The European Open Cloud for e-Science towards automation, service composition, big data analytics and new frontiers in data management

Davide Salomoni (davide@infn.it)

ISGC 2018

Taipei, 16-23 March 2018

Disclaimer

This talk was prepared by me in my personal capacity. The opinions expressed here are my own, and do not necessarily reflect the view, policy or position of anybody else.

My background: INFN (National Institute for Nuclear Physics) – www.infn.it

INFN & Computing

- A long tradition, from the first small clusters to GRID and Cloud based large scale computing
- INFN not interested in computing per-se, but as a mean for its research related activities
- In the last 10 years, this has principally meant supporting the Experiments @ CERN (LHC)
- Currently, INFN operates:
  - 9 medium size centers (Tier-2s in the Worldwide LHC Computing GRID hierarchy)
  - 1 large WLCG Tier-1 center, at CNAF (Bologna) – certified ISO-27001
- All the centers are connected with at least 10 Gbit/s dedicated connections, currently being upgraded to 100 Gbit/s
- Collectively, our main centers have about 50,000 CPU cores, 40PB of enterprise-level disk space, 60PB of tape storage

Agenda

- I'll cover first some general points on the European Open Science Cloud (EOSC)…
- Then I'll show some concrete steps toward the implementation of the EOSC (a few bits on the EOSC-hub project)…
- But mostly I will deal with technical topics on Cloud automation, service composition and data management.

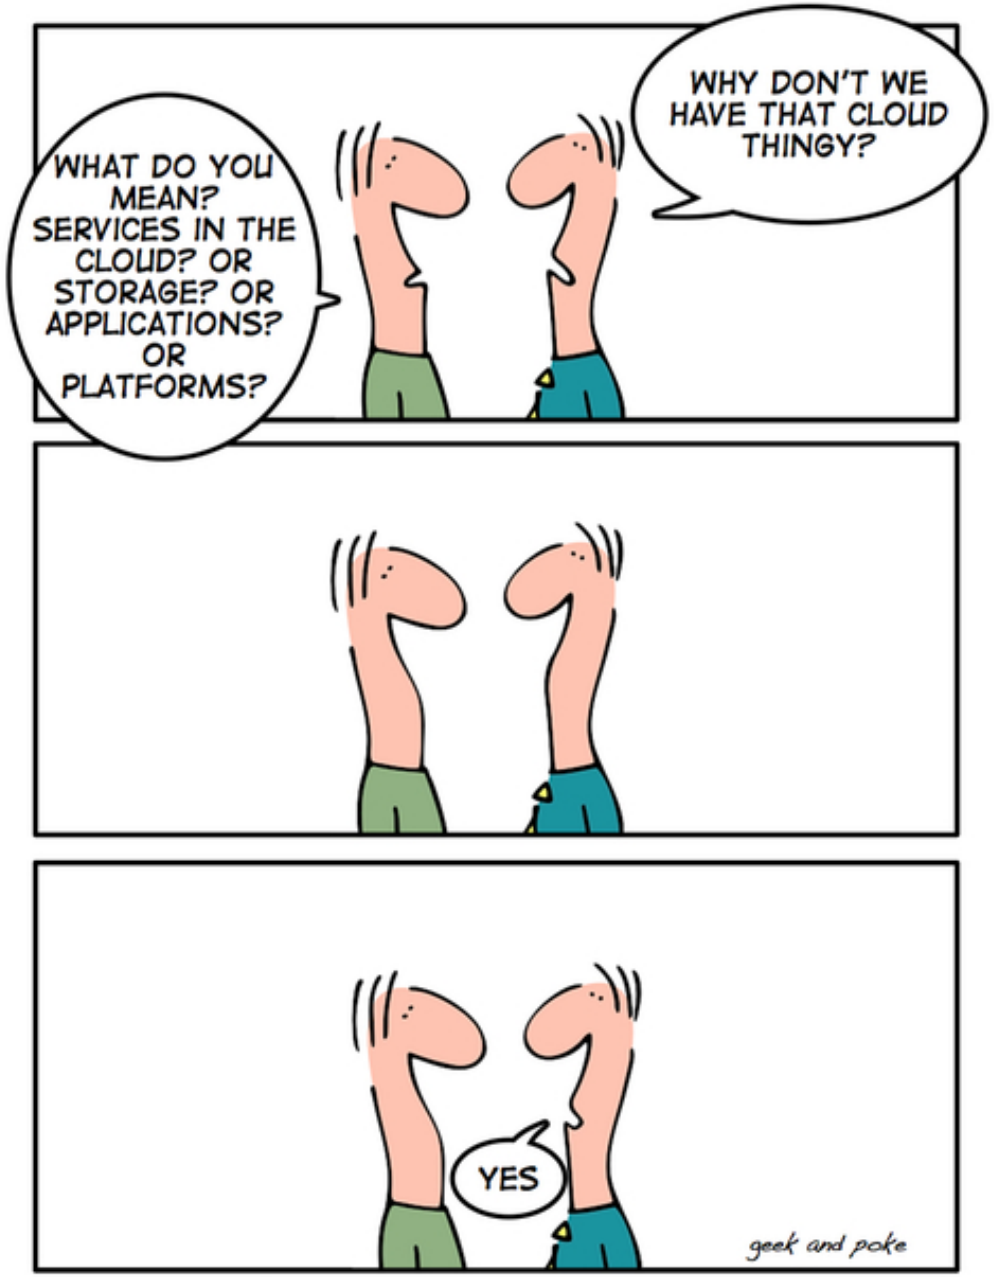

THE CLOUD THINGY

The European Open Science Cloud (EOSC)

- The EOSC: "**a model for the use of a cloud in the private and public sectors**" (European Parliament resolution on the European Cloud Initiative, Feb 2017)
- **Why the EOSC?**
  - To facilitate scientific developments and make the EU a center for global research
  - EO**S**C: "S" as in "Science", but with its user base to be extended to industry and governments
  - To foster the growth of the European Digital Economy → competiveness, global market positioning (esp. for SMEs)
  - To accelerate work on standards & interoperability, sharing of open data, creation of an open environment for storing, sharing and re-using scientific data and results, with the overall goal of removing fragmentation.

Key points of the EOSC

- **A data infrastructure common and a federation of existing resources**, which will:
  - Re-use existing building blocks & state-of-the-art services and solutions whenever possible
  - Develop and offer services based on user needs
- In summary: leverage past investment in research data infrastructures to **add value in terms of scale, interdisciplinarity and faster innovation, with a clear business model for sustainability.**

The EOSC Roadmap

- **Actions toward the EOSC** were included in the Horizon 2020 (H2020) Work Programme (WP) 2016-2017 and in the WP 2018-2020 for an aggregate budget of **about €600m**.
- An **EOSC roadmap working document** was published by the EC on March 14, 2018. According to this roadmap, **six action lines** are identified:
  1. Architecture → reduce fragmentation via interoperable federated infrastructures
  2. Data → FAIR data management and tools
  3. Services → Create a service catalogue starting from users' needs
  4. Access & Interface → how to access the EOSC
  5. Rules → rules of participation, legal & technical frameworks
  6. Governance → long-term orientation, priorities, compliance, implementation, monitoring, reporting

What then?

- **From vision to action: the EOSC-hub project**
- EOSC-hub, funded with about 30M€, mobilizes providers from 20 major digital infrastructures, EGI, EUDAT CDI and INDIGO-DataCloud, jointly offering services, software and data for advanced data-driven research and innovation through a **unified service catalogue.**
- Factoids:
  - 100 Partners, 75 funded beneficiaries
  - 3874 PMs, 108 FTEs, more than 150 technical and scientific staff involved
  - 36 months: Jan 2018 – Dec 2020

The EOSC-hub service structure

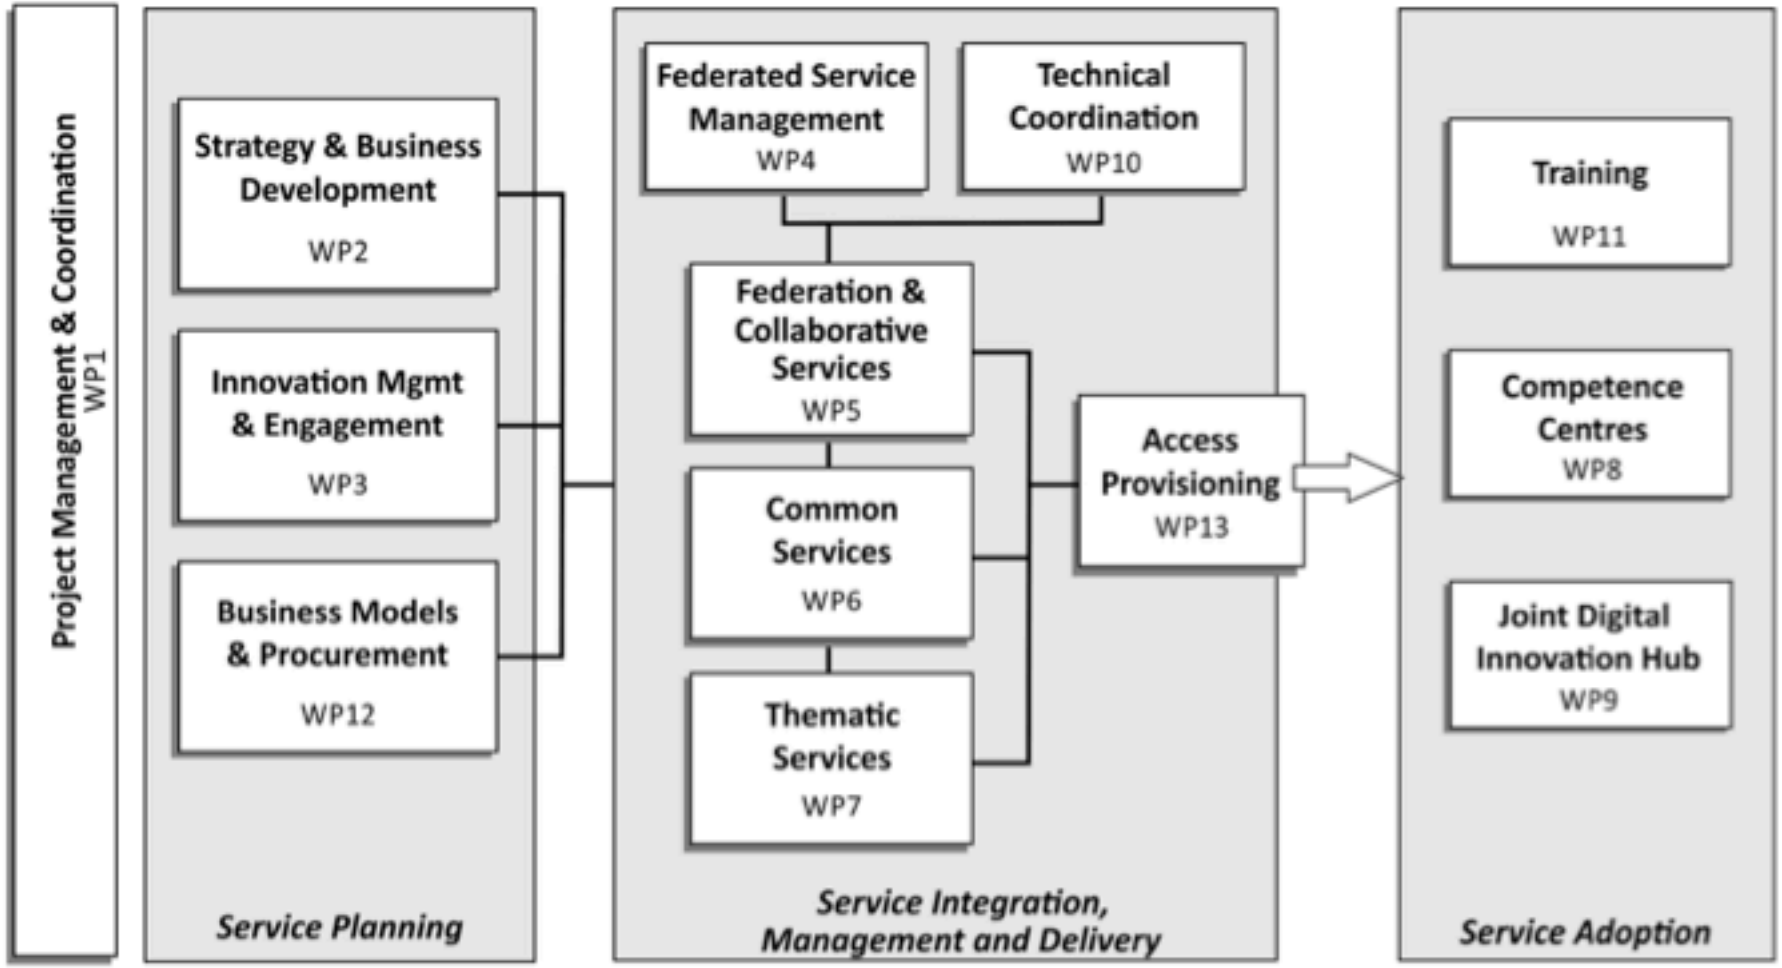

The EOSC-hub Architecture

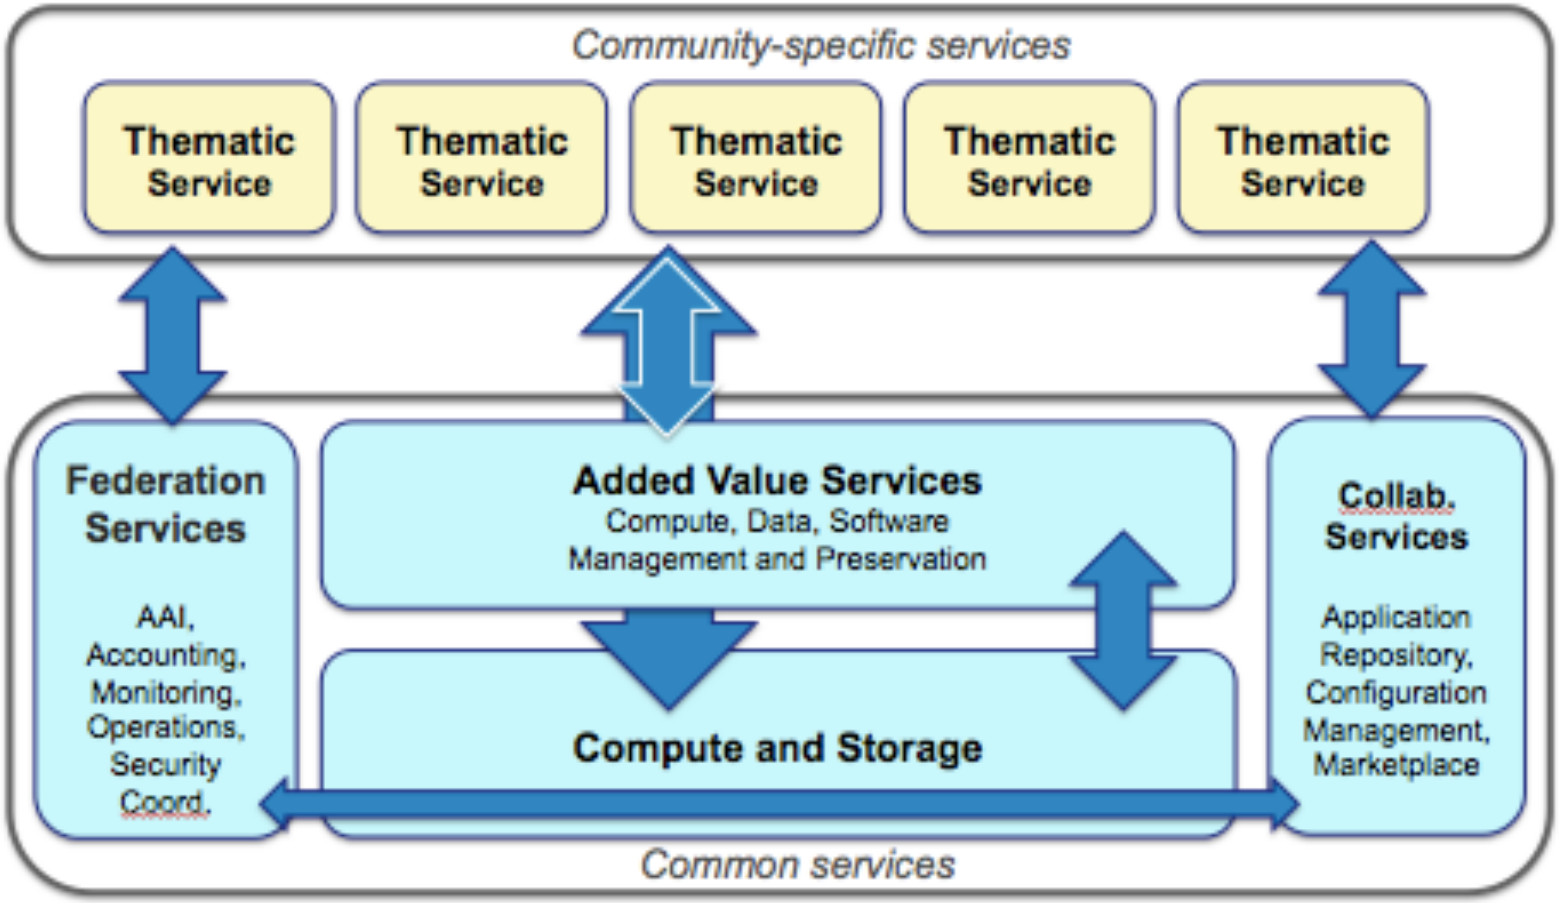

Some of the EOSC-hub services

Federation and Collaborative Services

- Identification, Authentication, Authorization and Attribute Management (EUDAT B2ACCESS, EGI CheckIn, INDIGO WaTTS)
- Marketplace and Order Management (Service Portfolio Management Tool, Data Project Management Tool)
- Integrated Business and Operations Support Systems (Configuration Management Data Base [GOCDB], EGI Operations Portal)
- Monitoring, Accounting, Messaging, Security Tools (ARGO, EGI Security monitoring tool, EUDAT accounting repository, APEL)
- Helpdesk Services and Tools (EGI & EUDAT)
- Application store, Software Repositories and other Collaboration Tools

Common Services: Integration and Maintenance

- Discovery and Access (INDIGO IAM, Onedata, CDMI storage services; EGI DataHub; EUDAT B2SHARE, B2FIND, B2STAGE, B2HANDLE, B2DROP; OpenStack Swift [external])
- Federated Compute (INDIGO advanced IAAS services [Tosca Heat translator, Cloud spot instances, udocker]; CREAM-CE, BDII, VOMS, STORM, CVMFS, ARGUS, DIRAC)
- Processing and Orchestration (INDIGO PaaS [TOSCA-based templates], FutureGateway, Open Mobile Toolkit)
- Data and Metadata management (EUDAT B2NOTE, B2SAFE, B2HANDLE)
- Preservation (Trusted Digital Repository [by DANS-KNAW])
- Sensitive Data (TSD [by SIGMA2], ePouta [by CSC])

The EOSC-hub Thematic Services and Competence Centers

Thematic Services

- **ECAS** – Climate Analytics Service (ECAS), provided by ENES
- **DARIAH** – Science gateway tailored for the digital arts and humanities communities
- **OPENCoastS** – On-demand Operational Coastal Circulation Forecast Service
- **GEOSS** – GEO DAB (Discovery and Access Broker), GEOSS portal
- **EO Pillar** – Earth observation services, coordinated by ESA
- **WeNMR** – Structural biology services
- **DODAS** – Dynamic On Demand Analysis Service (CMS & others)
- **LifeWatch** – Citizen science services, GBIF, Digital Knowledge preservation framework, remote monitoring and smart sensing
- **CMI** – The Component MetaData Infrastructure, including the Virtual Language Observatory and the Virtual Collection Registry, provided by CLARIN

Competence Centers

- **IFREMER**: unified data analytics platform for the marine science community wrt climatology and oceanography
- **LOFAR**: online platform for radioastronomy data storage and data analysis
- **Fusion/ITER**: Platform data storage and simulation & modelling for open data in fusion research
- **European Integrated Data Archive framework** (EIDA): tools for seismological data and services
- **ELIXIR**: dataset distribution service and tools for life science
- **EISCAT_3D**: data analysis for atmosphere and near-Earth space studies

Walking toward the EOSC…

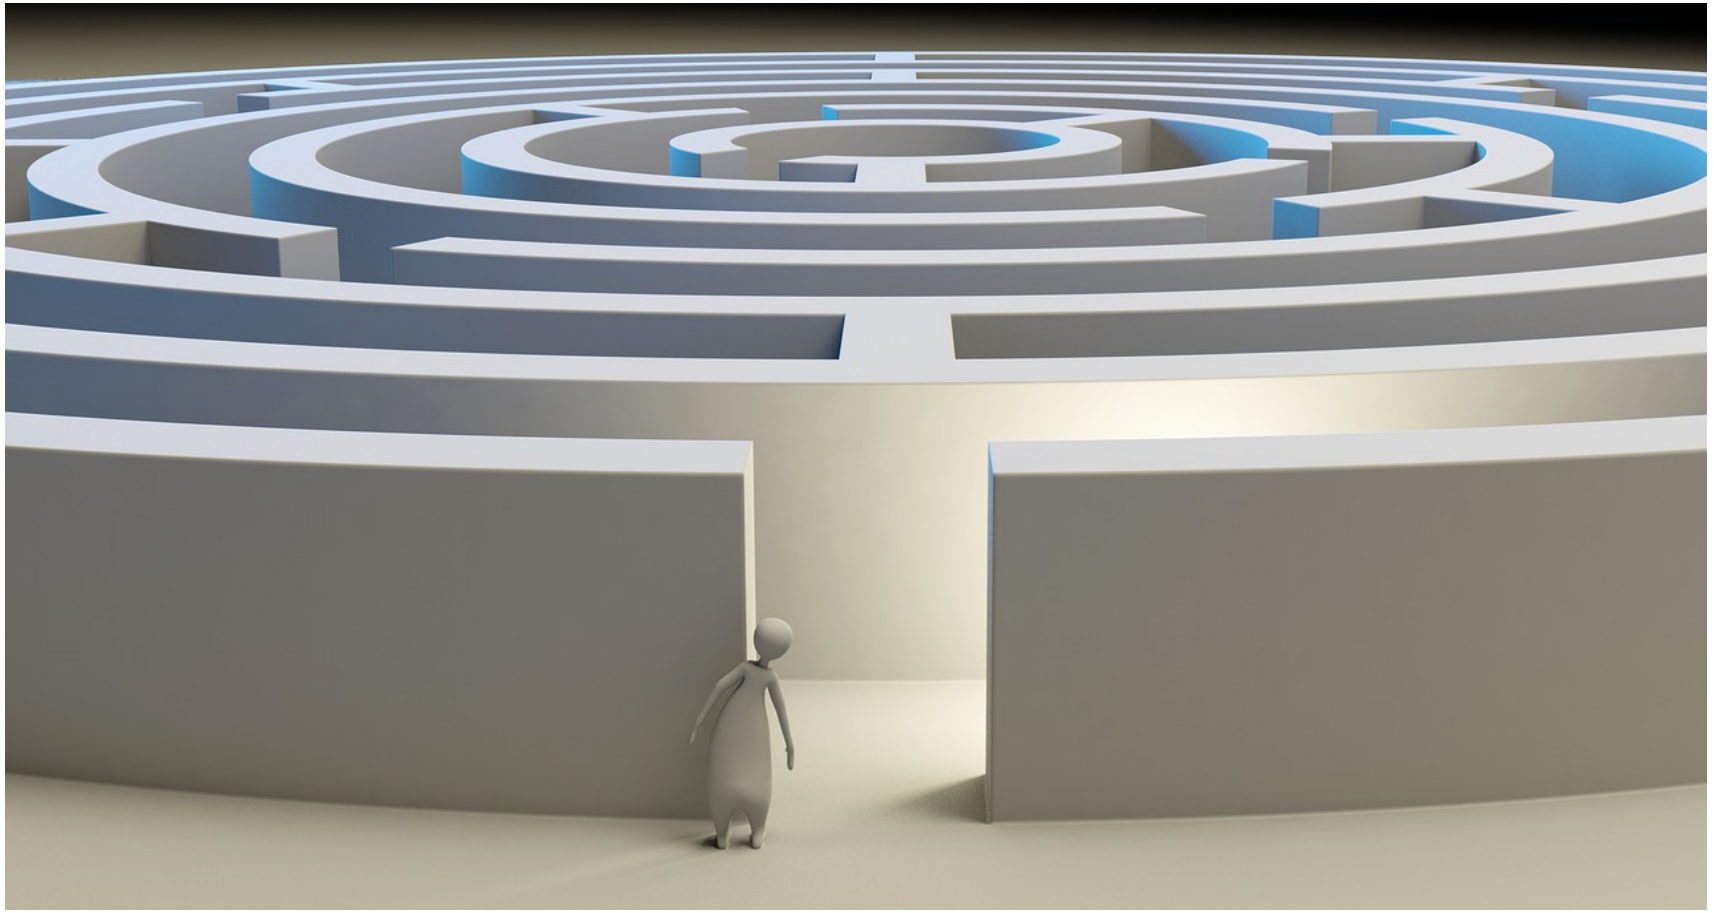

INDIGO-DataCloud

INDIGO - DataCloud

- **INDIGO-DataCloud** is one of the 3 H2020 initiatives (INDIGO-DataCloud, EGI-Engage and EUDAT2020) that jointly prepared the EOSC-hub project
  - 11.1M€, 30 months (from April 2015 to September 2017)
- **Who**: 26 European partners in 11 European countries, cordination by INFN
- **What**: develop an open source Cloud platform for computing and data ("DataCloud") tailored to science but applicable to other domains as well.
- **For whom**: multi-disciplinary scientific communities, such as structural biology, earth science, physics, bioinformatics, cultural heritage, astrophysics, life science, climatology.
- **Where**: deployable on hybrid (public or private) Cloud infrastructures
- **Why**: answer to the technological needs of scientists seeking to easily exploit distributed Cloud/Grid compute and data resources.

The INDIGO-DataCloud overall objective

To **develop open software components and solutions** to facilitate (or simply make possible) the **exploitation of distributed cloud and data resources** through **public or private infrastructures**.

“Better Software for Better Science.”

The INDIGO approach

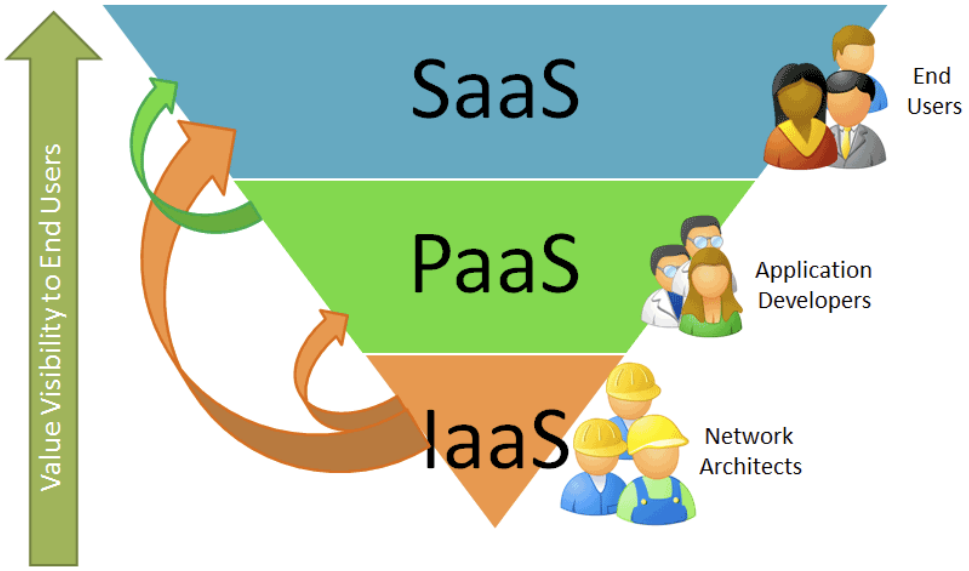

Scientific Users

Adopt, Use

INDIGO - DataCloud

INDIGO Advanced Components and Solutions

Deployed on

Private or Commercial Clouds (Public, PCP-based, etc.)

Publicly funded e-infrastructures (EGI, EUDAT, GEANT, PRACE, RI, etc.)

Exploiting

Datasets, Resources

To produce

Scientific Results

The INDIGO software

- INDIGO released two major software versions. The second and final one is called **ElectricIndigo.**
- **ElectricIndigo (https://www.indigo-datacloud.eu/service-component):**
  - 47 open source modular components, distributed via 170 software packages, 50 ready-to-use Docker containers
  - Supported operating systems: CentOS 7, Ubuntu 16.04
  - Supported cloud frameworks: OpenStack Newton, OpenNebula 5.x (plus connection to Amazon, Azure, Google)
  - Download it from the INDIGO-DataCloud Software Repository: http://repo.indigo-datacloud.eu/index.html

The ElectricIndigo Service Catalogues

ElectricIndigo is described in:

- A **high-level catalogue**, providing a bird's eye view of the INDIGO components
- A **detailed technical catalogue**

INDIGO - DataCloud
Better Software for Better Science

INDIGO-DataCloud High-level services

Explore INDIGO-DataCloud high-level services and discover how they have been applied to concrete use cases brought forward by scientific communities and resource providers.

www.indigo-datacloud.eu

INDIGO - DataCloud
Better Software for Better Science

INDIGO-DataCloud Service Catalogue

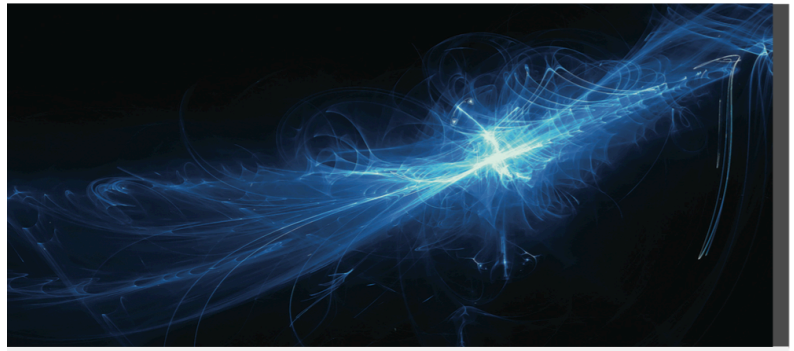

An unmatched open modular suite of software components for Data and Cloud computing is now available for resource providers and researchers from all disciplines, all around Europe

www.indigo-datacloud.eu

The ElectricIndigo Release

- The **ElectricIndigo modular software components** are organized around **5 areas**:

1. Application-level Interfaces to Cloud Providers and Automated Service Composition
   - For users **porting their apps to the Cloud**
2. Flexible Identity and Access Management
   - For users needing to **handle AAI**
3. Data Management and Data Analytics Solutions
   - For users **managing distributed [big] data**
4. Programmable Web Portals, Mobile Applications
   - For the **creation of front-ends**
5. Enhanced and Scalable Services for Data Centers and Resource Providers
   - For providers wishing to **optimize/enhance their service offerings**

The INDIGO-DataCloud Reference Architecture

- See details at http://arxiv.org/abs/1603.09536 or https://www.indigo-datacloud.eu/documents-deliverables

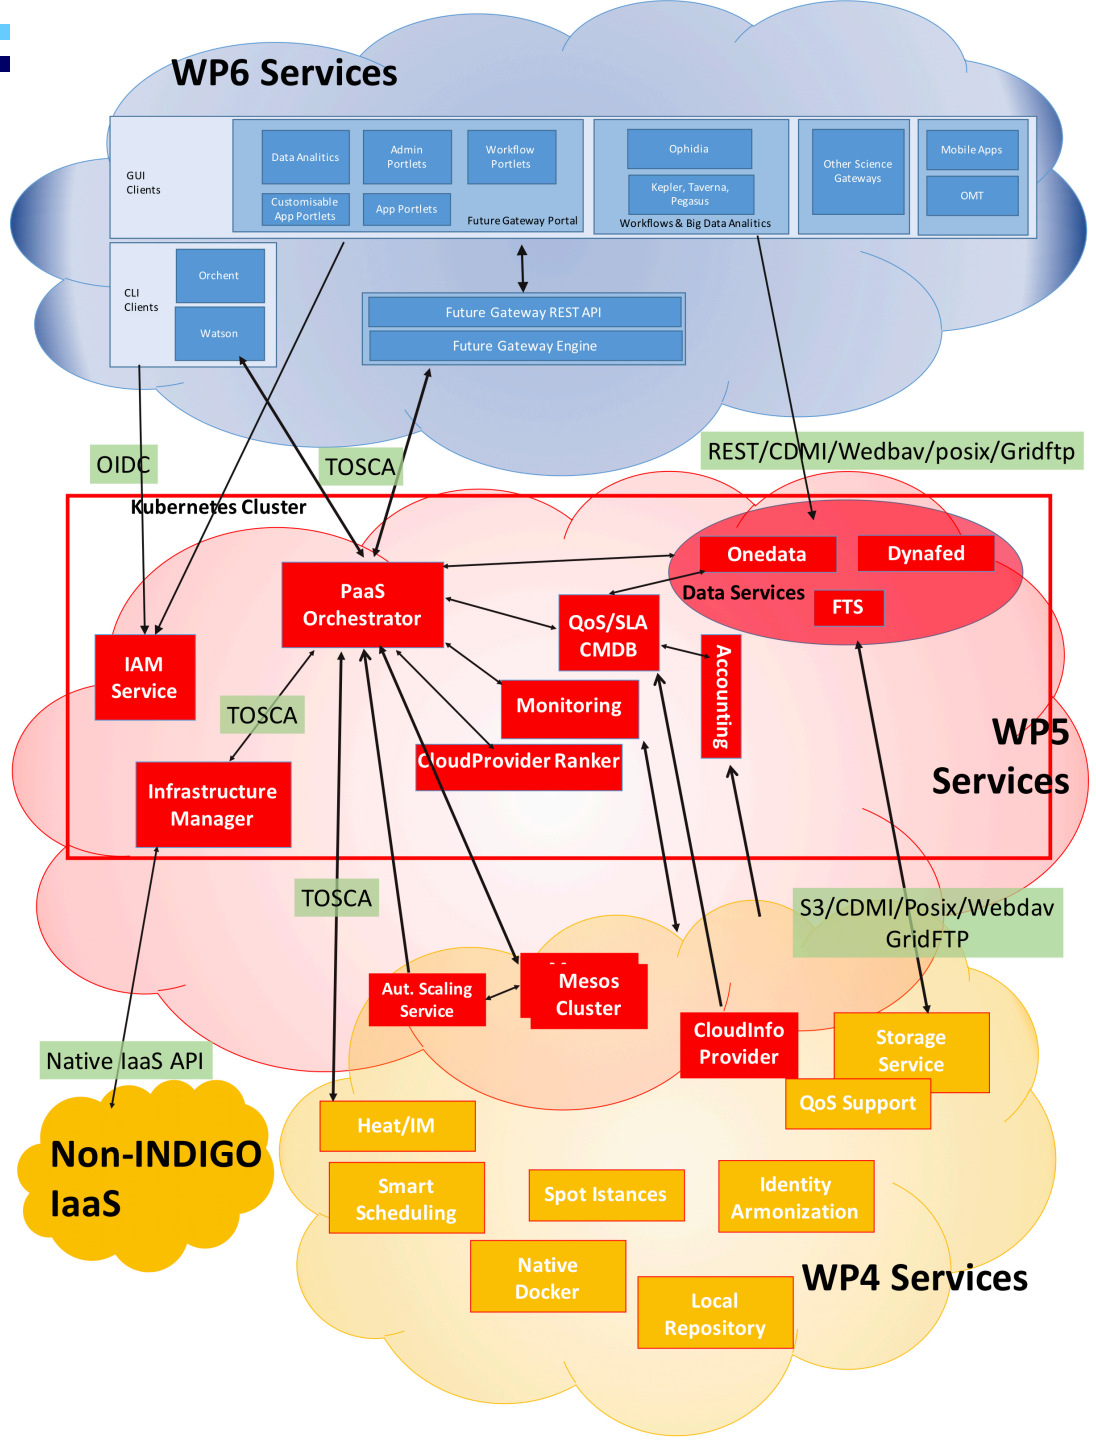

Moving beyond IaaS

IaaS vs PaaS vs SaaS

	IaaS	PaaS	SaaS
What you get	You get the infrastructure. Freedom to use or install any OS or software.	You get what you demand: software, hardware, OS, environment.	You don't have to worry about anything. A pre-installed, pre-configured package as per your requirement is given.
Deals with	Virtual Machines, Storage (Hard Disks), Servers, Network, Load Balancers etc.	Runtimes (like java runtimes), Databases (like MySQL, Oracle), Web Servers (tomcat, etc.)	Applications like email (Gmail, Yahoo mail, etc.), Social Networking sites (Facebook, etc.)
Popularity	Highly skilled developers, researchers who require custom configuration as per their requirement or field of research.	Most popular among developers as they can directly focus on the development of their possibly complex apps or scripts.	Most popular among normal consumers or companies which rely on software such as email, file sharing, social networking as they don't have to worry about the technicalities.

See https://goo.gl/ZwZtMQ

A foundation block: AAI

- AAI = Authentication and Authorization Infrastructure

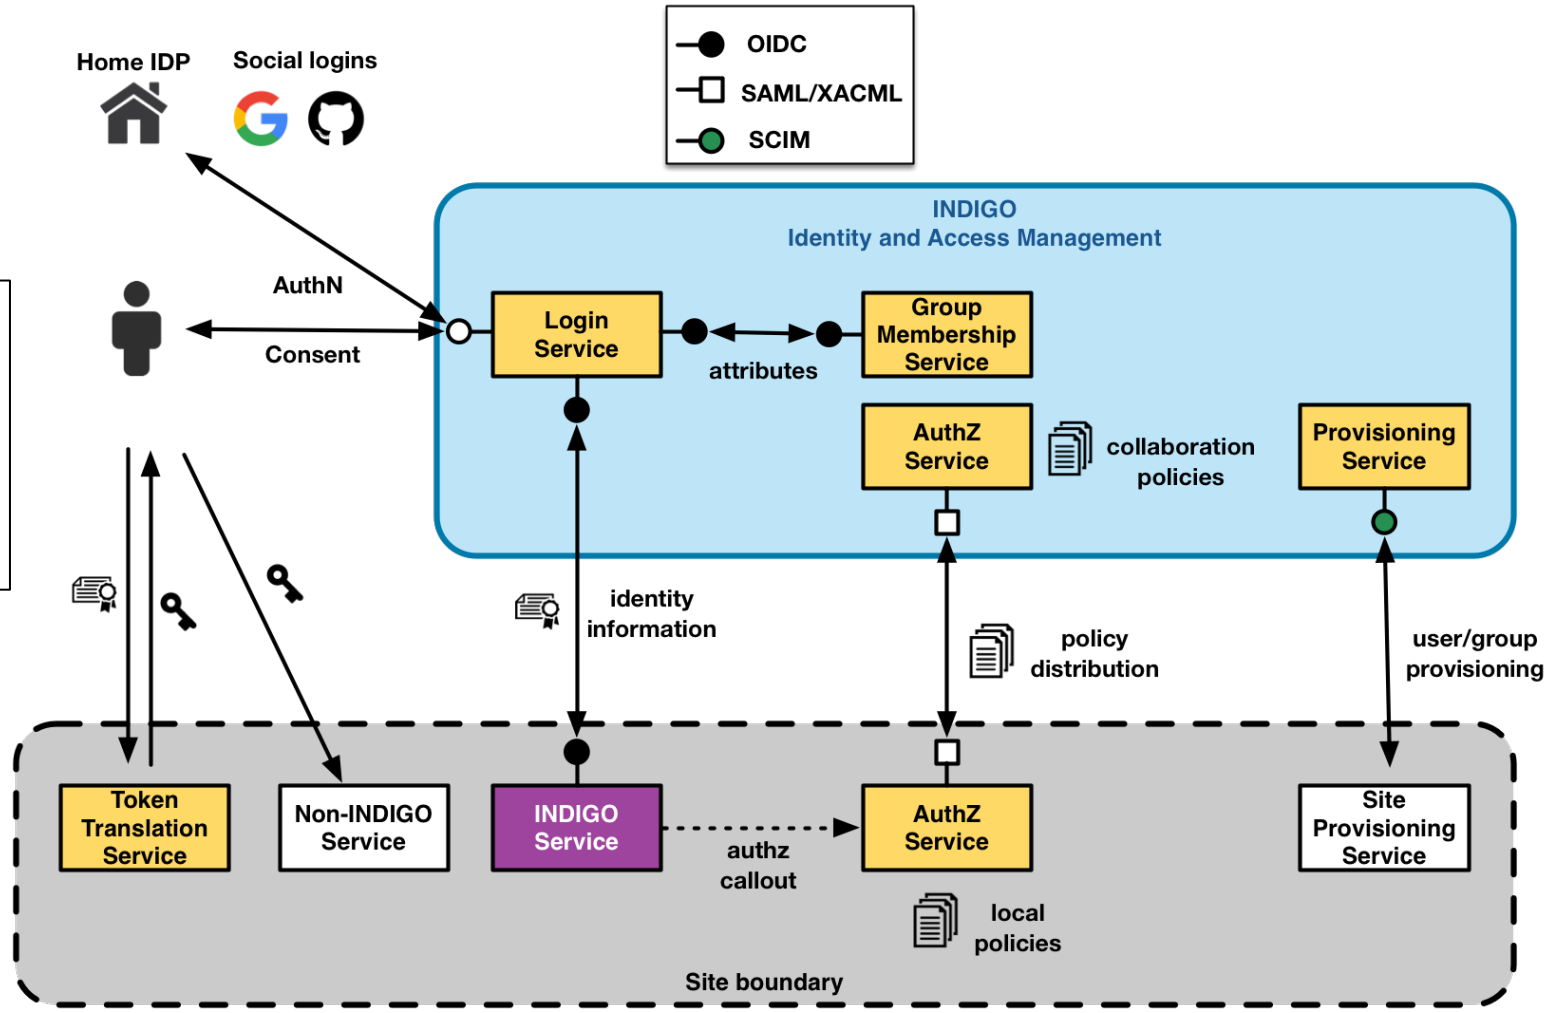

AAI example: the INDIGO IAM

- ***User Authentication***
  - username/password, **X.509** certificates, **OpenID Connect**, **SAML**
- ***User registration and onboarding***
  - Vetted registration flow where users **request membership** in the organization and requests for membership are approved by administrators
  - **Automated registration** flow for on-demand account creation for user authenticated through **trusted SAML IdPs**
- ***Account linking***
  - Link X.509 certificates, SAML and OpenID connect accounts to the IAM account
  - SSH key management
- ***Provisioning***
  - **SCIM 2** compliant provisioning and management APIs
- ***Organisation management:***
  - User, **group** and other attributes **management**
- ***Delegation, OpenID Connect/OAuth client and token management***
  - Client registration
  - Token management
  - Token exchange

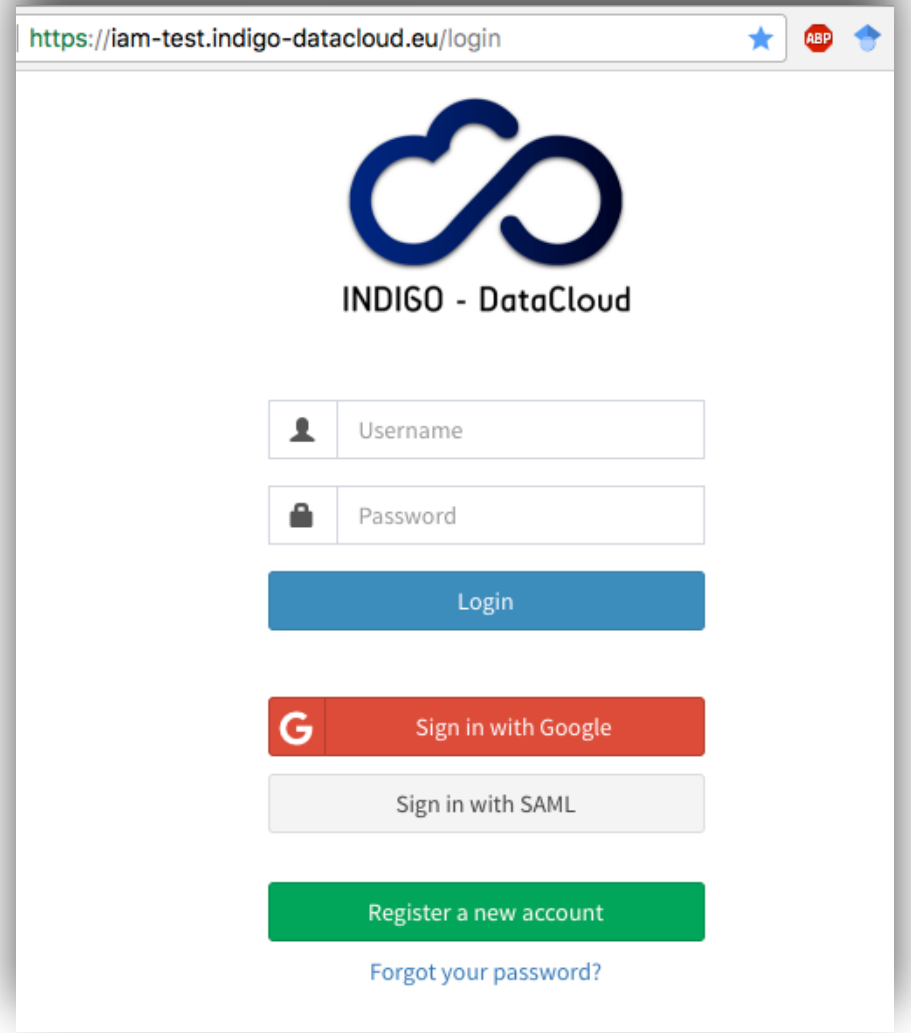

Harmonizing AAI

- Two main points:
  1. allow users from one AAI to be authenticated and authorized at resources managed with another AAI
  2. allow a site to support more than one organization possibly managed with different AAI systems

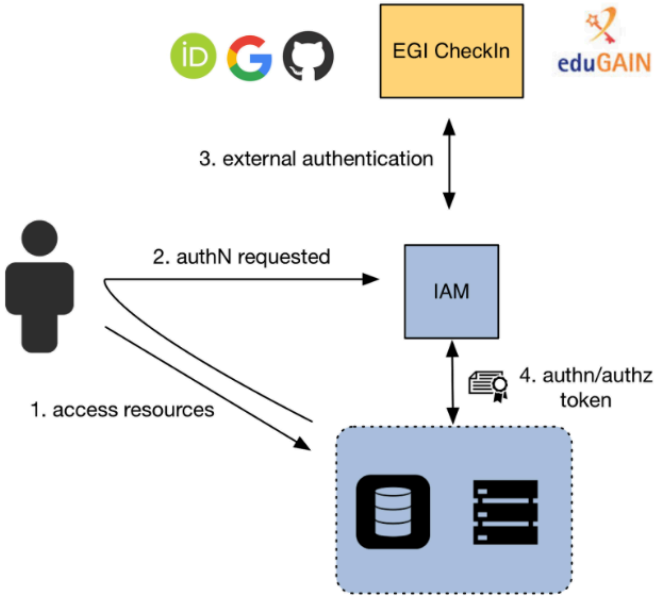

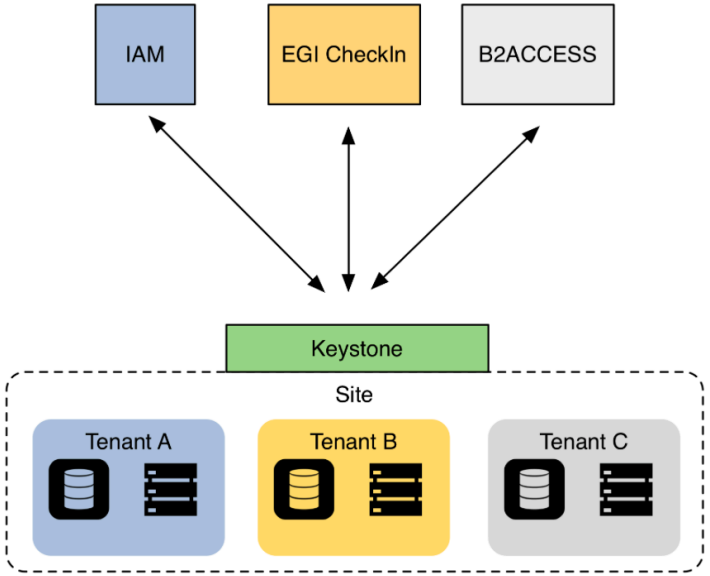

INDIGO WaTTS : the Token Translation Service

- WaTTS allows to authenticate to resources using any legacy service with federated identities, such as eduGain or Google.
- WaTTS accepts federated identities (via OpenID Connect) and uses a plugin scheme to generate credentials for your service. This extension mechanism allows you to provide services that do not normally support federated identities to federated users.
- Many plugins for common services already exist, for example:
  - SSH public/private keypair creation
  - S3 authentication tokens
  - OpenNebula credentials
  - X.509 certificates

Automation and service composition

- Very rarely we just need “bare bone VMs”. In the end, we want to run **applications**, not VMs.
- Applications may require several services to run, scale up or down, work with data, interact with the environment, etc.
- “In the Cloud” (e.g. in the EOSC), applications could be deployable over multiple cloud providers – public or private. We therefore also need an **orchestrator**.
- In INDIGO, we describe service composition and pass it to an orchestrator through a **templating language: TOSCA**
- In practice: **pick the services you really need** for your use-case, and **build your own platform**.

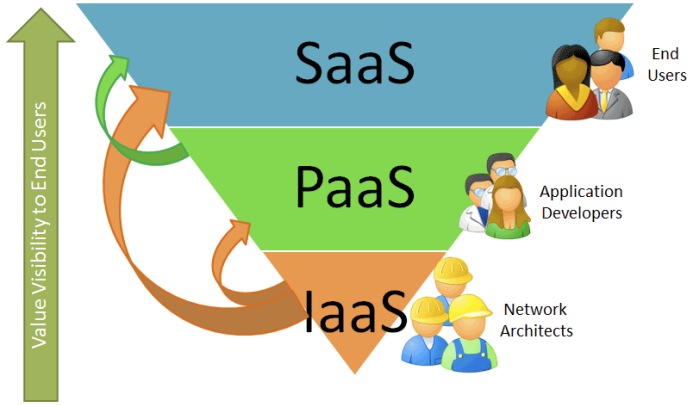

TOSCA, in short

- **T**opology and **O**rchestration **S**pecification for **C**loud **A**pplications
- Standardizes the language to describe:
  - *The structure* of an IT service (its **topology** model)
  - *How to orchestrate operational behavior* (plans such as build, deploy, patch, shutdown, etc.)
- It is a declarative model that spans applications, virtual and physical infrastructures.
- INDIGO supports TOSCA declarations at the IaaS, PaaS and SaaS levels to automatize the definition & instantiation of services.

Composing Services through the INDIGO PaaS (1)

- The INDIGO PaaS can automatize the provisioning of a complex cluster of services through a TOSCA template
- This allows end users to (relatively) easily specify dependencies and relationships between the services.
- It is possible to request and access simple IaaS resources (Virtual Machines or Docker containers), configuring services automatically using e.g. Ansible.
- However, and this is much more powerful, **we can also exploit the Mesos framework** to support long-running services and the monitored execution of applications in Docker containers.
- The INDIGO PaaS is also able to automatically mount virtual file system spaces (Onedata).

Composing Services through the INDIGO PaaS (2)

- **Apache Mesos** is used by INDIGO to provide:
  - auto-scaling of the allocated resources
  - monitoring of the long-running services
  - storage persistency for the long-running services
  - monitoring and re-execution in case of failure of an application running on a Docker container
  - dependency of “jobs” both in Marathon (long-running services) and in Chronos (execution of applications in Docker)
- Through the INDIGO PaaS layer, a user can submit a TOSCA template that describes an arbitrarily long list of application execution dependencies
  - The PaaS layer then submits the requested “jobs” to Chronos, that manages execution and monitoring.

Use case: a web portal using a batch system to run applications

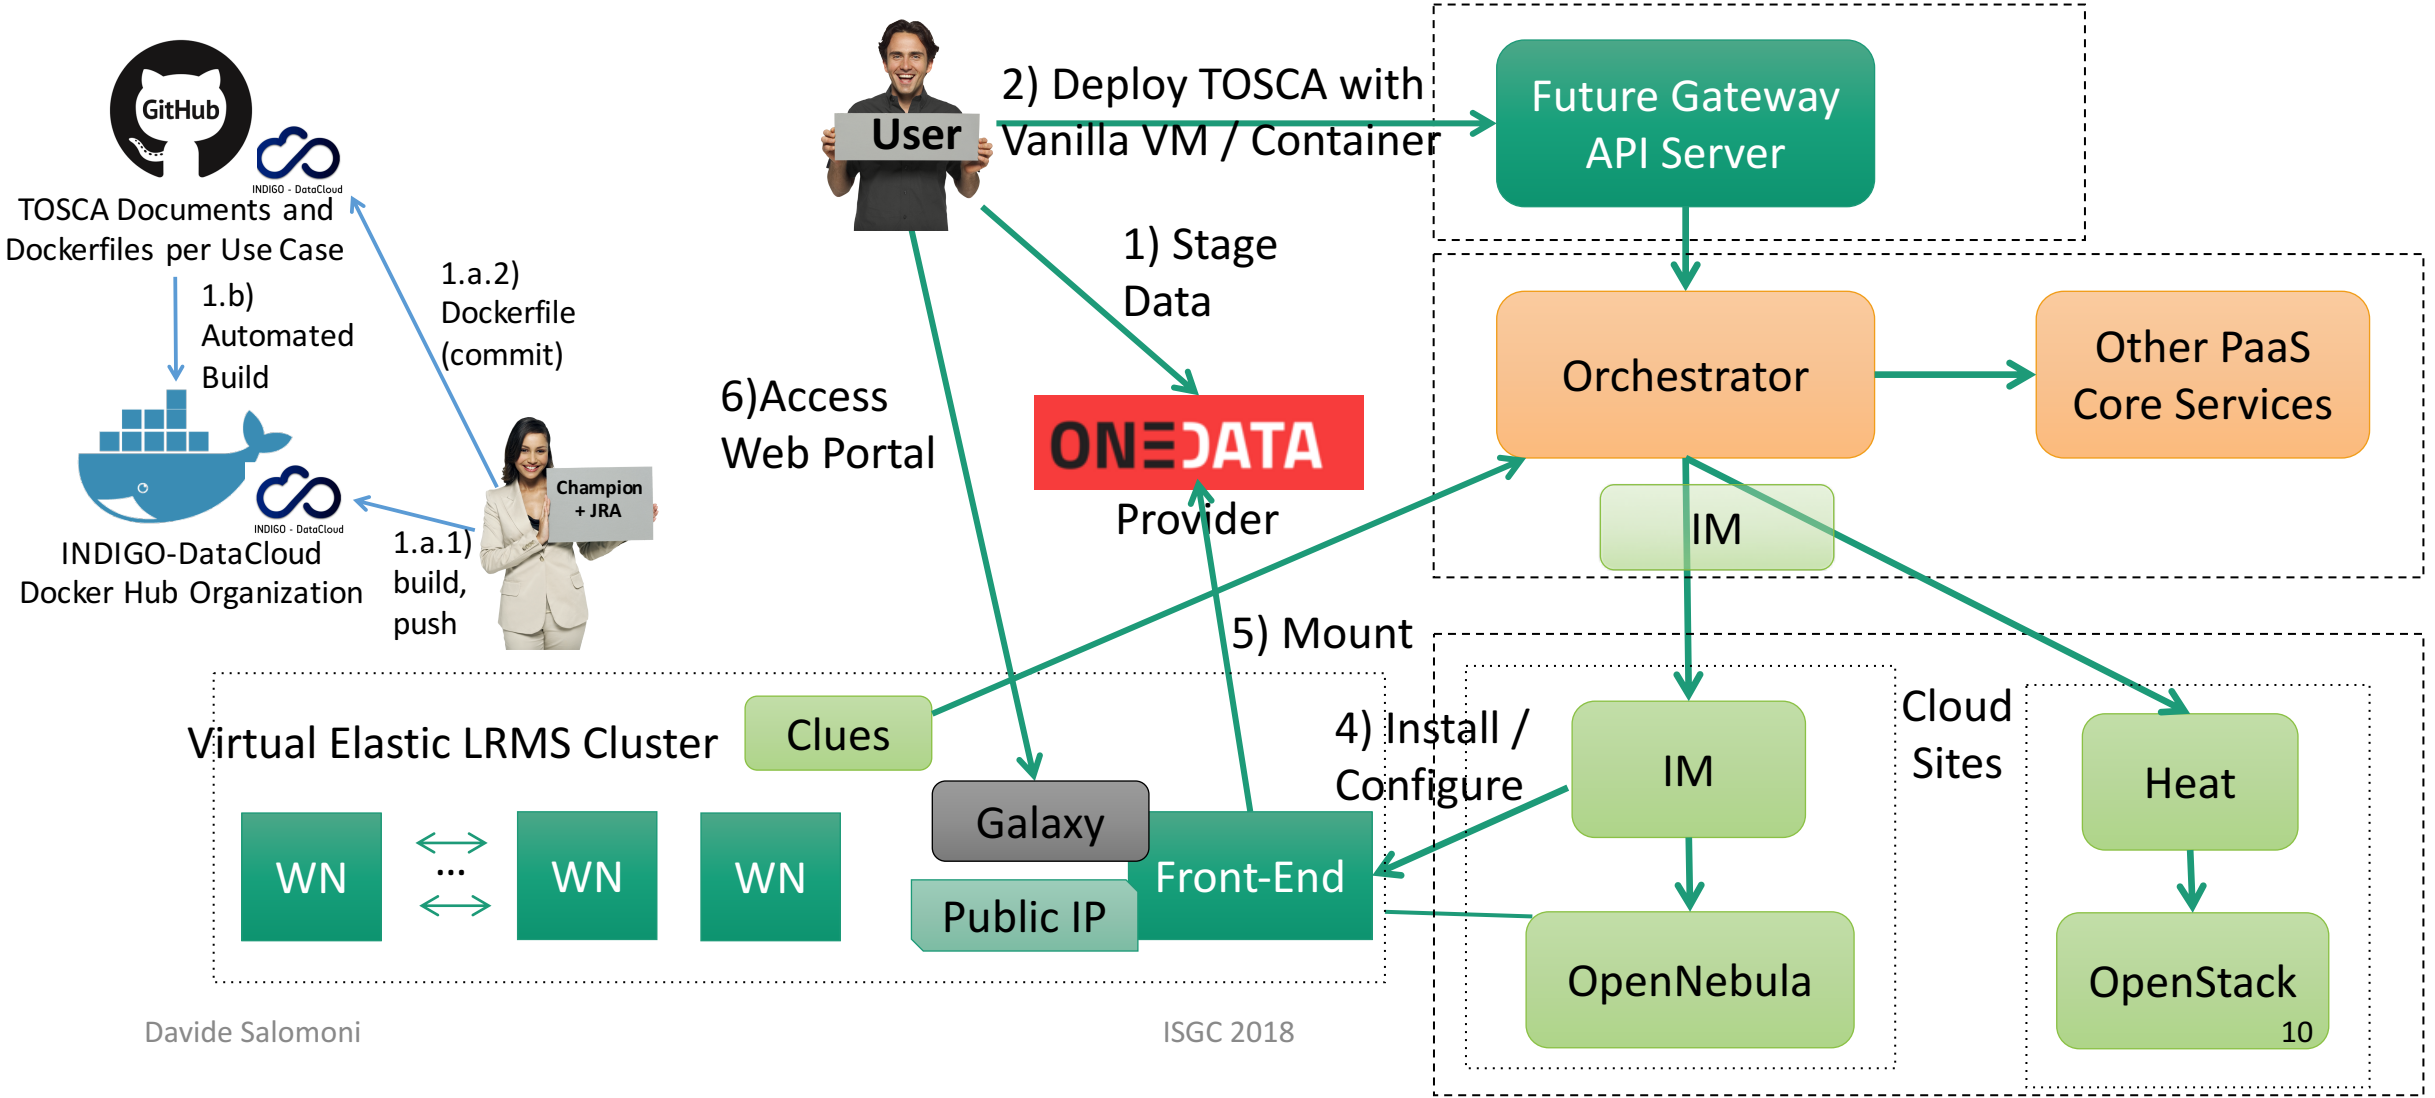

Big Data processing in the Cloud

- With the PaaS technology shown above, we can e.g. **instantiate on-demand Spark clusters on hybrid Clouds**.
- However, we also want to support **computing techniques for the analysis of very large datasets, exploiting specialized hardware components:**
  - GPUs…
  - … low-latency interconnects…
  - … and other resources usually accessed as “bare metal”.
- Also ensuring interoperability with the existing EOSC platforms and their services.
- These topics are being directly tackled by the **DEEP-Hybrid DataCloud project**
  - An INDIGO-DataCloud follow-on project, started in Nov 2017.

DEEP-Hybrid DataCloud

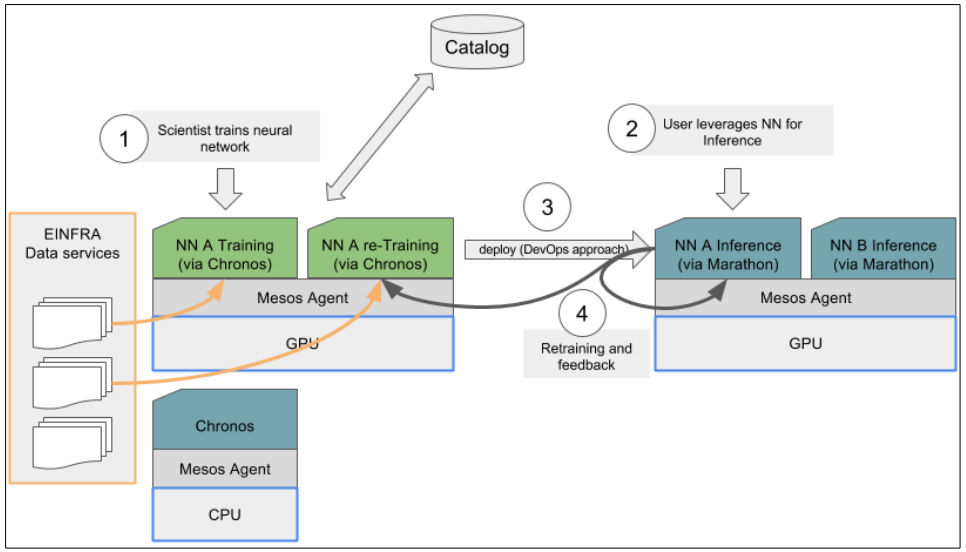

Figure 3: Initial architecture proposed to support a deep learning application using private cloud services.

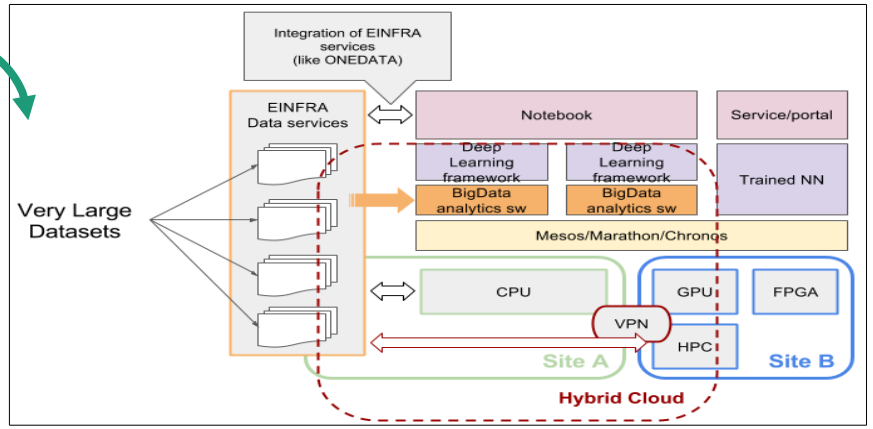

Figure 4: Initial architecture using a Hybrid Cloud approach.

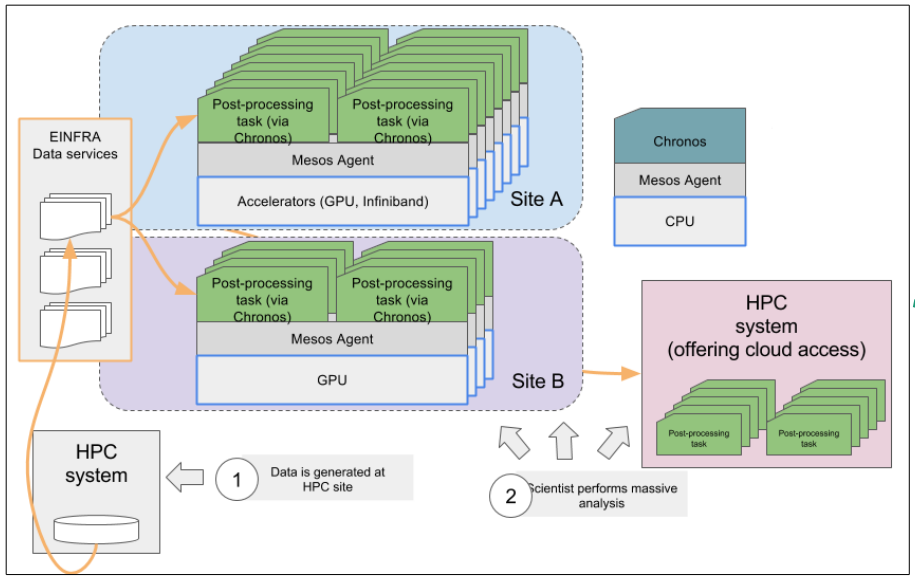

Figure 5: Massive post-processing of very large datasets using containers.

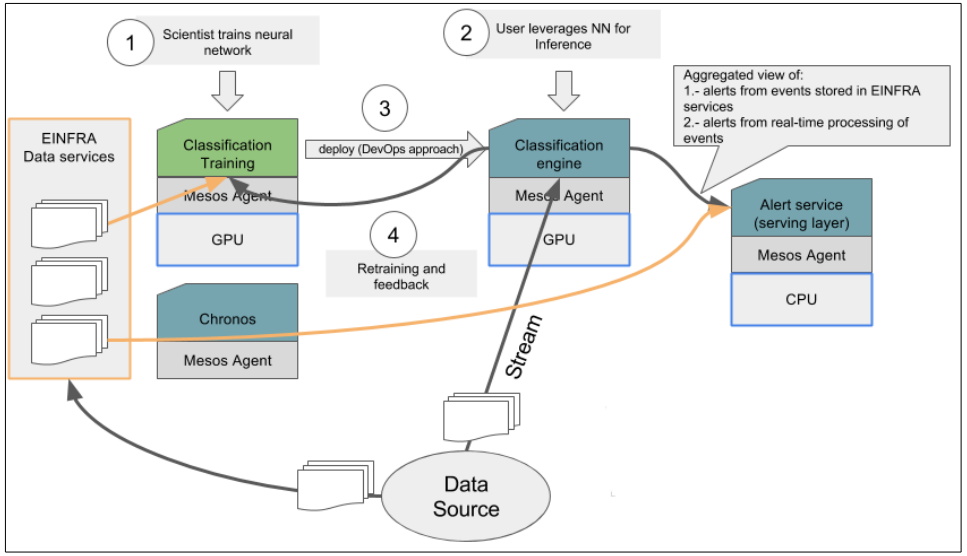

Figure 6: Analysis of massive real-time data streams following a lambda architecture.

Docker everywhere?

Pros

- Easy software distribution
- All dependencies bundled together
- Isolation
- Immutable images
- Distribution independence
- "All the cool kids use containers, why are we stuck in the '80s?"

Cons

- Special kernel facilities required
- Recent kernel and distributions required for optimal support
- Complex setup for efficient storage back-ends
- Cumbersome access to shared storage
- Daemons required to run containers and perform bookkeeping
- Access to the container daemon means root access to the host
- "Can't you just rewrite your tools to work in our environment?”
- How about HPC?

Emulation: udocker

- **Udocker**, developed by the INDIGO project, is a "userland docker": a tool to run contents of Docker images without special kernel facilities
- No special dependencies (aside from Python and libc)
- Used by unprivileged users
- No special daemon required
- Can select several execution modes (PRoot, Fakechroot, runC and Singularity, if installed on the host machine).
- Available today: https://indigo-dc.gitbooks.io/indigo-datacloud-releases/content/indigo2/udocker2.html
- **Can be used to run containers on most infrastructures** (local systems, Grids, Clouds, and HPC included).

Evolutions in Data Management

- Especially in the HEP world, there has been a recurring reference to the concept of **Data Lake**.
  - I.e., a seamless, worldwide interconnection of large data and compute centers.
- Not really a new concept, but now made possible by the evolution of technology

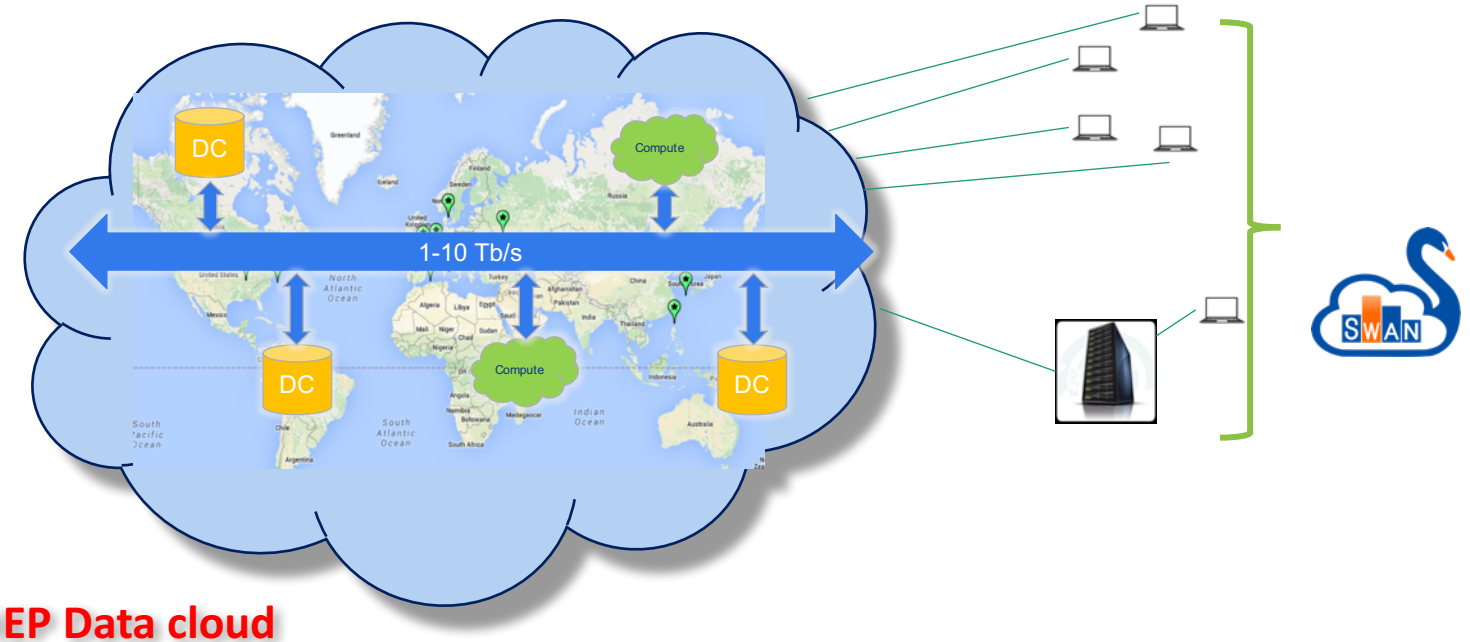

HEP Data cloud

Storage and compute

A flashback from a while ago…

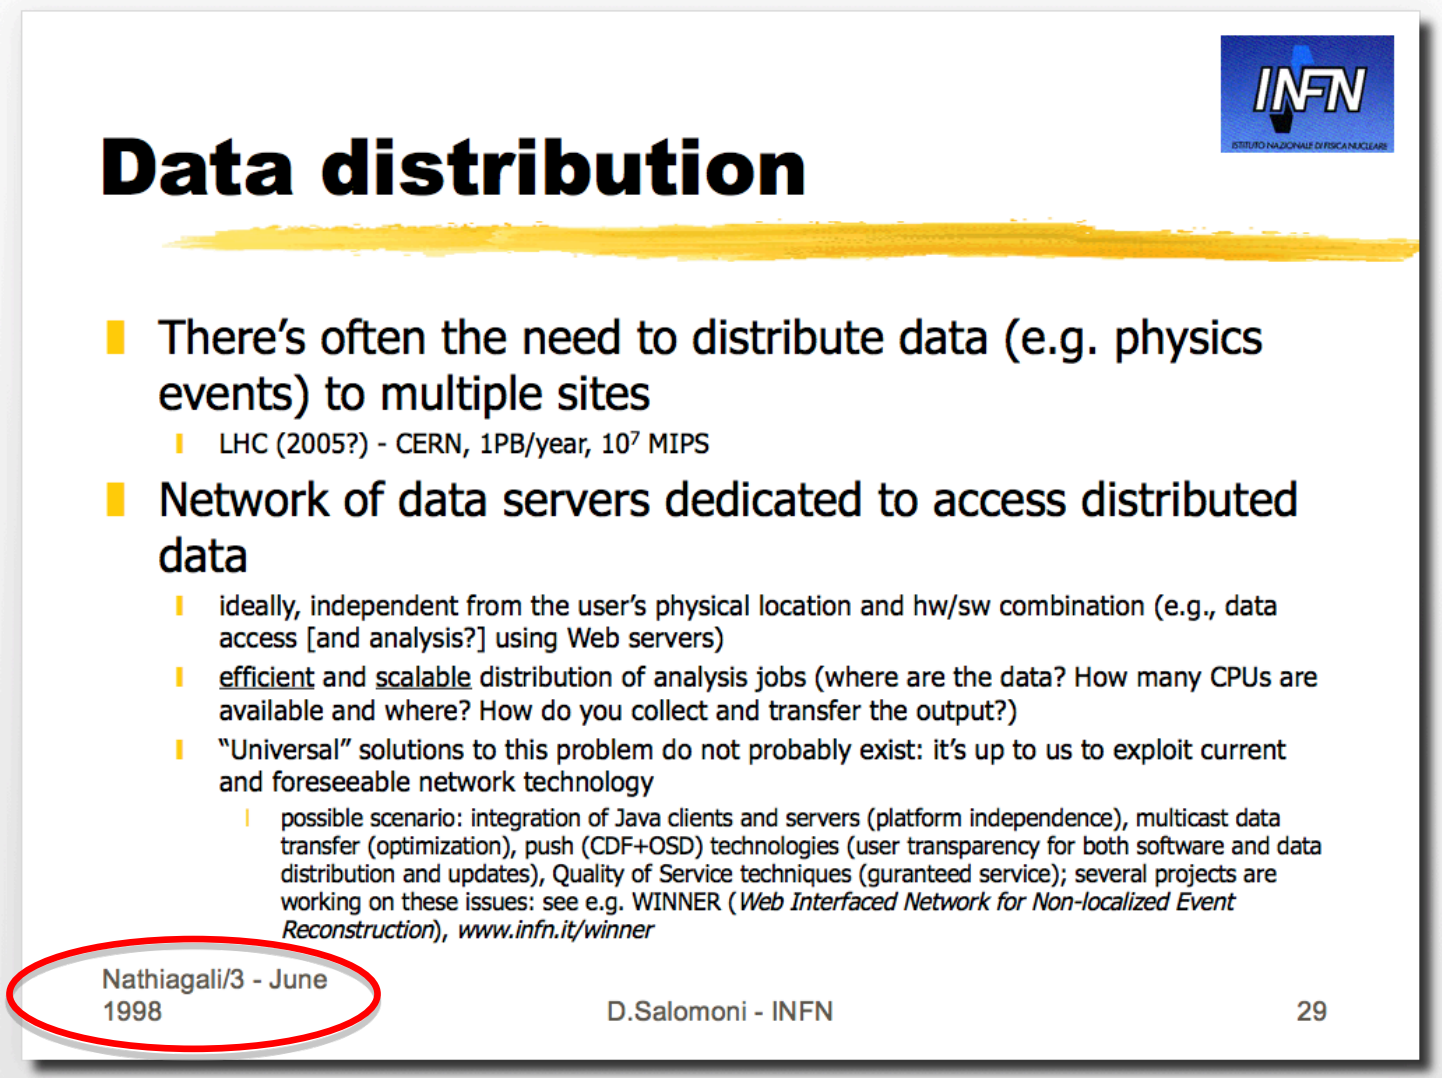

Data distribution

- There's often the need to distribute data (e.g. physics events) to multiple sites
  - LHC (2005?) - CERN, 1PB/year, 10^7 MIPS
- Network of data servers dedicated to access distributed data
  - ideally, independent from the user's physical location and hw/sw combination (e.g., data access [and analysis?] using Web servers)
  - efficient and scalable distribution of analysis jobs (where are the data? How many CPUs are available and where? How do you collect and transfer the output?)
  - "Universal" solutions to this problem do not probably exist: it's up to us to exploit current and foreseeable network technology
    - possible scenario: integration of Java clients and servers (platform independence), multicast data transfer (optimization), push (CDF+OSD) technologies (user transparency for both software and data distribution and updates), Quality of Service techniques (guranteed service); several projects are working on these issues: see e.g. WINNER (*Web Interfaced Network for Non-localized Event Reconstruction*), *www.infn.it/winner*

Nathiagali/3 - June 1998 D.Salomoni - INFN 29

What next then?

- **Policy-driven data management**, with storage QoS and intelligent & automated dataset distribution
- **Data pre-processing** at ingestion time
- **Smart caching** algorithms and tools
- **PaaS-level unified data access** with encryption services and metadata management
- These topics are being directly tackled by the **eXtreme-DataCloud project**
  - An INDIGO-DataCloud follow-on project, started in Dec 2017

eXtreme-DataCloud (XDC)

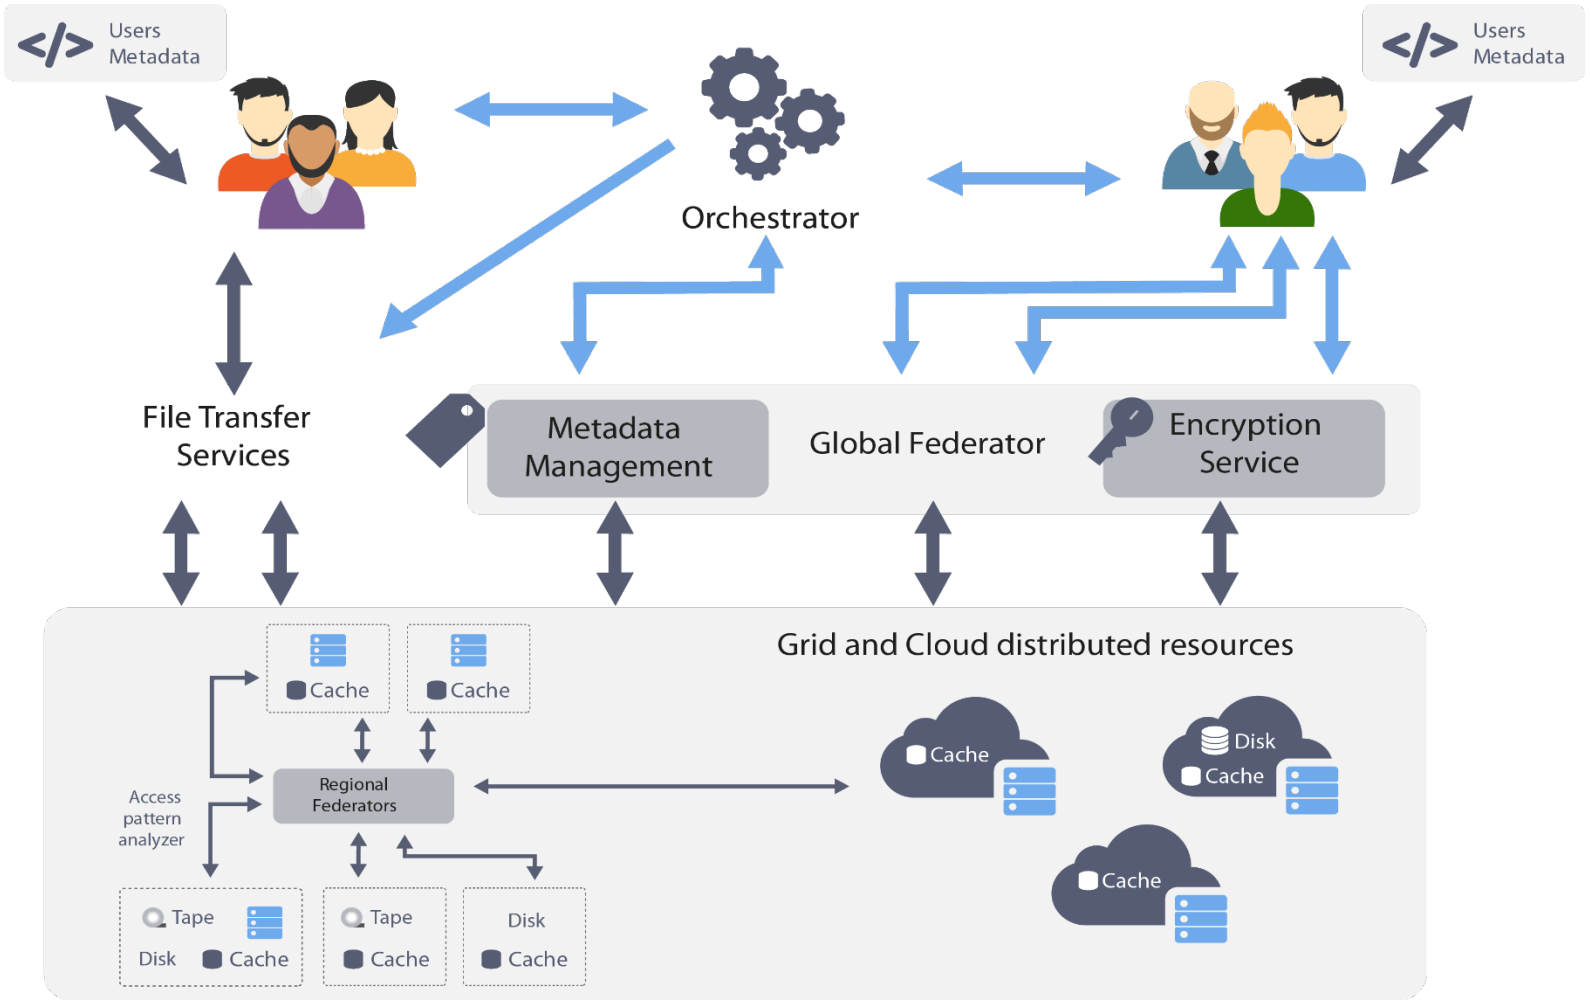

Conclusions

- Europe has been investing a relatively large amount of money to create a **European Open Science Cloud (EOSC)**.
- Not an easy process! But several concrete steps are being taken, for example through the EOSC-hub project, through initiatives such as INDIGO, DEEP, XDC and through several upcoming calls.
- **Cloud "lakes" are here to stay**, possibly with public/private partnerships.
- **We still need to work** on simplifying access, moving up the Cloud stack (from IaaS to PaaS and SaaS), and hide technology as much as we can, **focusing on applications, on open, transparent and scalable composition of services, and on the ingestion, analysis and integration of distributed big data via Cloud and HPC resources**.
  - Topics such as deep learning, policy-driven data management and QoS, unified data & metadata management are all over the place. Which is good.

Related talks at ISGC 2018

- **Automated service composition for batch and big data processing via the Dynamic On-Demand Analysis Service (DODAS)**, by Daniele Spiga
  - March 20, 11:30-12:00, Media Conference Room
- **Next Generation Data Management Services: the eXtreme DataCloud project**, by Daniele Cesini
  - March 20, 14:00-14:30, Conference Room 1
- **Smart Policy Driven Data Management and Data Federations, enabled by the H2020 eXtreme-DataCloud Project**, by Patrick Fuhrmann
  - March 20, 14:30-15:00, Conference Room 1

Acknowledgments

- Thanks especially to Luciano Gaido, Isabel Campos, Jesus Marco, Ludek Matyska, Jorge Gomes, Tiziana Ferrari, Patrick Fuhrmann, Marcus Hardt, Giacinto Donvito, Lukasz Dutka, Marcin Plociennik, Roberto Barbera, Andrea Ceccanti, Daniele Spiga, Cristina Duma, Marco Tangaro, Marica Antonacci, Daniele Cesini, Alvaro Lopez.
- And thanks collectively to the **INDIGO-DataCloud** (https://www.indigo-datacloud.eu), **EOSC-hub** (https://www.eosc-hub.eu), **eXtreme-DataCloud** (https://www.extreme-datacloud.eu) and **DEEP-Hybrid Datacloud** (https://www.deep-hybrid-datacloud.eu) projects, as well as to the **EC** for its support.

Thank you

“Stat rosa pristina nomine, nomina nuda tenemus”