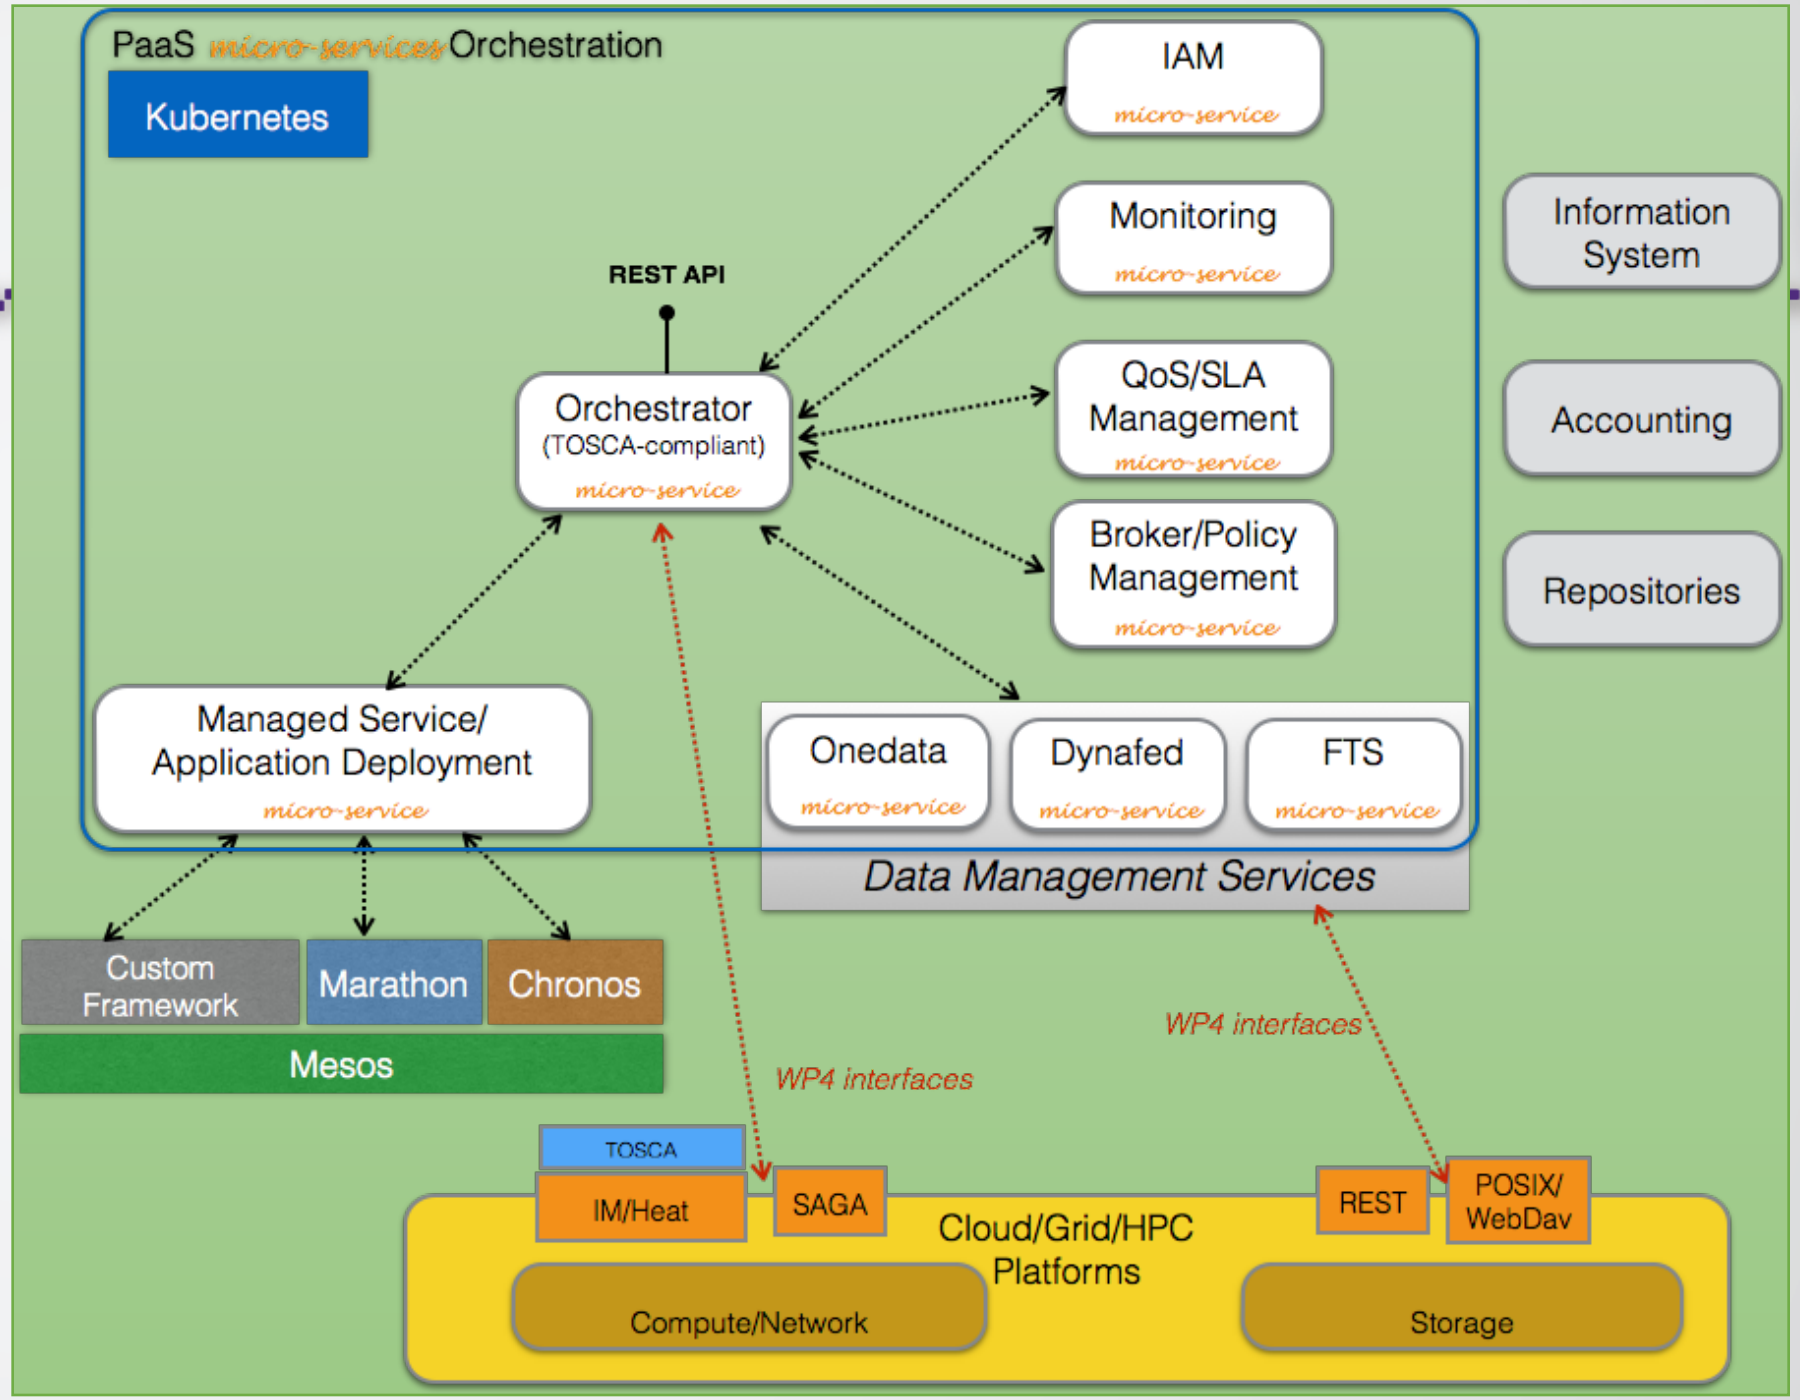

INDIGO - DataCloud

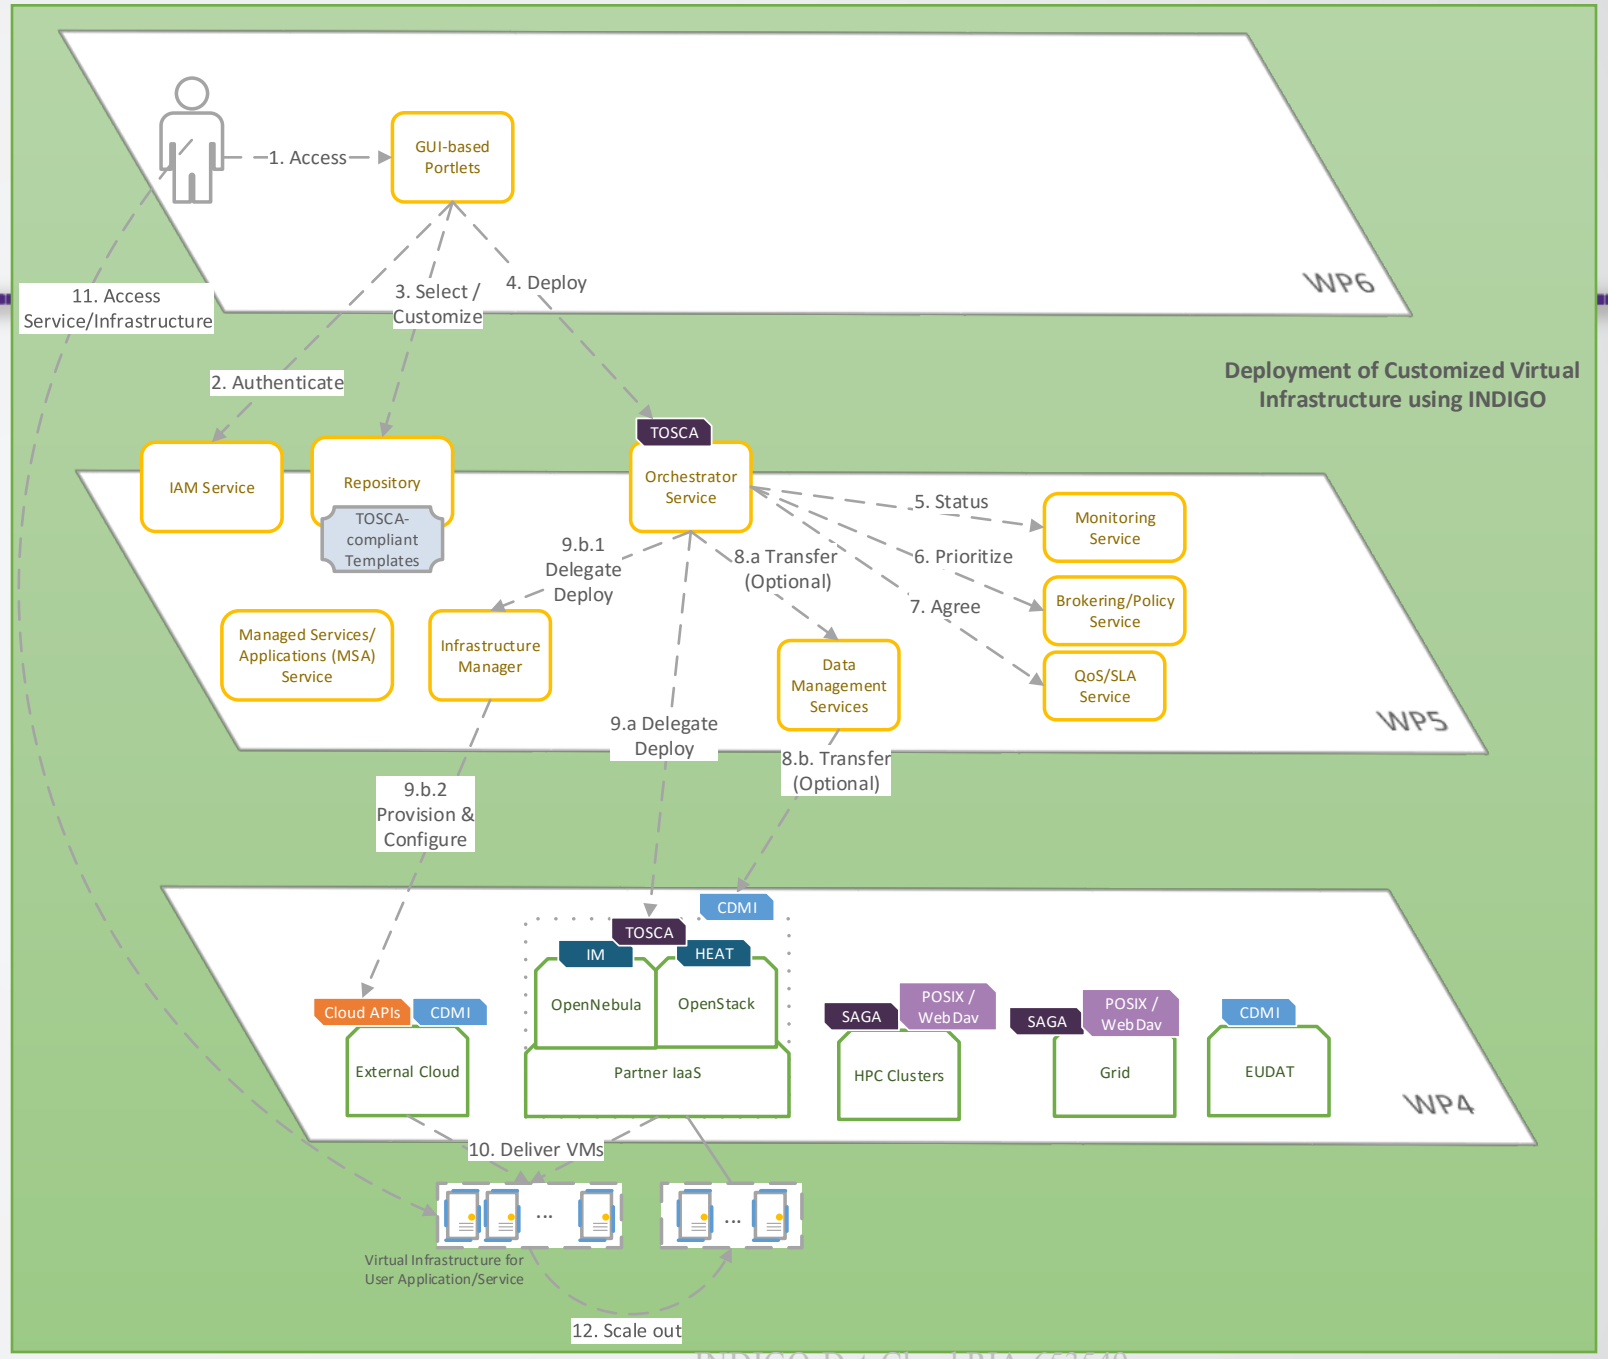
1. Access
GUI-based Portlets
WP6
11. Access Service/Infrastructure
3. Select / Customize
4. Deploy
2. Authenticate
Deployment of Customized Virtual Infrastructure using INDIGO
TOSCA
IAM Service
Repository
TOSCA-compliant Templates
Orchestrator Service
5. Status
Monitoring Service
9.b.1 Delegate Deploy
8.a Transfer (Optional)
6. Prioritize
Brokering/Policy Service
7. Agree
Managed Services/ Applications (MSA) Service
Infrastructure Manager
Data Management Services
QoS/SLA Service
WP5
9.a Delegate Deploy
8.b. Transfer (Optional)
9.b.2 Provision & Configure
CDMI
TOSCA
IM
HEAT
Cloud APIs
CDMI
OpenNebula
OpenStack
SAGA
POSIX / WebDav
SAGA
POSIX / WebDav
CDMI
External Cloud
Partner IaaS
HPC Clusters
Grid
EUDAT
WP4
10. Deliver VMs
Virtual Infrastructure for User Application/Service
12. Scale out

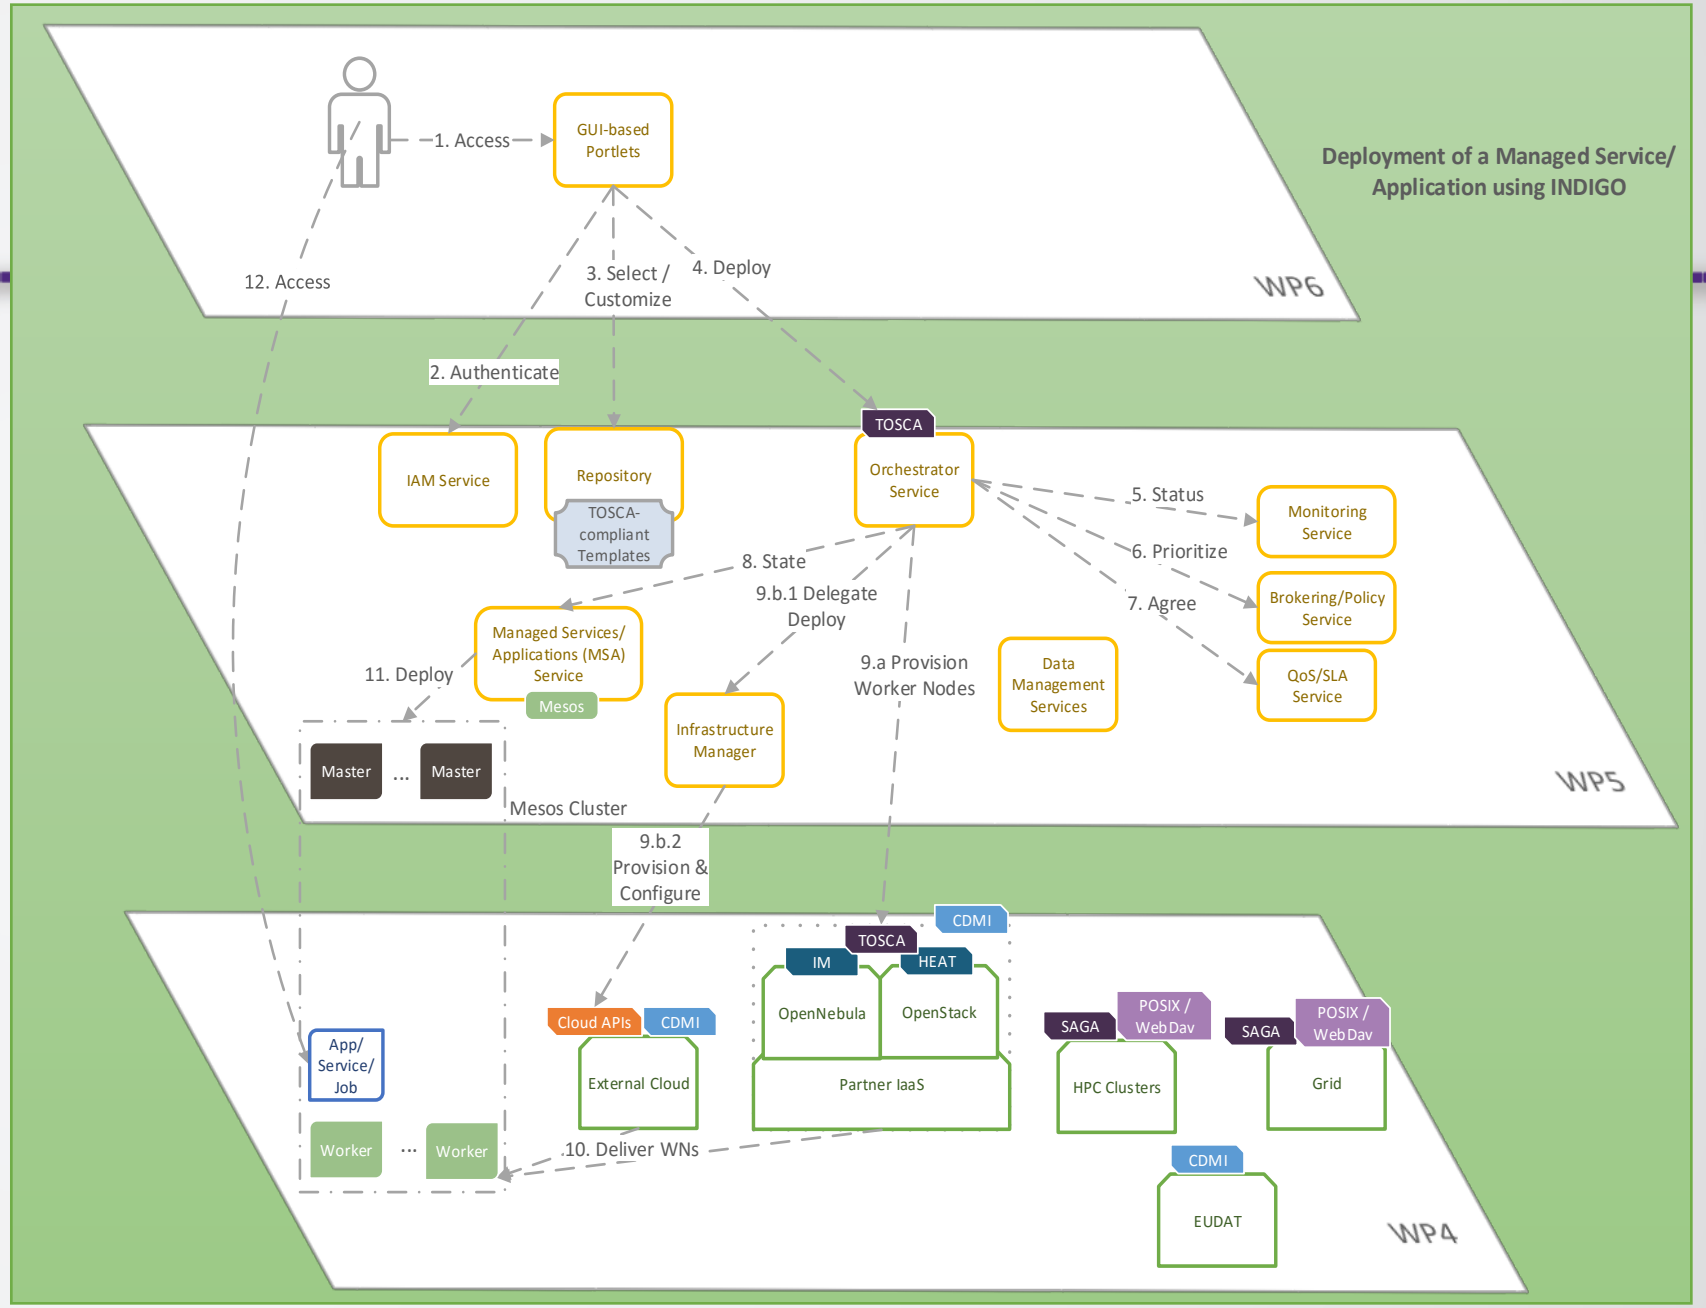
Deployment of a Managed Service/ Application using INDIGO
1. Access
GUI-based Portlets
12. Access
3. Select / Customize
4. Deploy
WP6
2. Authenticate
IAM Service
Repository
TOSCA-compliant Templates
TOSCA
Orchestrator Service
5. Status
6. Prioritize
7. Agree
Monitoring Service
Brokering/Policy Service
QoS/SLA Service
8. State
9.b.1 Delegate Deploy
9.a Provision Worker Nodes
Data Management Services
Managed Services/ Applications (MSA) Service
Mesos
11. Deploy
Infrastructure Manager
Master
Master
Mesos Cluster
WP5
9.b.2 Provision & Configure
CDMI
TOSCA
IM
HEAT
OpenNebula
OpenStack
Partner IaaS
Cloud APIs
CDMI
External Cloud
SAGA
POSIX / WebDav
HPC Clusters
SAGA
POSIX / WebDav
Grid
App/ Service/ Job
Worker
Worker
10. Deliver WNs
CDMI
EUDAT
WP4

WP5: Storage Federation

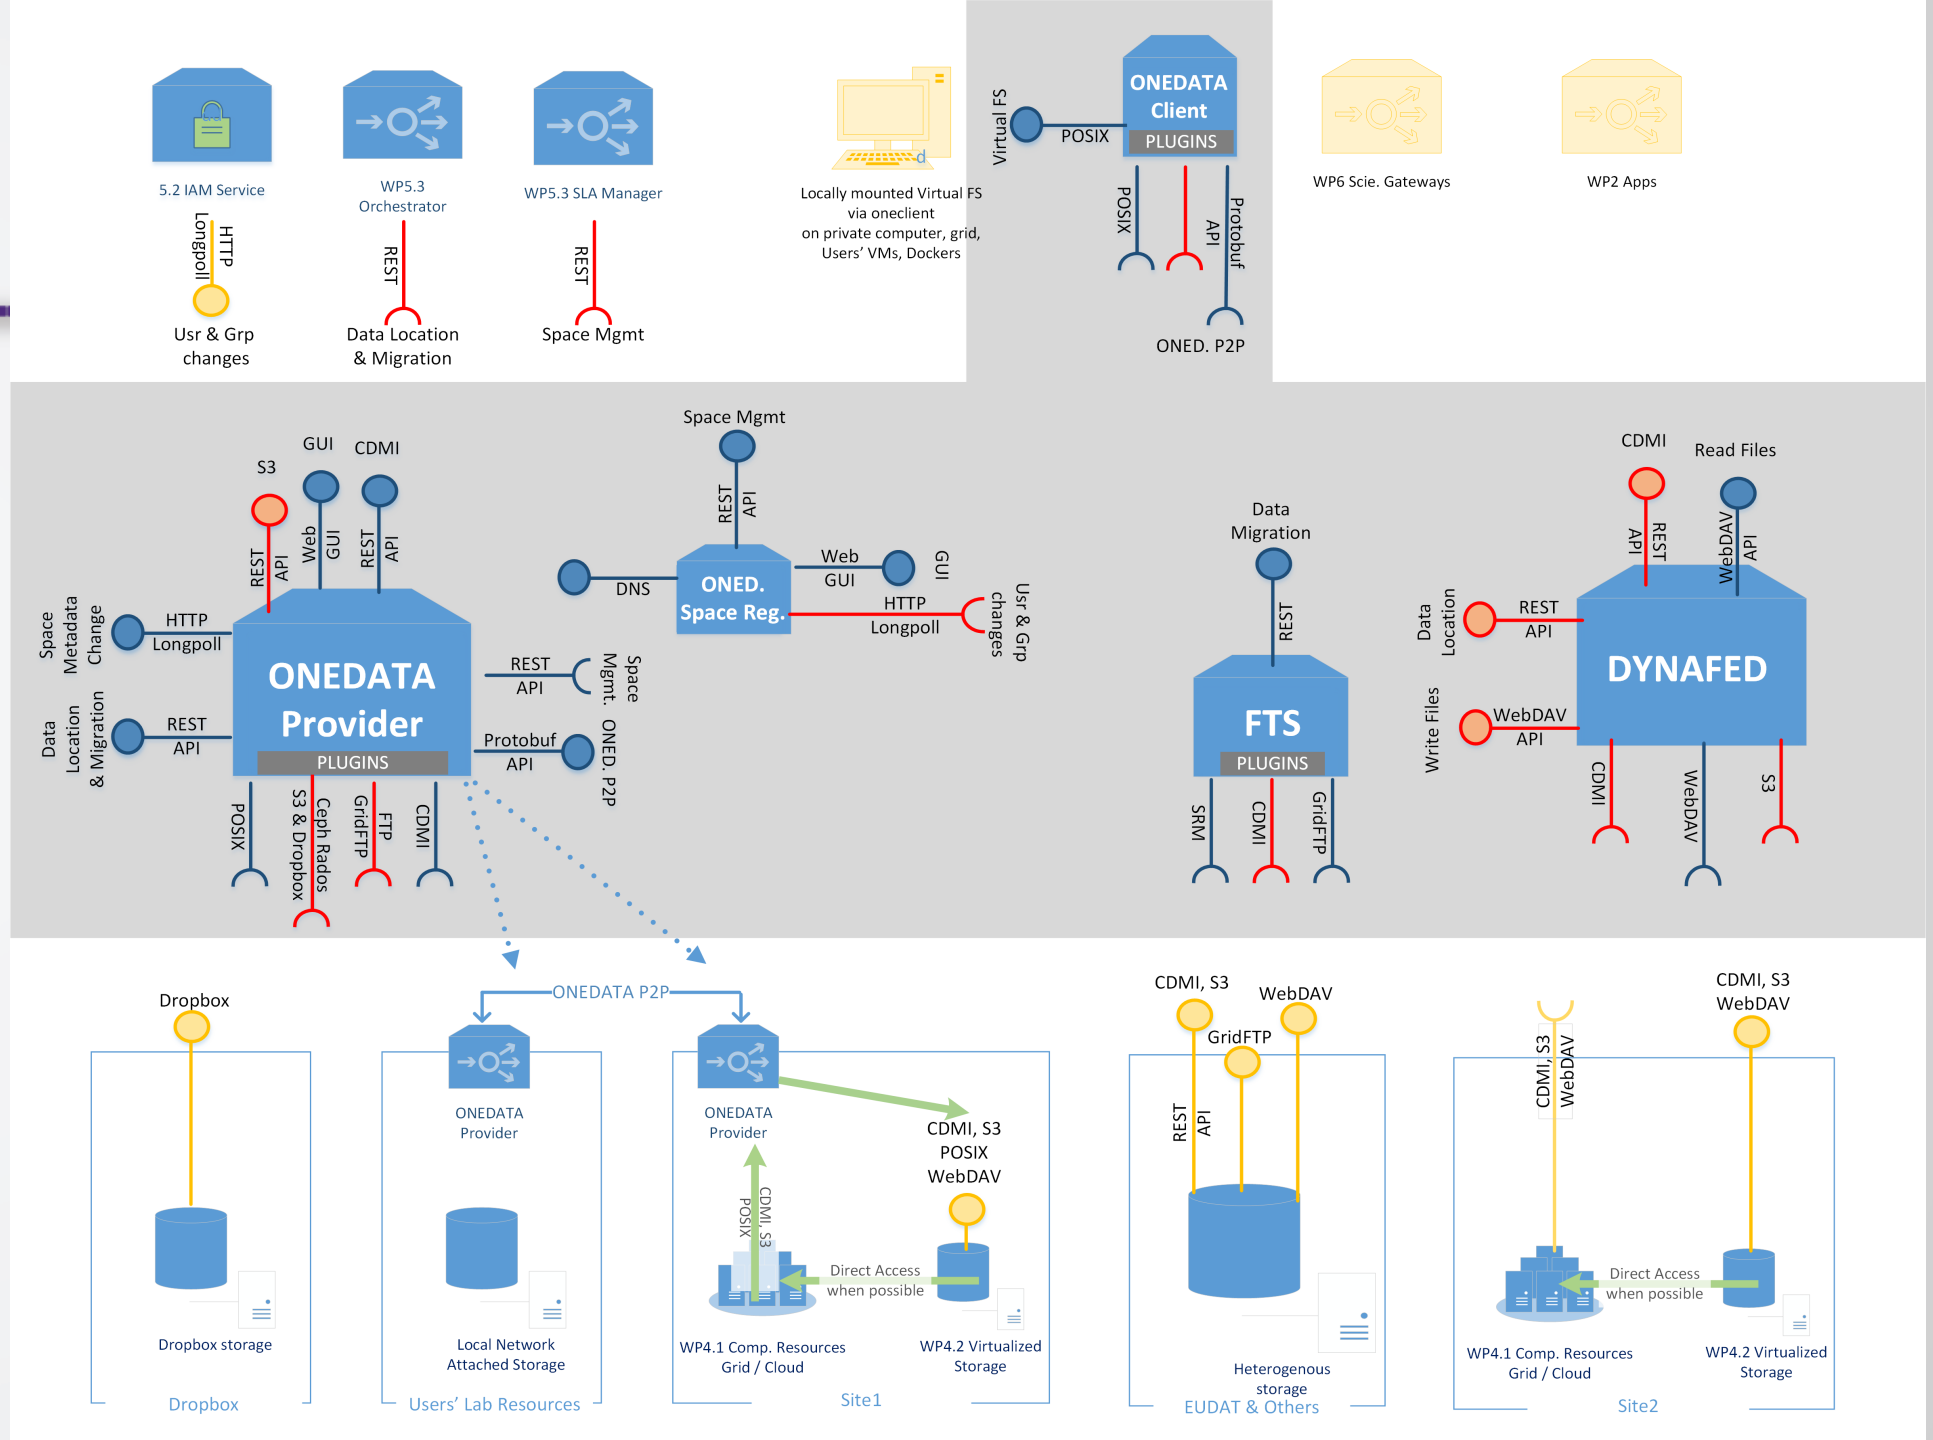

Conclusions

- First official release will be: end of July next year
- We will start make available some services as soon as they are ready enough to be tested
- All the changes on the already available projects, will be pushed back to the official releases.
  - OpenStack, OpenNebula, dCache, Onedata, Mesos, Accounting, QoS/SLA, etc

https://www.indigo-datacloud.eu

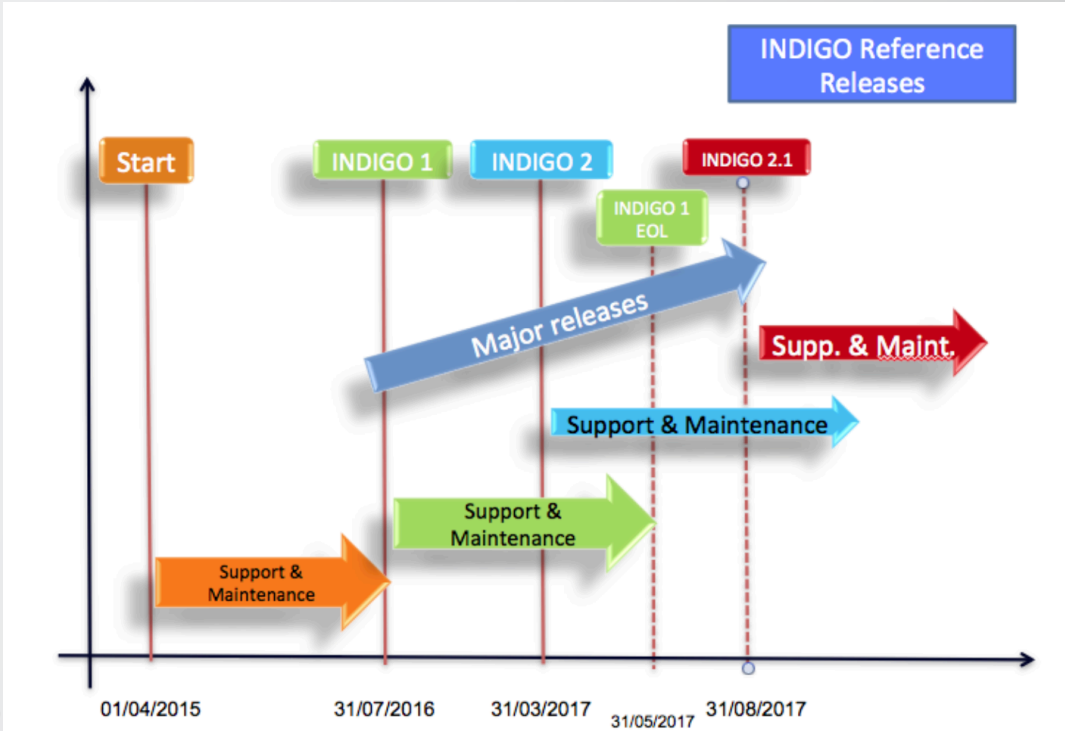

Figure 8: INDIGO-DataCloud Release Timeline

Questions?

Giacinto DONVITO

giacinto.donvito@ba.infn.it

Back-up Slides

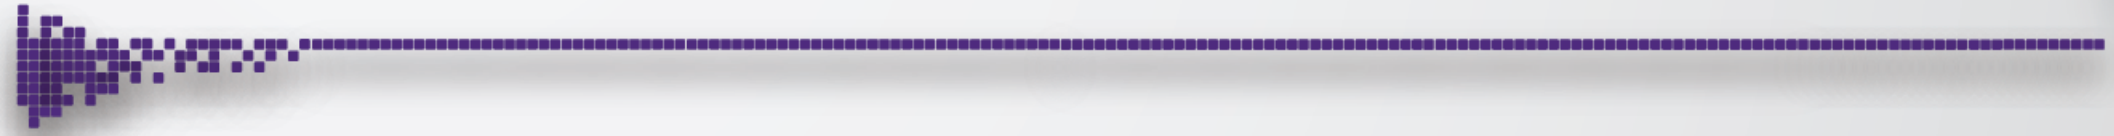

INDIGO Topic of interest

- Since the very beginning we identified key issues with both Grid and Cloud technologies that prevented scientific communities an easy and optimal exploitation of data and compute resources.
- We therefore decided to propose the development of a software platform centered around two of the EINFRA-1-2014 pillars:
  - Large scale virtualization of data/compute center resources. This became the focus of INDIGO WP4.
  - Development and adoption of a standards-based computing platform (with an open software stack). This became the focus of INDIGO WP5.

How do we see distributed computing in the future

1. Ease of access and use for small and big collaborations alike.
2. Software and economic sustainability.
3. Robustness (no single points of failure).
4. Modular, scalable architecture.
5. Open source software, vendor independence, hybrid infrastructures.

Users communities

SIMPLIFIED IMPACT TABLE SELECTED OBJECTIVES versus REQUESTS/ POTENTIAL IMPACT FOR COMMUNITIES O1: Development of the INDIGO Platform based on open software without restrictions on the e-Infrastructure	Life Sciences	Physical Sciences & Astronomy	Social Sciences & Humanities	Environmental Sciences
Research Communities & Initiatives , including ESFRIs	ELIXIR INSTRUCT/ WeNMR EuroBioImaging	CTA LBT WLCG	DARIAH DCH-RP	EMSO LIFEWATCH ENES
Examples of Applications	HADDOCK GROMACS AMBER GALAXY	MIDAS, IRAF, IDL, Geant4 ROOT/PROOF Geant4	Fedora Digital Libraries	Delft3D R-Studio TRUFA MATLAB
Design and development of a Platform providing advanced users and community developers a powerful and modern environment for development work. This includes programming and scripting tools, and composition of custom applications and software deployment	RELEVANT	**CRITICAL**	RELEVANT	**CRITICAL**
Developing a framework to enable the transparent execution on remote e-infrastructures of existing popular applications like MATLAB / OCTAVE, ROOT, MATHEMATICA, or R-STUDIO.	RELEVANT	**CRITICAL**	MINOR	**CRITICAL**
Provide the services and tools needed to enable a secure composition of services from multiple providers in support of scientific applications.	**CRITICAL**	**CRITICAL**	RELEVANT	RELEVANT
Develop and implement a solution that is able to deploy in a transparent and powerful way both services and applications in a distributed and heterogeneous environment made by several different infrastructures (EGI Grid and Federated Cloud, IaaS Cloud, Helix Nebula, HPC clusters)	**CRITICAL**	RELEVANT	MINOR	RELEVANT
Develop the capability in the PaaS to provide unified data access despite geographical location of data, including APIs access, based on existing standards, or virtually mount like a POSIX device to worker node, cloud virtual machines, personal computer etc.	**CRITICAL**	RELEVANT	**CRITICAL**	RELEVANT

INDIGO WPs

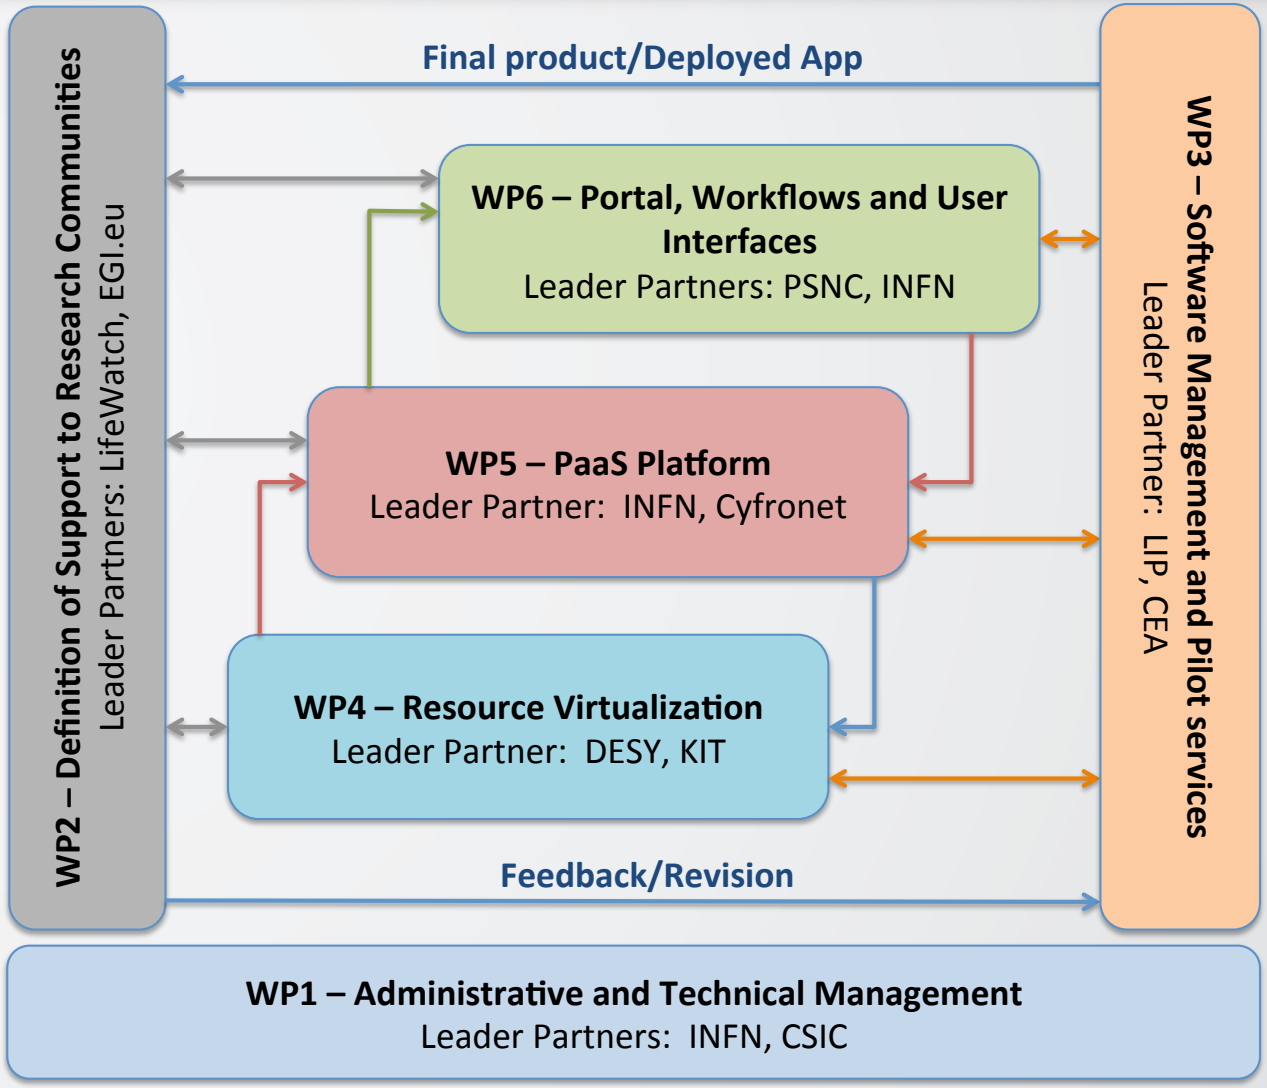
Final product/Deployed App
WP2 – Definition of Support to Research Communities
Leader Partners: LifeWatch, EGI.eu
WP6 – Portal, Workflows and User Interfaces
Leader Partners: PSNC, INFN
WP5 – PaaS Platform
Leader Partner: INFN, Cyfronet
WP4 – Resource Virtualization
Leader Partner: DESY, KIT
WP3 – Software Management and Pilot services
Leader Partner: LIP, CEA
Feedback/Revision
WP1 – Administrative and Technical Management
Leader Partners: INFN, CSIC

INDIGO approach

- Based on **Open Source** solutions
- widely **supported** by big communities
- whenever possible exploit **general solutions** instead of specific tools/services
  - or put effort in **increasing the generality** of tools developed in a given community
  - this will be important for **sustainability** of the architecture
- ensure that the framework offered to final users, as well as to developers, will have a **low learning curve**
  - **existing software suites** like ROOT, OCTAVE/MATLAB, MATHEMATICA or R-STUDIO, **will be supported** and offered in a transparent way