INDIGO - DataCloud

RIA-653549

INDIGO-DataCloud Outcomes

Better Software for Better Science.

Davide Salomoni – INFN/CNAF

INDIGO-DataCloud Project Coordinator

davide@infn.it

GDB, 13/12/2017

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

Outline

- INDIGO-DataCloud
- Selected examples
- Results and Impact

CAVEAT: not many technical details here!

See https://www.indigo-datacloud.eu for more information

The INDIGO-DataCloud Project

- **An H2020 project** approved in January 2015 in the EINFRA-1-2014 call
  - 11.1M€, 30 months (**from April 2015 to September 2017**)
- **Who**: **26 European partners** in 11 European countries
  - Coordination by the Italian National Institute for Nuclear Physics (INFN)
- **What**: **develop an open source Cloud platform** for computing and data ("DataCloud") tailored to science but applicable to other domains as well.
- **For**: **multi-disciplinary scientific communities**
  - E.g. structural biology, earth science, physics, bioinformatics, cultural heritage, astrophysics, life science, climatology
- **Where**: deployable on **hybrid (public or private) Cloud infrastructures**
  - INDIGO = **IN**tegrating **D**istributed data **I**nfrastructures for **G**lobal Expl**O**itation
- **Why**: answer to the technological **needs of scientists** seeking to easily exploit distributed Cloud/Grid compute and data resources.

The INDIGO-DataCloud Consortium Members

- Software developers
- Industrial partners
- Research institutes
- Universities
- e-infrastructures
- Scientific communities

INFN Istituto Nazionale di Fisica Nucleare
Universiteit Utrecht
CESNET
CSIC Consejo Superior de Investigaciones Científicas
PSNC
cea energie atomique • energies alternatives
cmcc Centro Euro-Mediterraneo sui Cambiamenti Climatici
AGH University of Science and Technology
Reply
ICCU
DESY
CERM Firenze
egi
KIT Karlsruher Institut für Technologie
Science & Technology Facilities Council
RB
cnrs
LIP Laboratório de Instrumentação e Física Experimental de Particulas
T-Systems
CERN
Consiglio Nazionale delle Ricerche
indra
AtoS
Istituto Nazionale di Geofisica e Vulcanologia
INAF Istituto Nazionale di Astrofisica National Institute for Astrophysics
Universitat Politècnica de València

The INDIGO-DataCloud overall objective

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

“Better Software for Better Science.”

Four general objectives

INDIGO - DataCloud

1. **Develop a Platform** (PaaS) based on open source software, without restrictions on the e-Infrastructures to be accessed (public or commercial, GRID/Cloud/HPC) or its underlying software.
2. Provide the **interface between e-Infrastructures and Platforms**, reducing also the performance impact generated by virtualization.
3. Provide **high-level access** to the platform services **in the form of science gateways, access libraries, mobile applications.**
4. Streamline the **adoption of the software products**.

INDIGO-DataCloud's Vision

- INDIGO:
  1. **Develops open, interoperable solutions for scientific data.**
  2. **Addresses requirements of heterogeneous scientific communities and providers.**
     - For distributed data and compute services.
  3. **Enables collaborations across diverse scientific communities worldwide.**

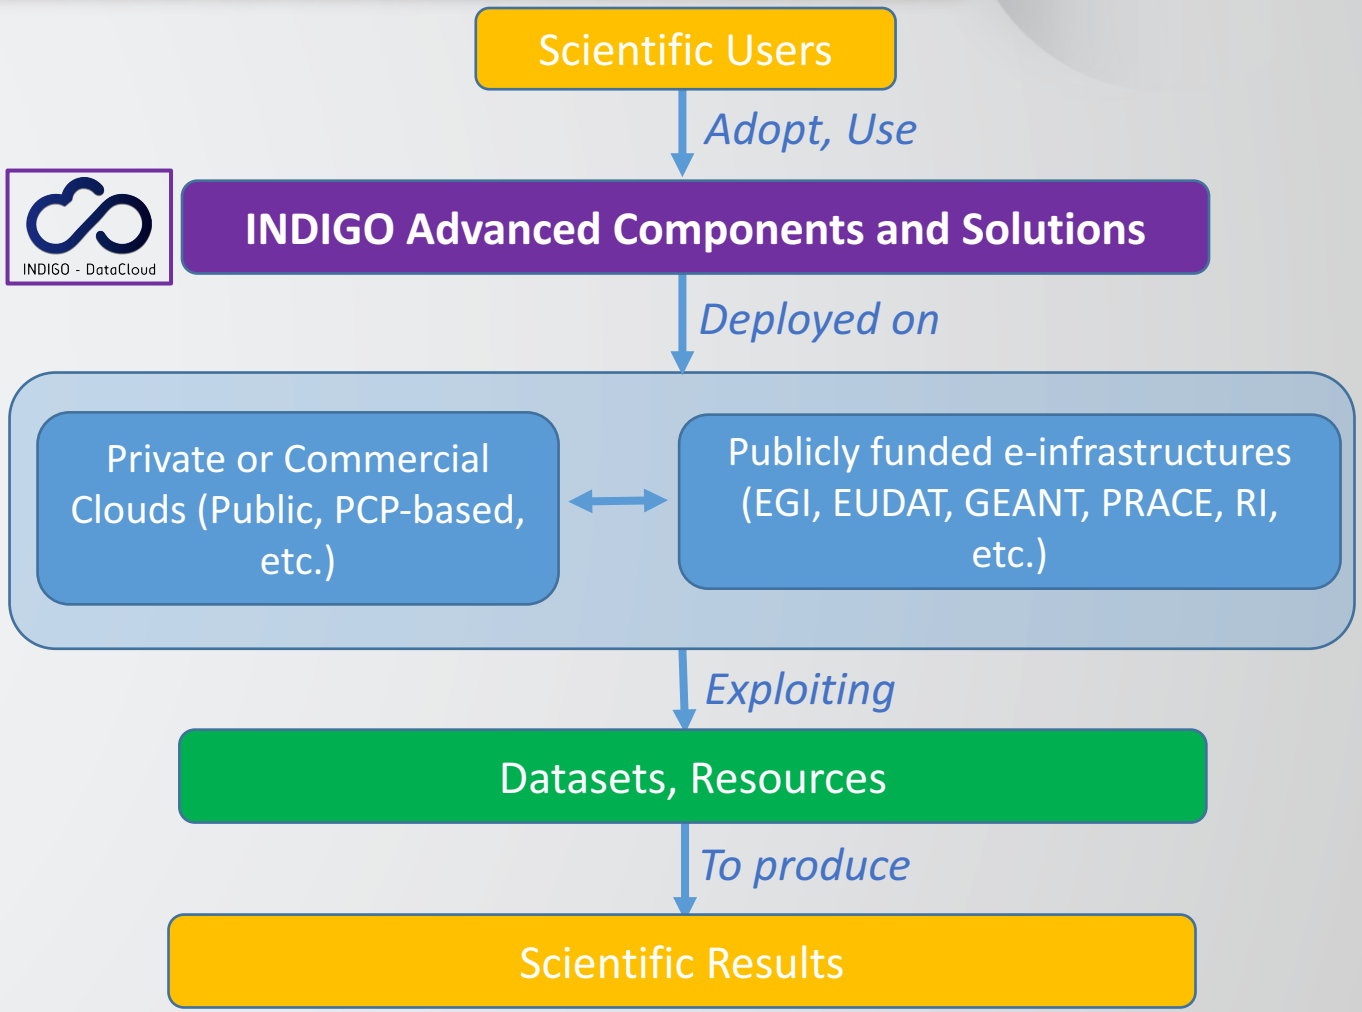

INDIGO-DataCloud's Approach

- We recognize that **value for users** (and hence, our main focus) is at the **upper layers**, not in the bare bone e-infrastructural services.
  - But we also provide ways to optimize e-infrastructural services for resource providers
- So, we abstract from underlying IaaS technologies and offer flexibility in choosing e-infra providers, resources and capabilities…
- … **giving users the possibility to easily express and implement requirements for their applications through enabling services and components.**

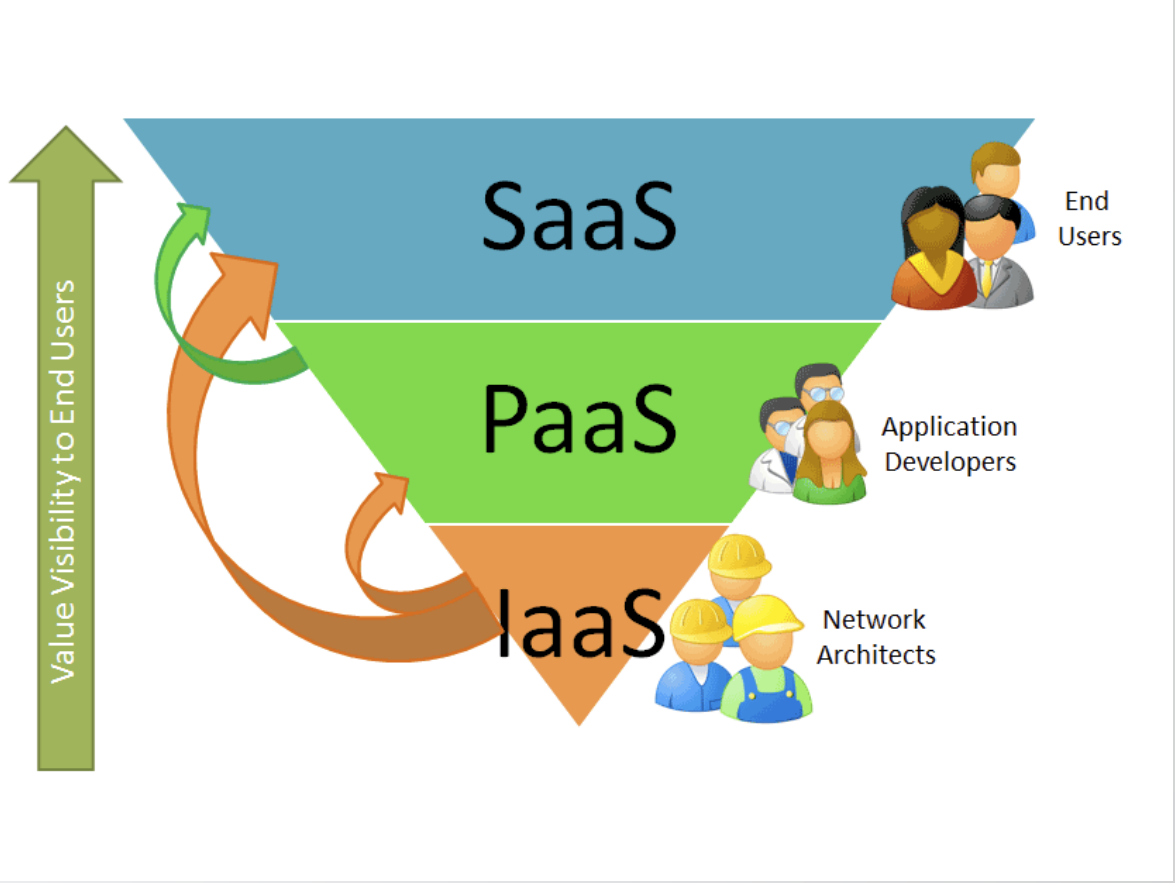

The INDIGO Software

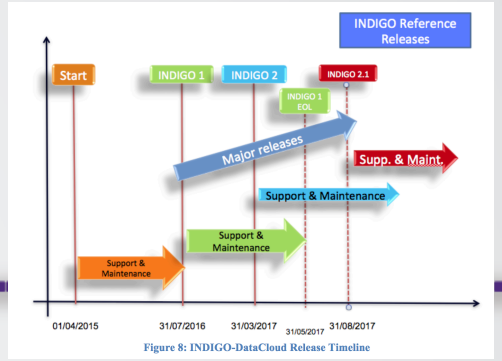

Figure 8: INDIGO-DataCloud Release Timeline

INDIGO - DataCloud

- We released two major software versions. Our 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

INDIGO - DataCloud

ElectricIndigo is described in:

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

The ElectricIndigo Release

INDIGO - DataCloud

- 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**

High-level view of the INDIGO architecture

This is the INDIGO-DataCloud General Architecture*

*: see details in http://arxiv.org/abs/1603.09536 or in https://www.indigo-datacloud.eu/documents-deliverables

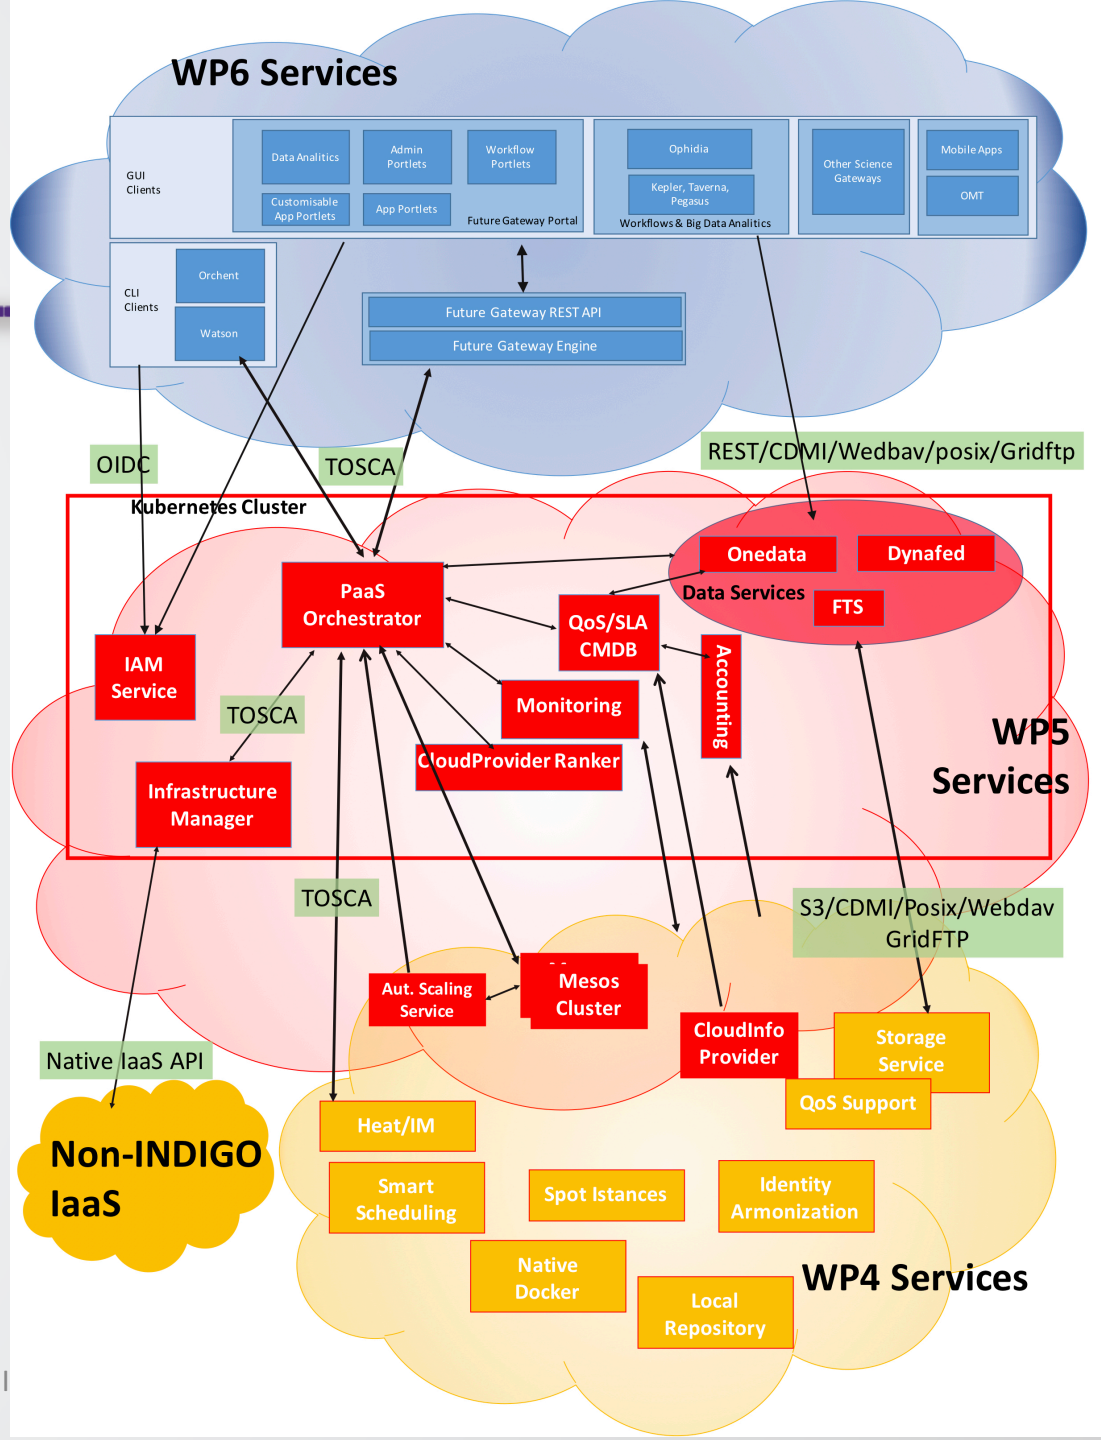

Selected examples

Some use case examples of production-level adoption of INDIGO components

- **LifeWatch**: monitor the evolution of the potential eutrophication of a water reservoir including data lifecycle management, hydrodynamic and water quality models for forecasting.
- **CMS/DODAS**: on-demand instantiation of Condor-based clusters for HEP.
- **ENES**: a worldwide-distributed multi-model analysis for climate modeling.
- **Structural biology**: virtualization of the DISVIS and POWERFIT tools with all their dependencies for execution on Clouds and Grids & GPU exploitation
- **Phenomenology**: running Mastercode (which glues together a lot of software, e.g. SoftSUSY, FeynHiggs, HiggsBound, HiggsSignal, SuFla, SuperIso, FeynWZ, DarkSUSY, Atom, Scorpion, SDECAY) on HPC or Grids
- **EuroBioImaging**: a self-contained elastic cloud platform to run image-processing components on POSIX-accessible medical image files.
- **Elixir**: a Galaxy provider platform, allowing full customization of virtual instances through a user-friendly web interface, ready to be used by life scientists and bioinformaticians.
- **DARIAH**: automated deployment of INVENIO-based data repositories in the cloud, with resource scaling and management.
- **Large Binocular Telescope (LBT)**: distributed archive with automated replication, data analysis, data and metadata management

Generic applicability

- The wide variety of use cases demonstrated by many diverse scientific communities (in INDIGO and elsewhere) show that **the INDIGO modular components are not linked to specific scenarios**.
- We have also provided **several "generic examples"** :
  - build an application encapsulated as a container running in an HPC system, using *udocker*
  - launch a containerized app in the Cloud using a web interface, using the *FutureGateway* and the *PaaS Orchestrator*
  - do the same, but through the command line, using *Orchent* and the *PaaS Orchestrator*
  - run a big-data SaaS application using *Ophidia*.
  - use INDIGO AAI to interface to existing apps

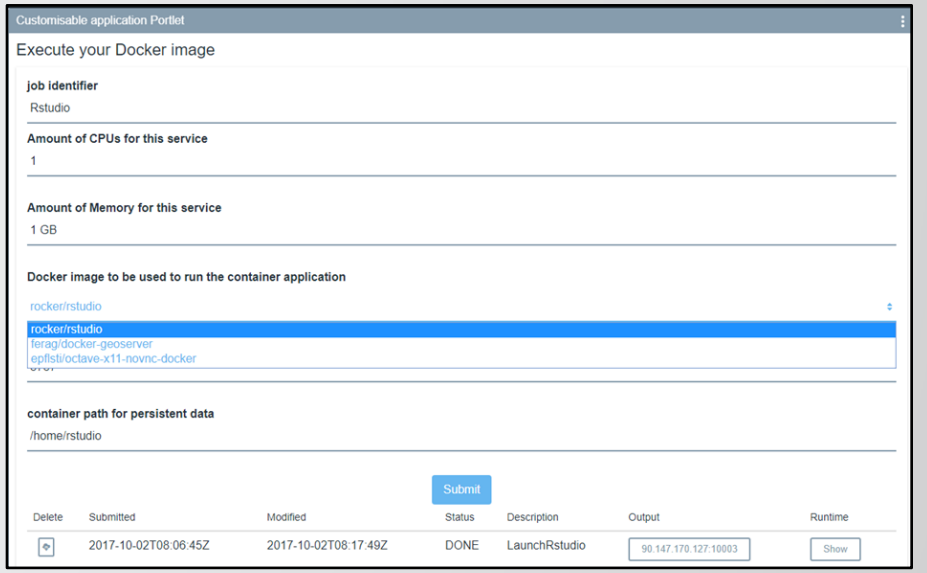

More complex solutions: a Virtual Elastic Cluster supporting a data intensive system

INDIGO - DataCloud

- Integrating several INDIGO-DataCloud technologies to automatically deploy a ready-to-use Galaxy production environment

M. Tangaro, CNR/Elixir-ITA

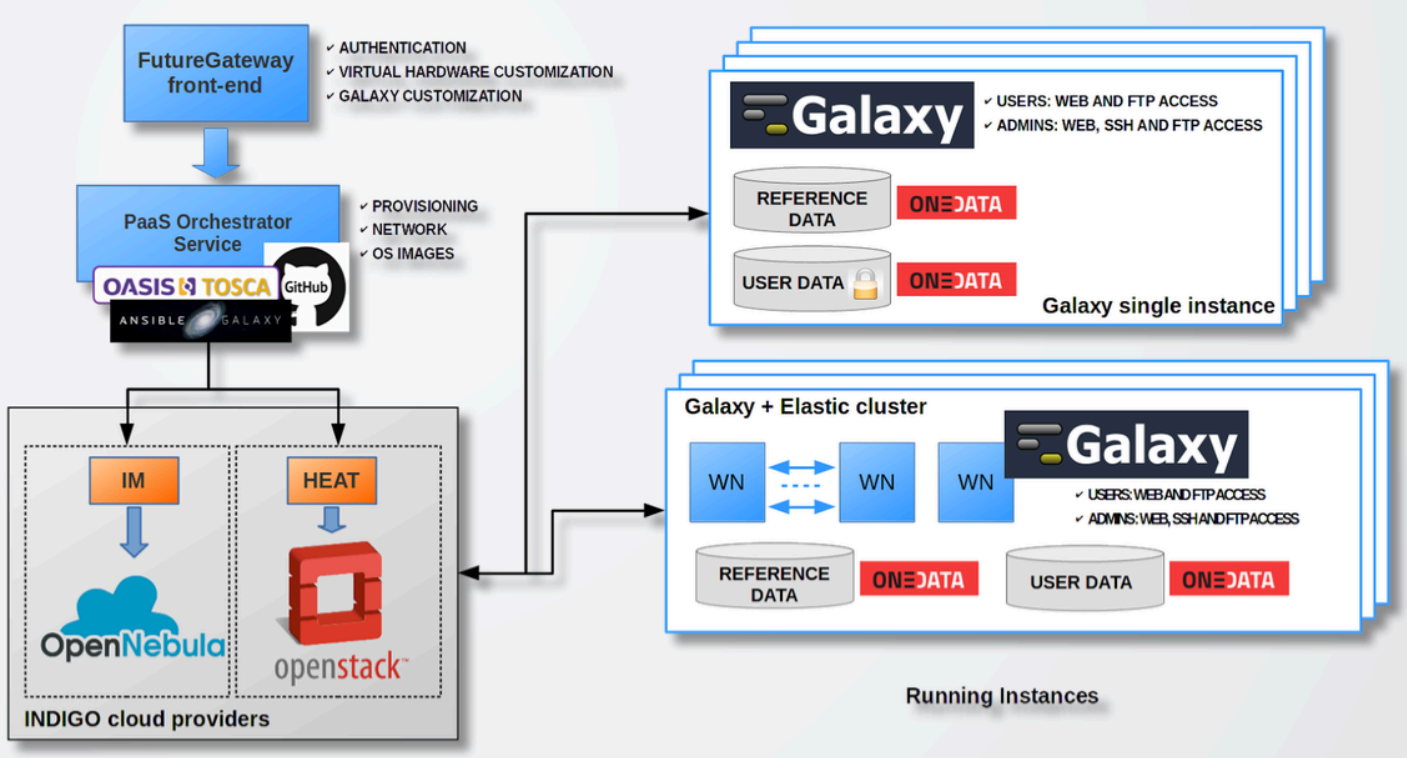

- Instance and Galaxy customization web front-end (FutureGateway).
- Full Galaxy production environment deployment (Orchestrator and Infrastructure Manager).
- Persistent storage (Onedata or IaaS block storage with/without LUKS filesystem encryption)
- Tools from ToolShed (Orchestrator and IM) with Conda support.
- Reference data availability (CernVM-FS, Onedata).
- Automatic elasticity (CLUES).
- Authentication service (FutureGateway and IAM).

The INDIGO AAI

- Provides the INDIGO **Identity and Access Management** (IAM) Service, responsible for
  - user **authentication**
  - **authorization**
  - provisioning

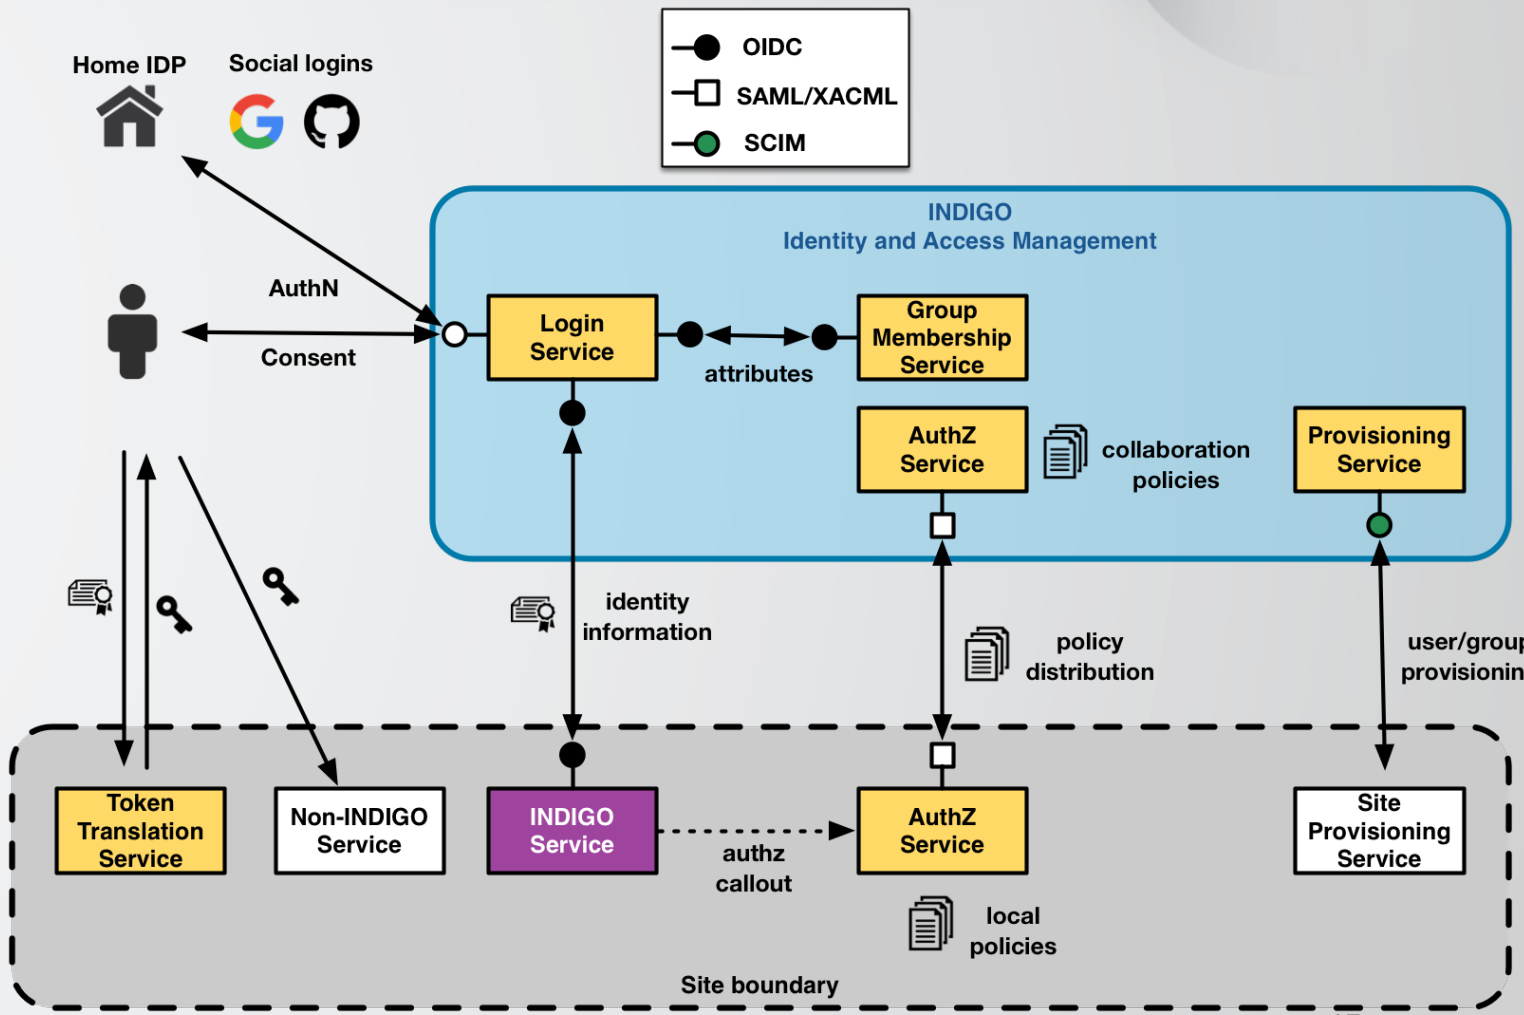

The INDIGO IAM

- ***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

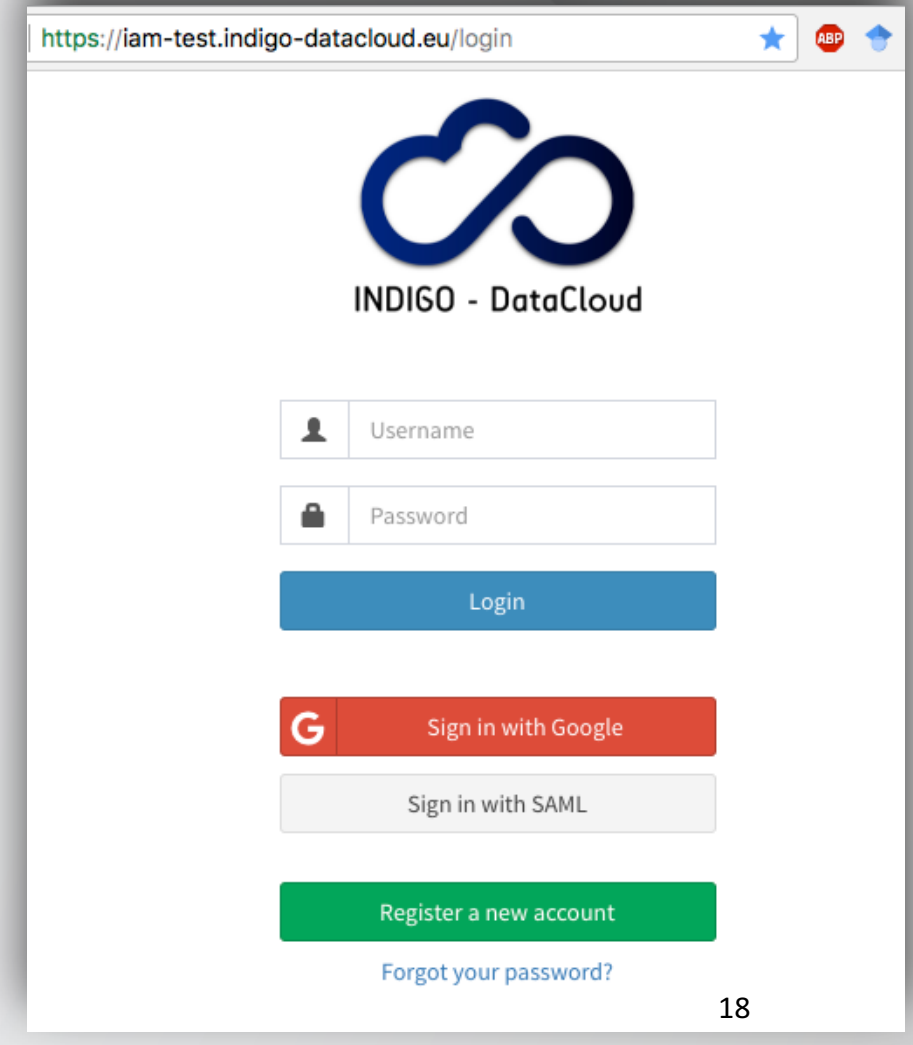

IAM in use

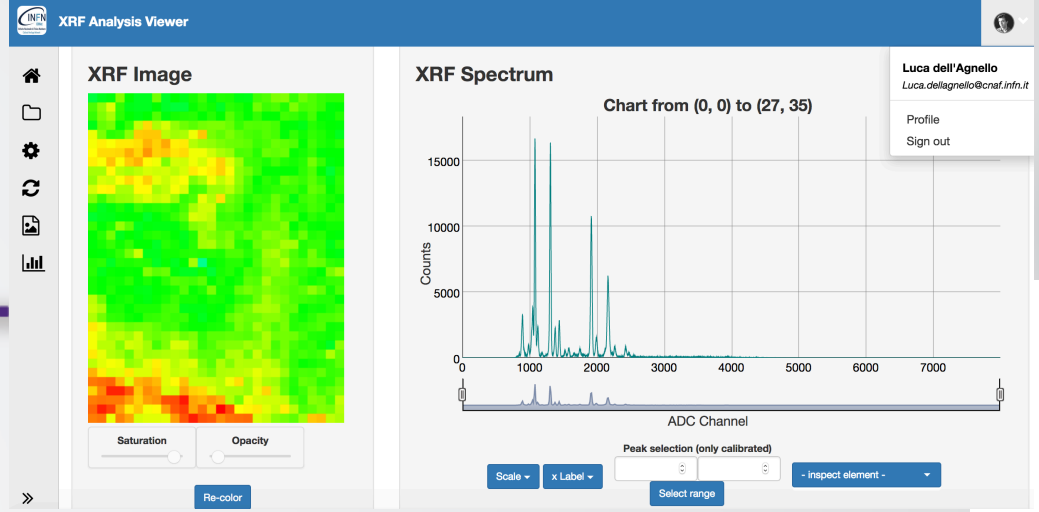

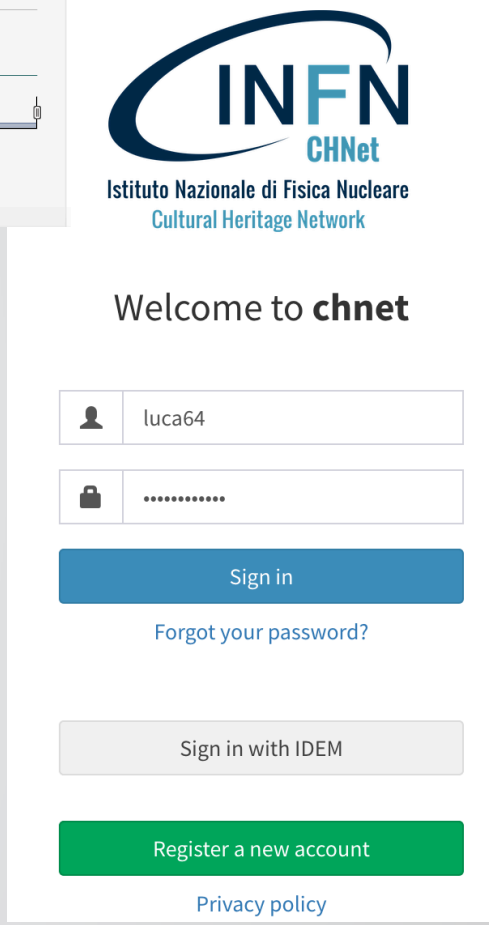

- **Federation of all types of access to resources** under a single system, in centers typically offering local access (username/password e.g. for user interfaces or batch systems), LDAP/Kerberbos-based authentication and autorization, Cloud access
- Authentication of users / groups using social logins and EduGain **plugged into standard OpenStack installations**
- Use of IAM into applications that have originally nothing to do with Clouds, such as the **Cultural Heritage Network (CHnet)**

INFN CHNet Istituto Nazionale di Fisica Nucleare Cultural Heritage Network

E-RIHS EUROPEAN RESEARCH INFRASTRUCTURE FOR HERITAGE SCIENCE

WaTTS : The Token Translation Service

Allow **any legacy service** to authenticate using IAM

- Web Frontend
- Support services via (extensible) shell scripts:
  - e.g. ssh, X.509, S3, mysql, ...

- Login via **INDIGO-IAM**, Google, Human-Brain, EGI-Checkin, EUDAT-b2access
- Official certification of OpenID-Connect standardization body
- Implementation is there: **WaTTS**
  - Production and development instances working
  - B2access and EGI froze plans after demos on WaTTS
    - **Both will evaluate WaTTS before pursuing own plans**
  - Extension to provide drop-in-replacement functionality added (Master-Portal)
  - Used in INDIGO:
    - CMS-group
    - OpenNEbula (OpenStack supported through keystone)
  - Used externally:
    - Masterportal for RCauth.eu
    - S3 tokens for HBP
    - Will be used in German Helmholtz Data Federation

Heat translator as a service

- TOSCA: The Orchestration Standard.
  - "tosca-parser" and "heat-translator": fixes for supporting custom types. This work has been contributed upstream.
  - TOSCA support in ONE through the *Infrastructure Manager* (IM).
  - TOSCA support in OpenStack through the "*heat-translator*".
  - Work with IBM : "*HEAT Translator as a Service*" (Paper)

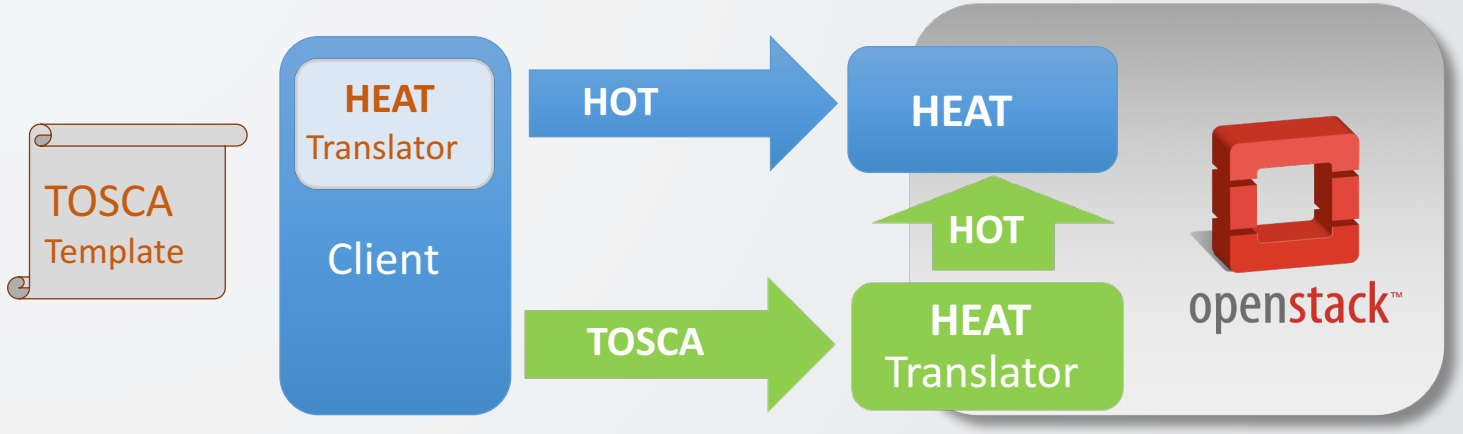

uDocker (https://github.com/indigo-dc/udocker/)

INDIGO - DataCloud

Running Docker containers in userspace, without requiring root privileges (nor Docker)

- Enables the execution of Docker containers without Docker with four different engines based on PRoot, Fakechroot, runC and Singularity.
- It can be used to execute containers in batch or interactive systems where Docker cannot be used or is unavailable.
- None of the uDocker engines requires root privileges being therefore adequate for deployment and use by end-users.
- By using uDocker it is possible to use GPU computing and MPI and Infiniband without software installation or system administration intervention, making it adequate to execute containers in batch systems and HPC clusters. uDocker grew to become a tool that integrates several existing technologies to run Docker containers in unprivileged mode.
- E.g. case studies in Structural Biology, Lattice QCD, for MasterCode (SuSy) at CERN, Fermilab, DESY, etc.

MPI simulations: Open QCD

- Lattice QCD is a strongly computing-characterized discipline (hundreds of millions of CPU hours/year).
- Current simulations run spread over thousands of processor cores in parallel.
- OpenQCD is a very advanced GPL-licensed code to run lattice simulations.

Running MPI codes w/ udocker

- Download and install the container as with MasterCode.
- caveat: exactly the same version of MPI on the host and in the container.
- With udocker, the mpiexec of the host system is used to submit the MPI processes.

```
emanuele [0]> ${HOST_OPENMPI_PATH_BIN}/mpiexec \
 -np 128 udocker run \
 --hostenv --hostauth --user=${USERID} \
 --workdir=${OPENQCD_CONTAINER_DIR} \
 openqcd \
 ${OPENQCD_CONTAINER_DIR}/ym1 -i ym1.in
```

This is not specific to udocker: it applies to all container solutions (singularity, Docker itself) as containers run directly on the hosts using their kernels and drivers

Slides by Emanuele Bagnaschi - 11th Annual Meeting of the Helmholtz Alliance "Physics at the Terascale", 27-29/11/2017
https://indico.desy.de/indico/event/18681/session/8/contribution/109/material/slides/0.pdf

Scaling test

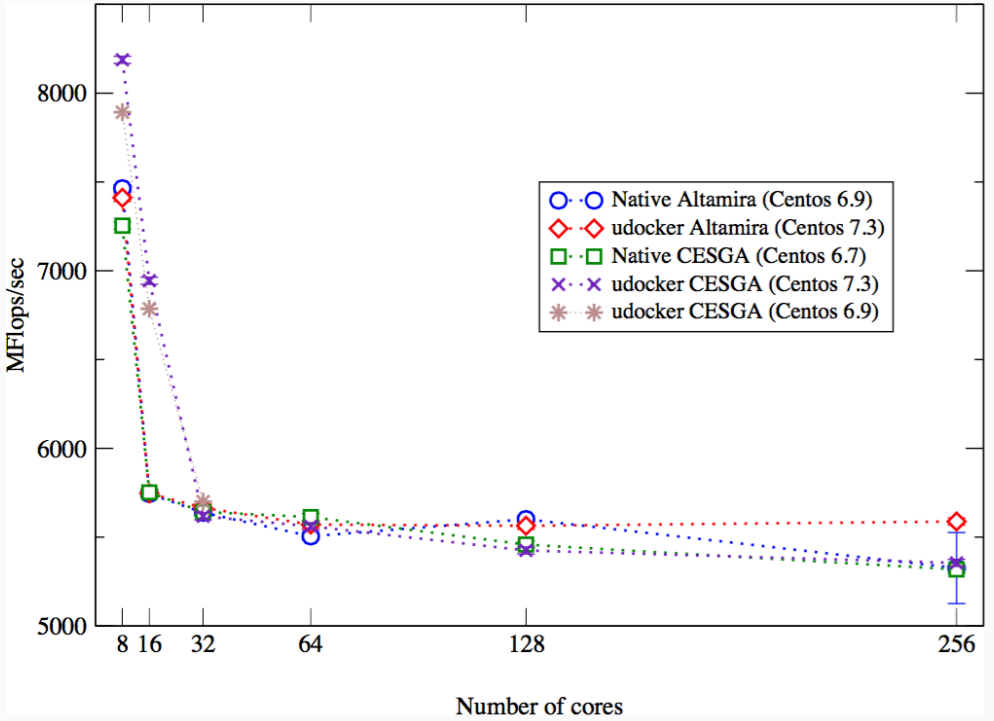

- Scaling performance as a function of the cores for the computation of application of the Dirac operator to a spinor field (Practically, it is a sparse-matrix-matrix $\times$ vector multiplication).
- `udocker` at *least* as fast as the host.
- At CESGA `udocker` *faster* than host because of newest libraries for 8 and 16 cores.

Biomolecular complexes: DisVis, Powerfit and Gromacs

- `DisVis` and `Powerfit` are MIT-licensed codes available on `GitHub` to model biomolecular complexes.
- They leverage GPUs through `OpenCL`, via `PyOpenCL`
- `Gromacs`is a molecular dynamics package for both biochemical and non-biochemical systems.

Running w/ `udocker`

- Download and install the container as with `MasterCode`.
- caveat: exactly the same version of the NVIDIA drivers and libraries needs to be installed on the host and the container.

This is not specific to udocker: it applies to all container solutions (singularity, Docker itself) as containers run directly on the hosts using their kernels and drivers

DisVis

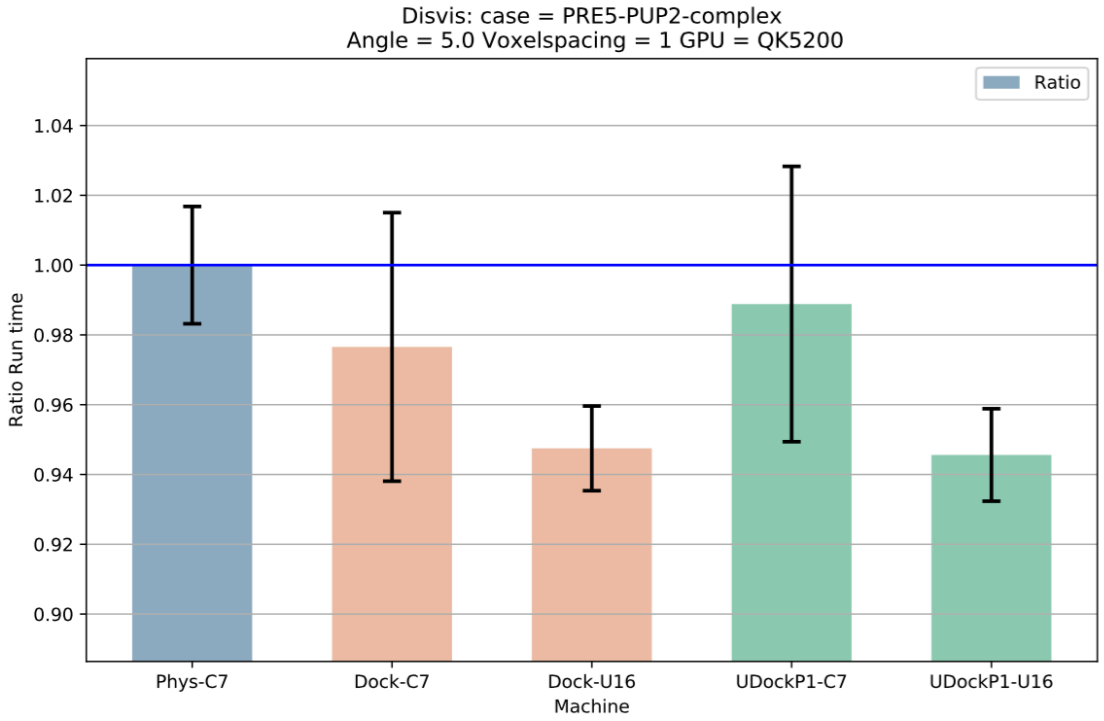

- `udocker` and `docker` have same performance as the host when using `CentOS7`.
- Improved performance due to newer userland libraries when using `Ubuntu16`.

Gromacs

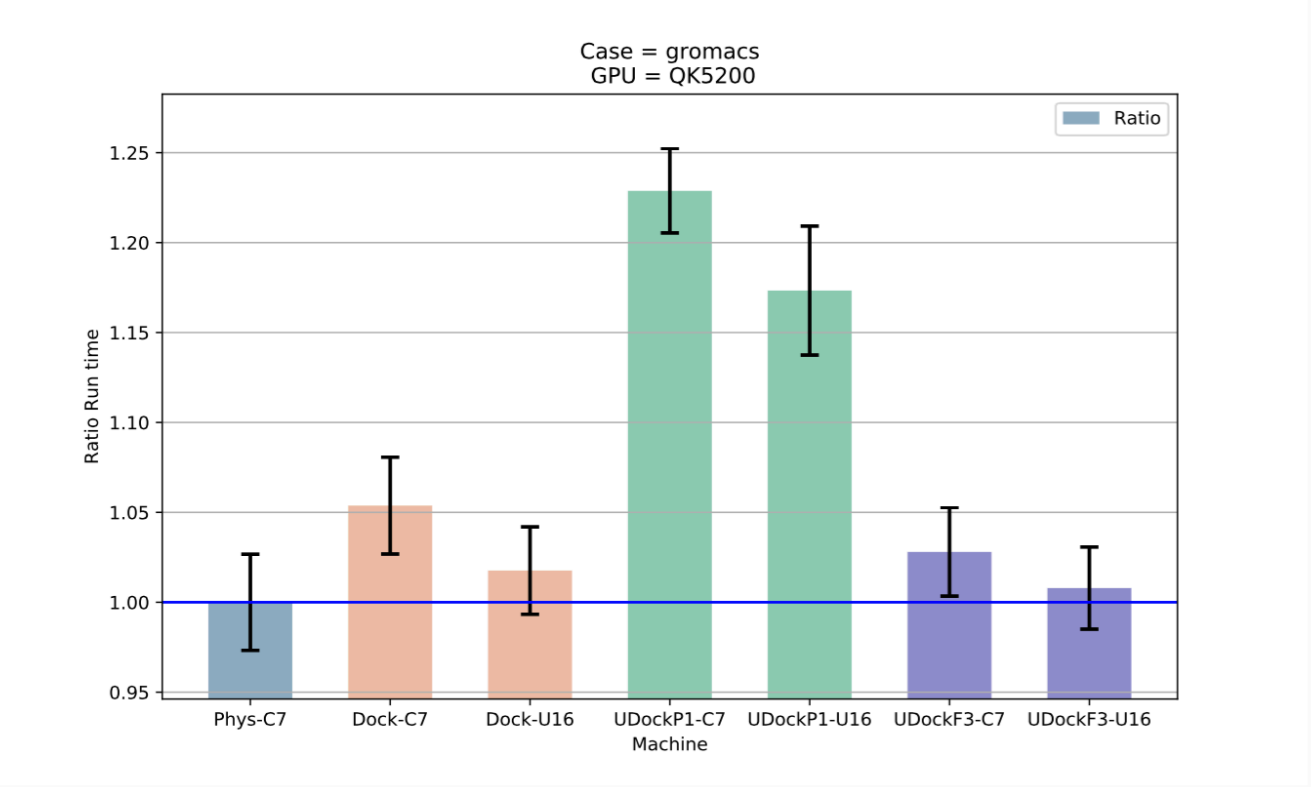

- `udocker` and `docker` are worse than the host when using `CentOS7` of $\mathcal{O}$(3-5%).
- Same performances when using `Ubuntu16`.
- Use of P1 mode results in $\mathcal{O}$(22%) performance hit (due to communication between the GPGPU and the CPU threads – Gromacs spawns 8 OpenMP threads / GPU).

Jorge Gomes (LIP)

udocker and Singularity

- Singularity as an additional udocker execution mode is a recent addition to udocker
  - udocker can use Singularity, if locally available, to execute DOCKER containers (not Singularity containers)
- This is interesting e.g. for CentOS 6 where the default mode of udocker using PTRACE can become slow, if many system calls are invoked
- Or for users that want to use Docker containers using udocker, that provides a Docker-like logic and interface, and would also like to profit from Singularity (when available) to run their containers.

Jorge Gomes (LIP)

udocker + Singularity

- For example, you can run a container in mode P1 (PRoot) to install software using the root emulation mode (you cannot do this with Singularity)…
- … and then you can move to S1 (Singularity) to execute the container:

```
./udocker setup  --execmode=P1  mycontainer
./udocker run mycontainer  dnf install some-packages
./udocker setup  --execmode=S1  mycontainer
./udocker run mycontainer  /bin/bash <<EOF
do-stuff
do-stuff
EOF
```

Deployed on private clouds (INFN @ multiple sites, IFCA), T-Systems' OTC, Microsoft Azure with ~0 configuration changes
Linked to the CMS ITB – moving to the CMS Global Pool

Multi-tenancy, scalability, auto-scaling

INDIGO - DataCloud

- These features are all **automatically managed by several components**, such as the Infrastructure Manager and the INDIGO PaaS

INFN Istituto Nazionale di Fisica Nucleare

CMS Data Analysis on private cloud

CMS

T··Systems···

1.The scale
- 258 running worker nodes (HTCondor)

2.Elasticity
- depending on load takes seconds.

3.Self-healing
- Site problem can be recovered by the system

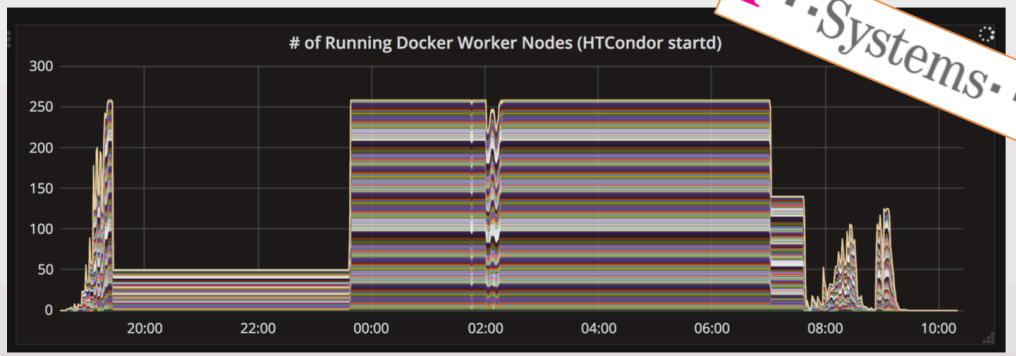

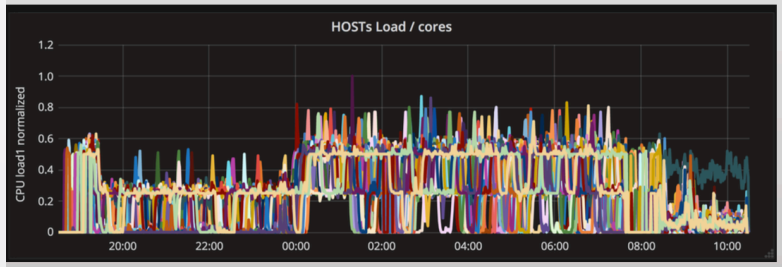

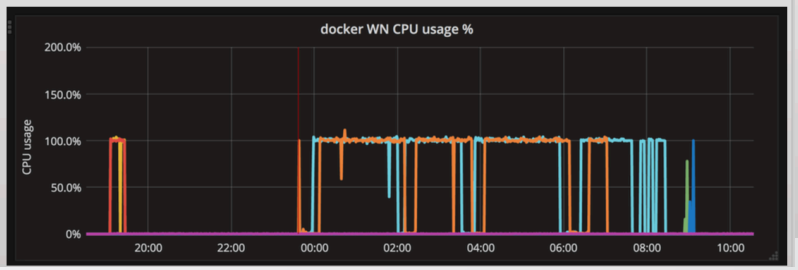

DODAS = **D**ynamic **O**n-**D**emand **A**nalysis **S**ervice

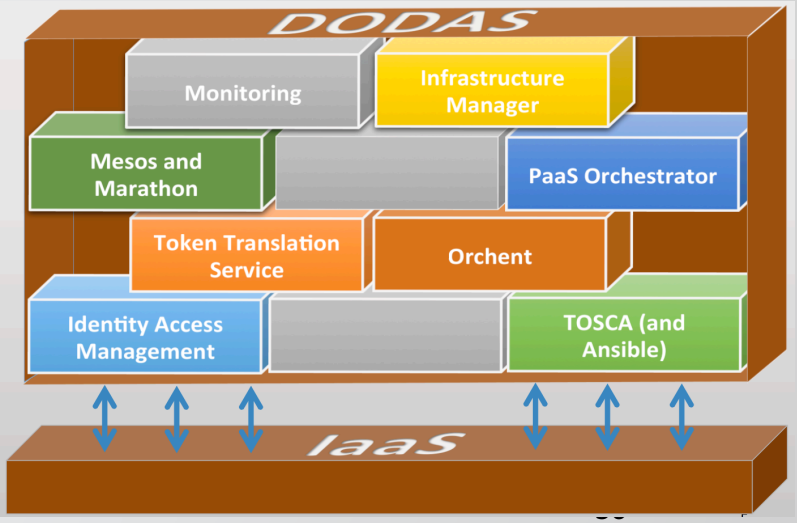

Deployed on private clouds (INFN @ multiple sites, IFCA), T-Systems' OTC, Microsoft Azure with ~0 configuration changes
Linked to the CMS ITB – moving to the CMS Global Pool

Multi-tenancy, scalability, auto-scaling

INDIGO - DataCloud

- These features are all **automatically managed by several components**, such as the Infrastructure Manager and the INDIGO PaaS

Scaling up to **700** worker nodes (1.4k cores, ~3TB RAM) in Nov 2017

INFN Istituto Nazionale di Fisica Nucleare

CMS Data Analysis on private cloud

CMS

DODAS = **D**ynamic **O**n-**D**emand **A**nalysis **S**ervice

1.The scale
- 258 running worker nodes (HTCondor)

2.Elasticity
- depending on load takes seconds.

3.Self-healing
- Site problem can be recovered by the system

T··Systems···

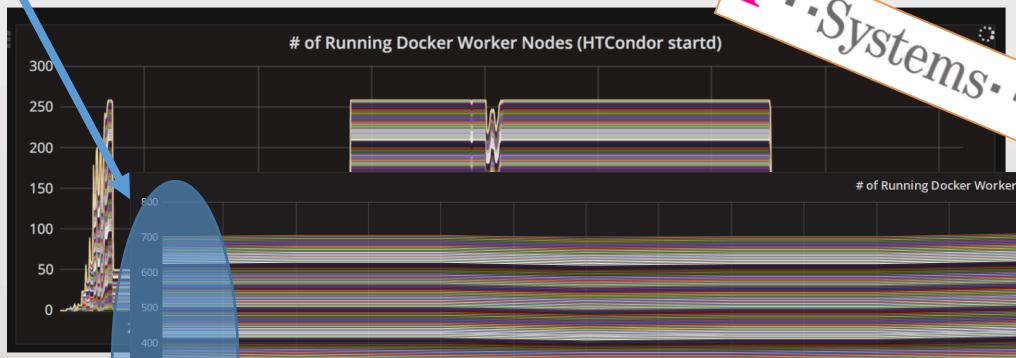

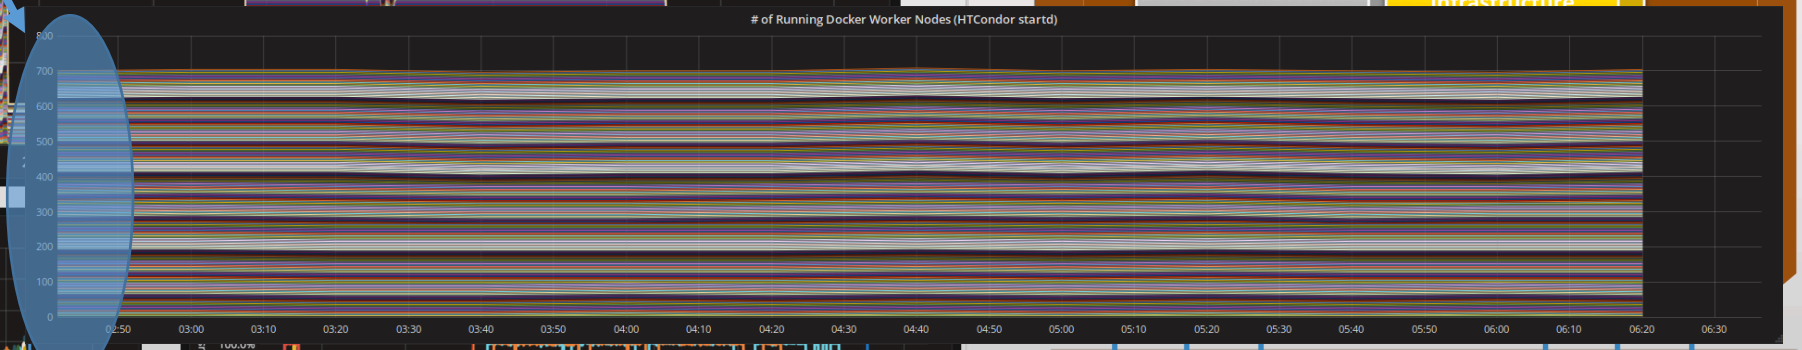

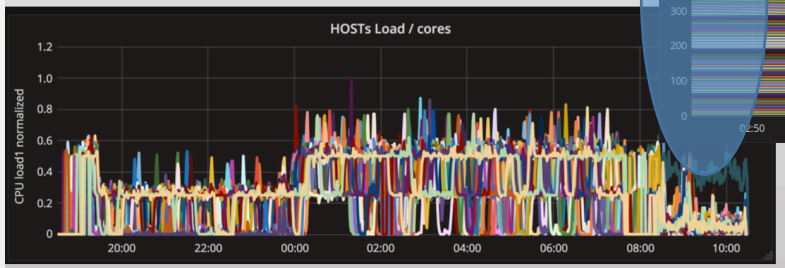

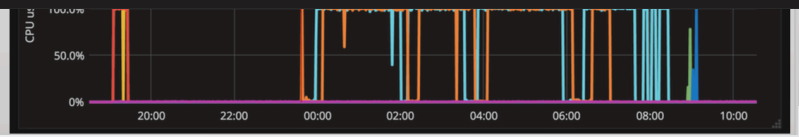

Validation

DODAS = **D**ynamic **O**n-**D**emand **A**nalysis **S**ervice

- ✓ We also worked on physics validation
  - ✓ running a real analysis workflow
    - D meson invariant mass reconstruction with Kpipi decays

CMS Official Tier2 @ Legnaro

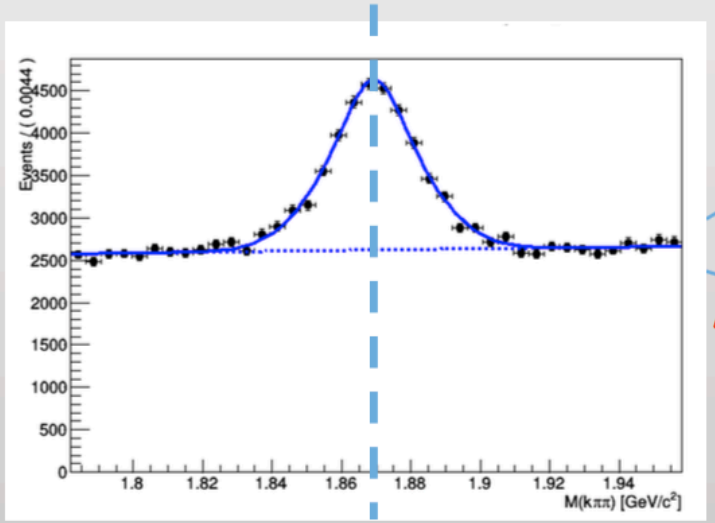

Full data statistics
Standard WLCG grid infrastructure

DODAS site on Private Cloud

DODAS site on Public Cloud

Openstack @ ReCaS INFN - Bari

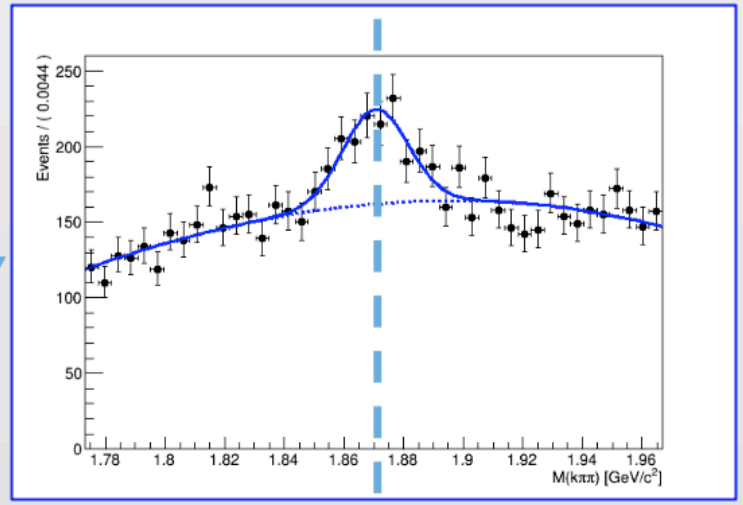

Open Telekom Cloud

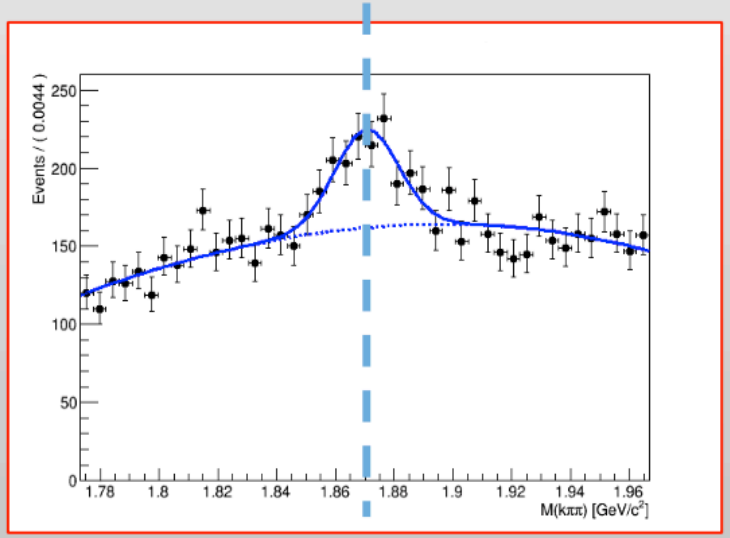

… and Performance Evaluation

DODAS = **D**ynamic **O**n-**D**emand **A**nalysis **S**ervice

- Measuring CPU Time/Job Time
- DODAS ephemeral Site performance VS official grid Sites

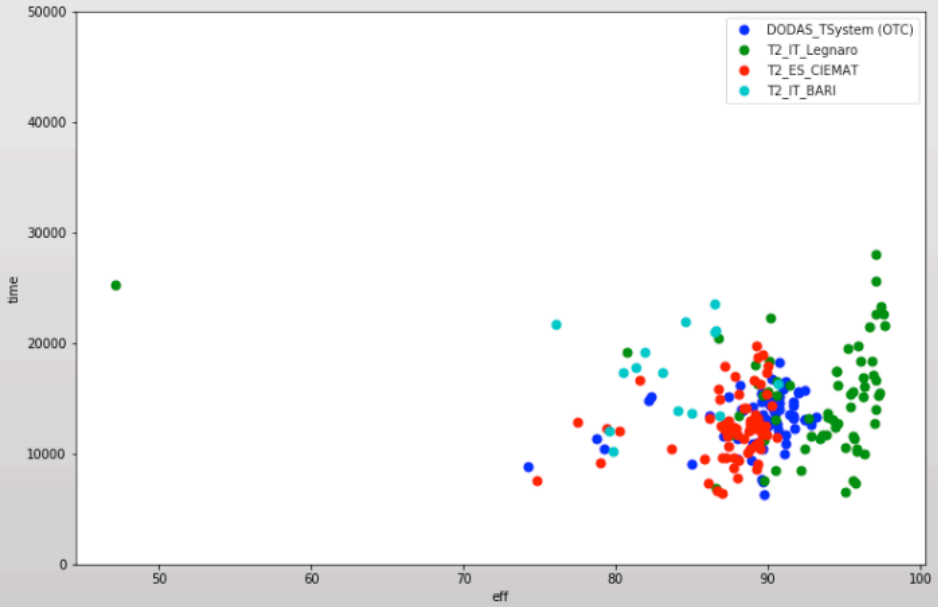

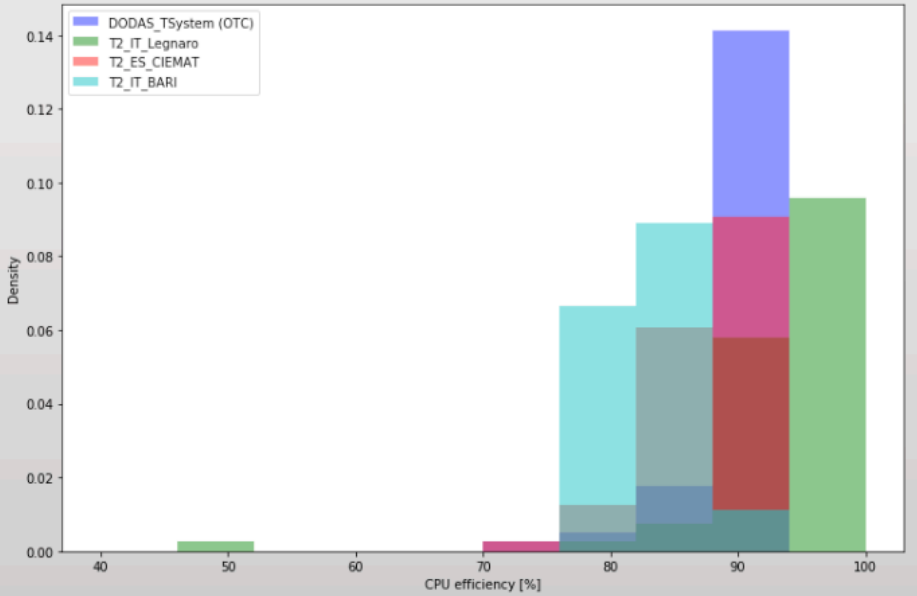

No overhead introduced by DODAS architecture!

Onedata @ T-Systems OTC

Functionality and performance evaluation:

- Run 1000 containers with Oneclient parallel I/O access to datasets
- Manage datasets up to 350 TB across Cyfronet (PL) and OTC (D) infrastructure with 10 Gps connect through GEANT
- I/O performance
  - Block storage (CEPH, Gluster FS, Lustre) and Object Storage
  - Each client can access data close to max. throughput
  - I/O overhead less than 10%
- Scalability: ongoing

I/O Performance

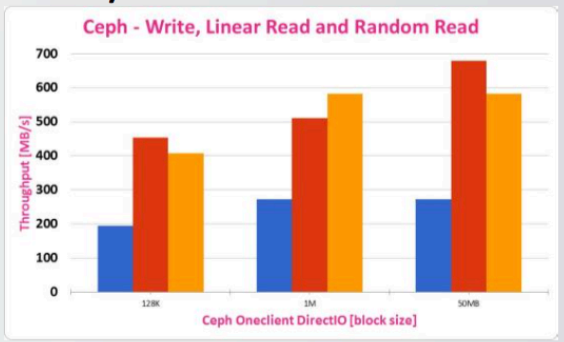

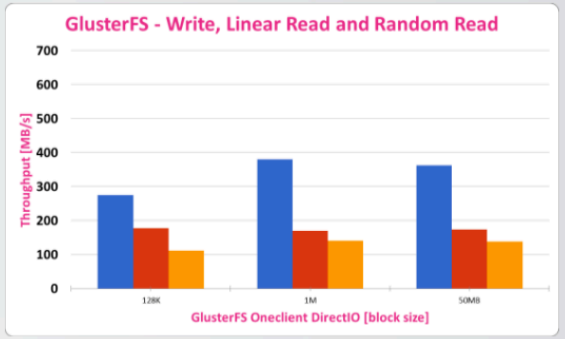

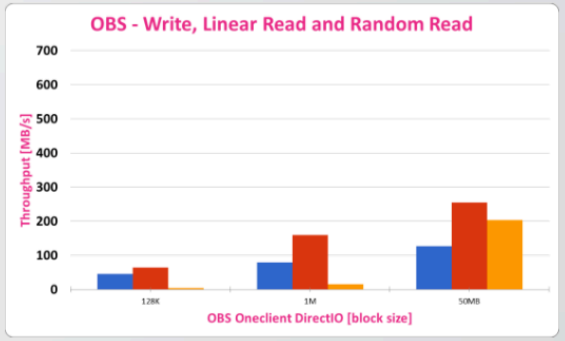

Patrick Fuhrmann (DESY) and Marcus Hardt (KIT)

Storage QoS: Motivation

INDIGO - DataCloud

Providing a common vocabulary of Storage Quality Services for

- Automatic or web based selection of storage qualities in the cloud
  - e.g. DISK, SSD, Tape, Blue Ray, Scratch, Long, Archive, Fast
- Common vocabulary for Data Management Plans
- Commonly accepted vocabulary in Cloud Storage Procurements

The rules on European standardisation allow the European Commission to identify information and communication technology (ICT) technical specifications - that are not national, European or international standards - to be eligible for referencing in public procurement. This allows public authorities to make use of the full range of specifications when buying IT hardware, software and services, allowing for more competition in the field and reducing the risk of lock-in to proprietary systems.

Patrick Fuhrmann (DESY) and Marcus Hardt (KIT)

QoS in storage (Ideal Timeline)

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

Time

Patrick Fuhrmann (DESY) and Marcus Hardt (KIT)

Sustainability through SNIA and RDA

- RDA
  - Working Group: Storage Service Definitions (former QoS and DLC)
  - Vocabulary for DMP, Procurement and automated storage quality selection.
  - Unfortunately only got up to speed right now, but will be continued by DESY, within XDC.
- SNIA
  - INDIGO Reference implementation now official branch in SNIA GitHub
  - Two INDIGO Products listed on SNIA web site as official servers
    - INDIGO Reference implementation
    - dCache.org server

Patrick Fuhrmann (DESY) and Marcus Hardt (KIT)

CDMI Ref. Implementation structure

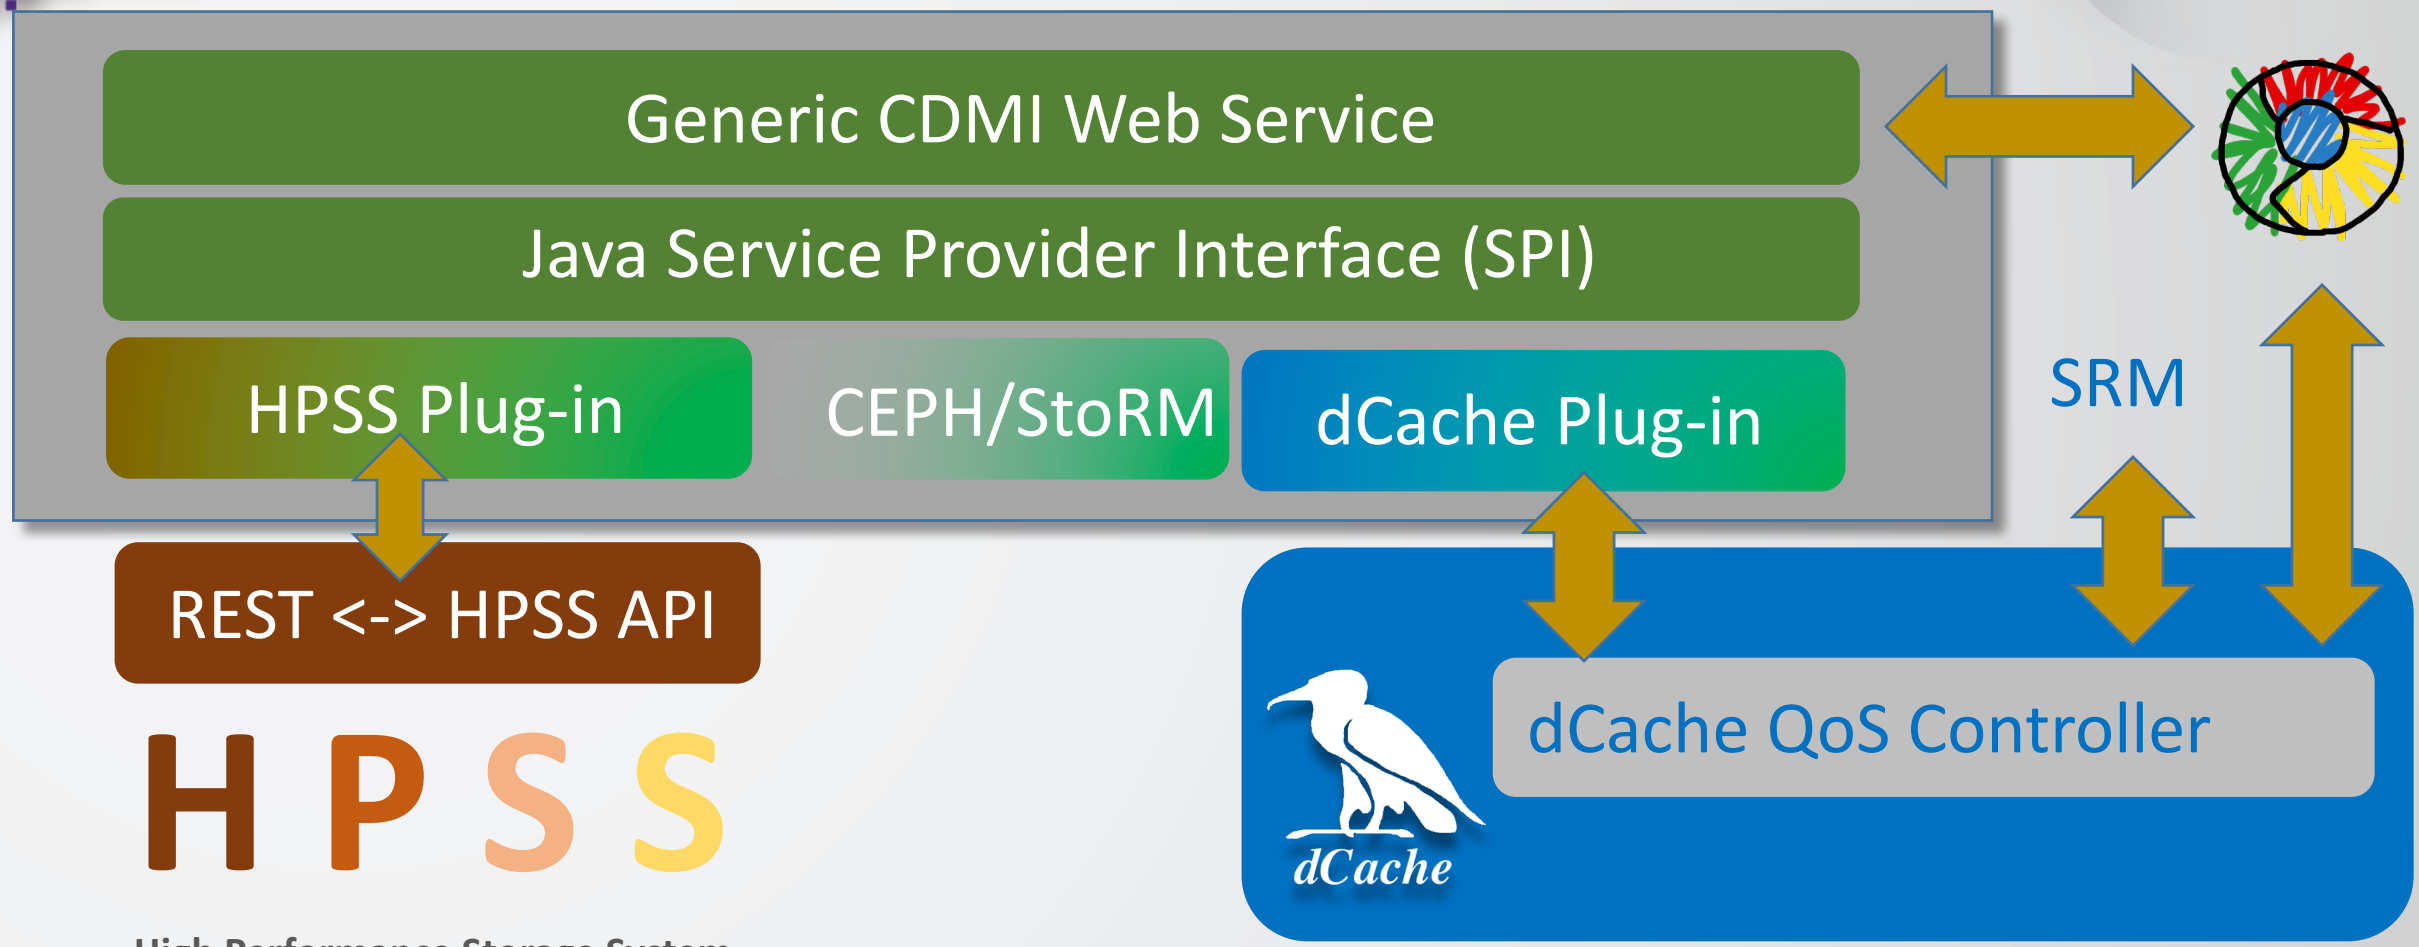

Patrick Fuhrmann (DESY) and Marcus Hardt (KIT)

CDMI Implementation in dCache

QoS interface in production in INDIGO/dCache

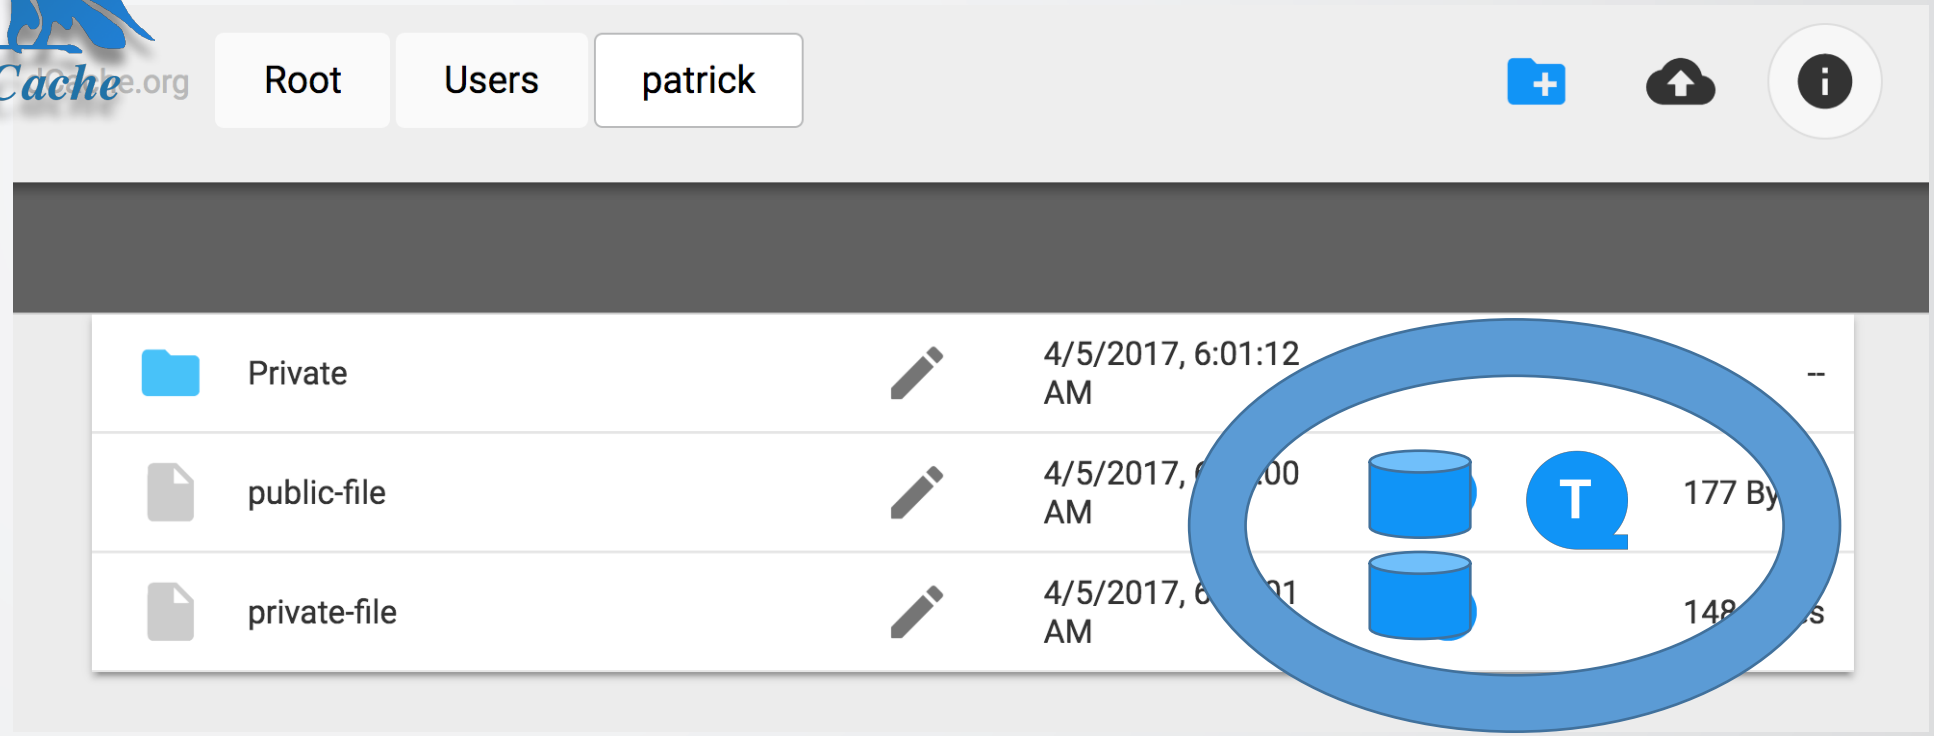

Patrick Fuhrmann (DESY) and Marcus Hardt (KIT)

CDMI Implementation in dCache

INDIGO - DataCloud

dCache

QoS interface in pro

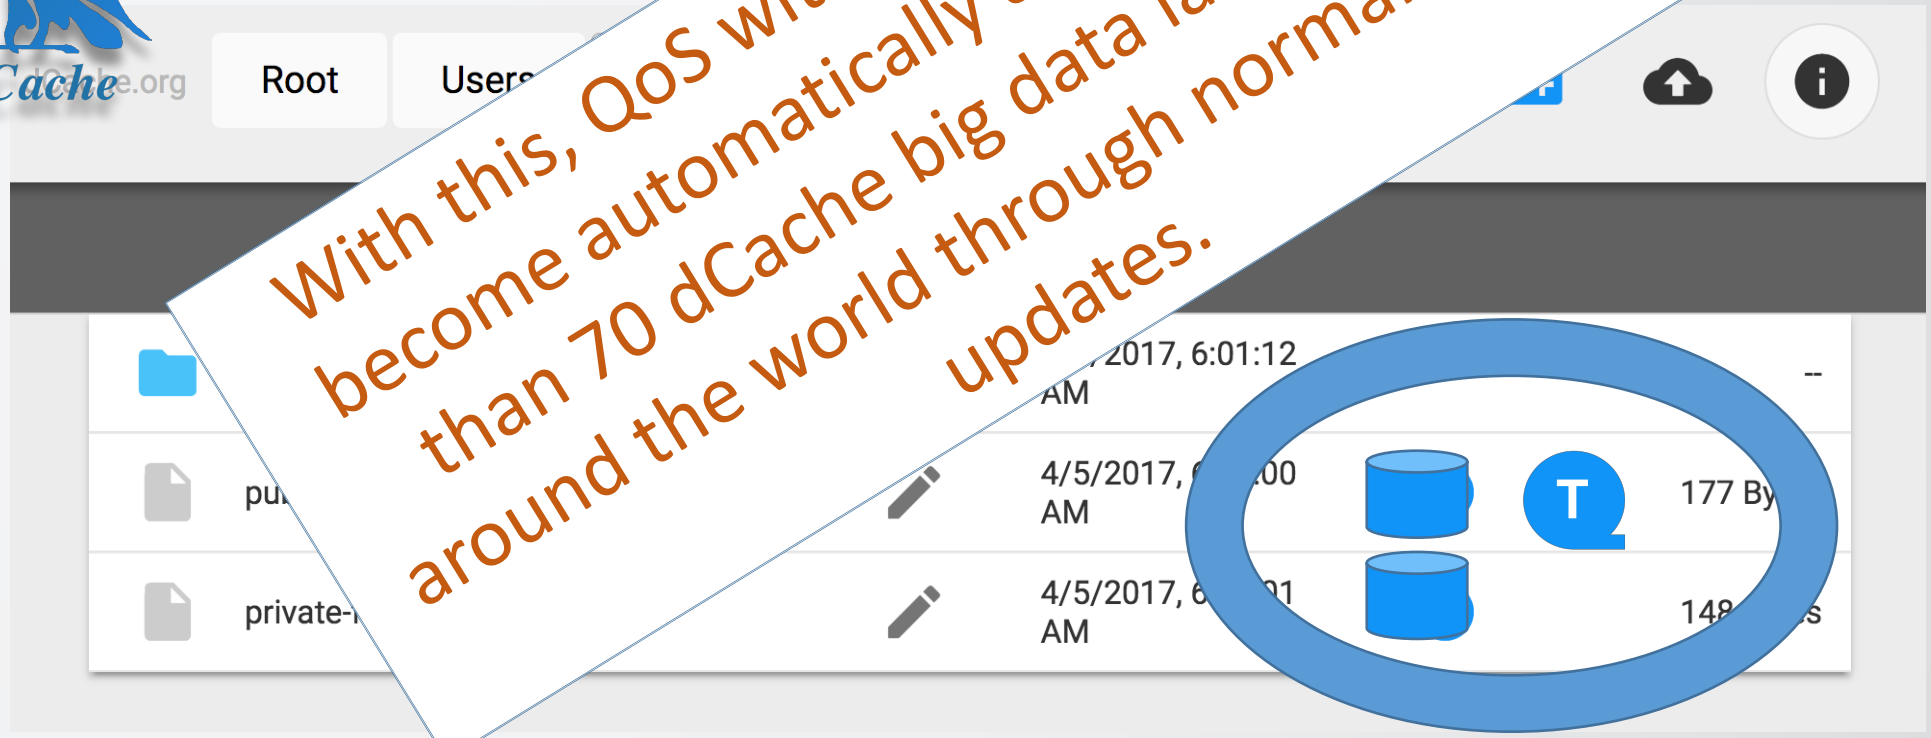

With this, QoS with CDMI and REST will become automatically available at more than 70 dCache big data laboratories around the world through normal release updates.

Patrick Fuhrmann (DESY) and Marcus Hardt (KIT)

CDMI evaluation setup in Europe

INDIGO - DataCloud

- Implementation of CDMI extensions
  - Including Staging
  - Web-Service at KIT continuously verifies 5 evaluation endpoints:
  - http://cdmi-web.data.kit.edu
    - KIT (GPFS, HPSS)
    - BARI (GPFS)
    - CNAF (GPFS, StoRM)
    - DESY (dCache)
    - POZNAN (CEPH)

Patrick Fuhrmann (DESY) and Marcus Hardt (KIT)

And now putting it all together.

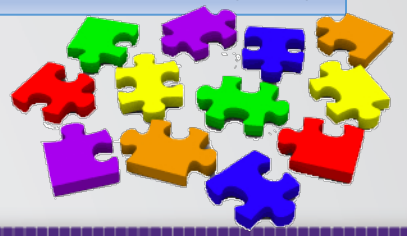

INDIGO - DataCloud

Storage Broker Webservice

curl://

CDMI

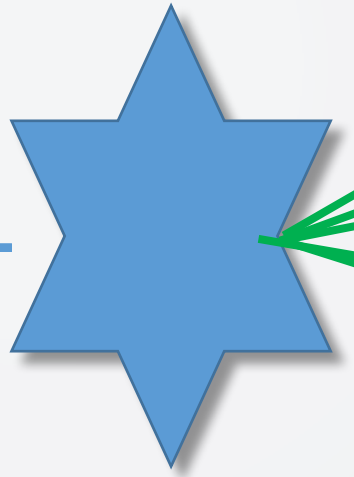

As example for a PaaS service

IBM GPFS / HPFS

ceph

StoRM

dCache

Patrick Fuhrmann (DESY) and Marcus Hardt (KIT)

The CDMI Broker Page (at KIT)

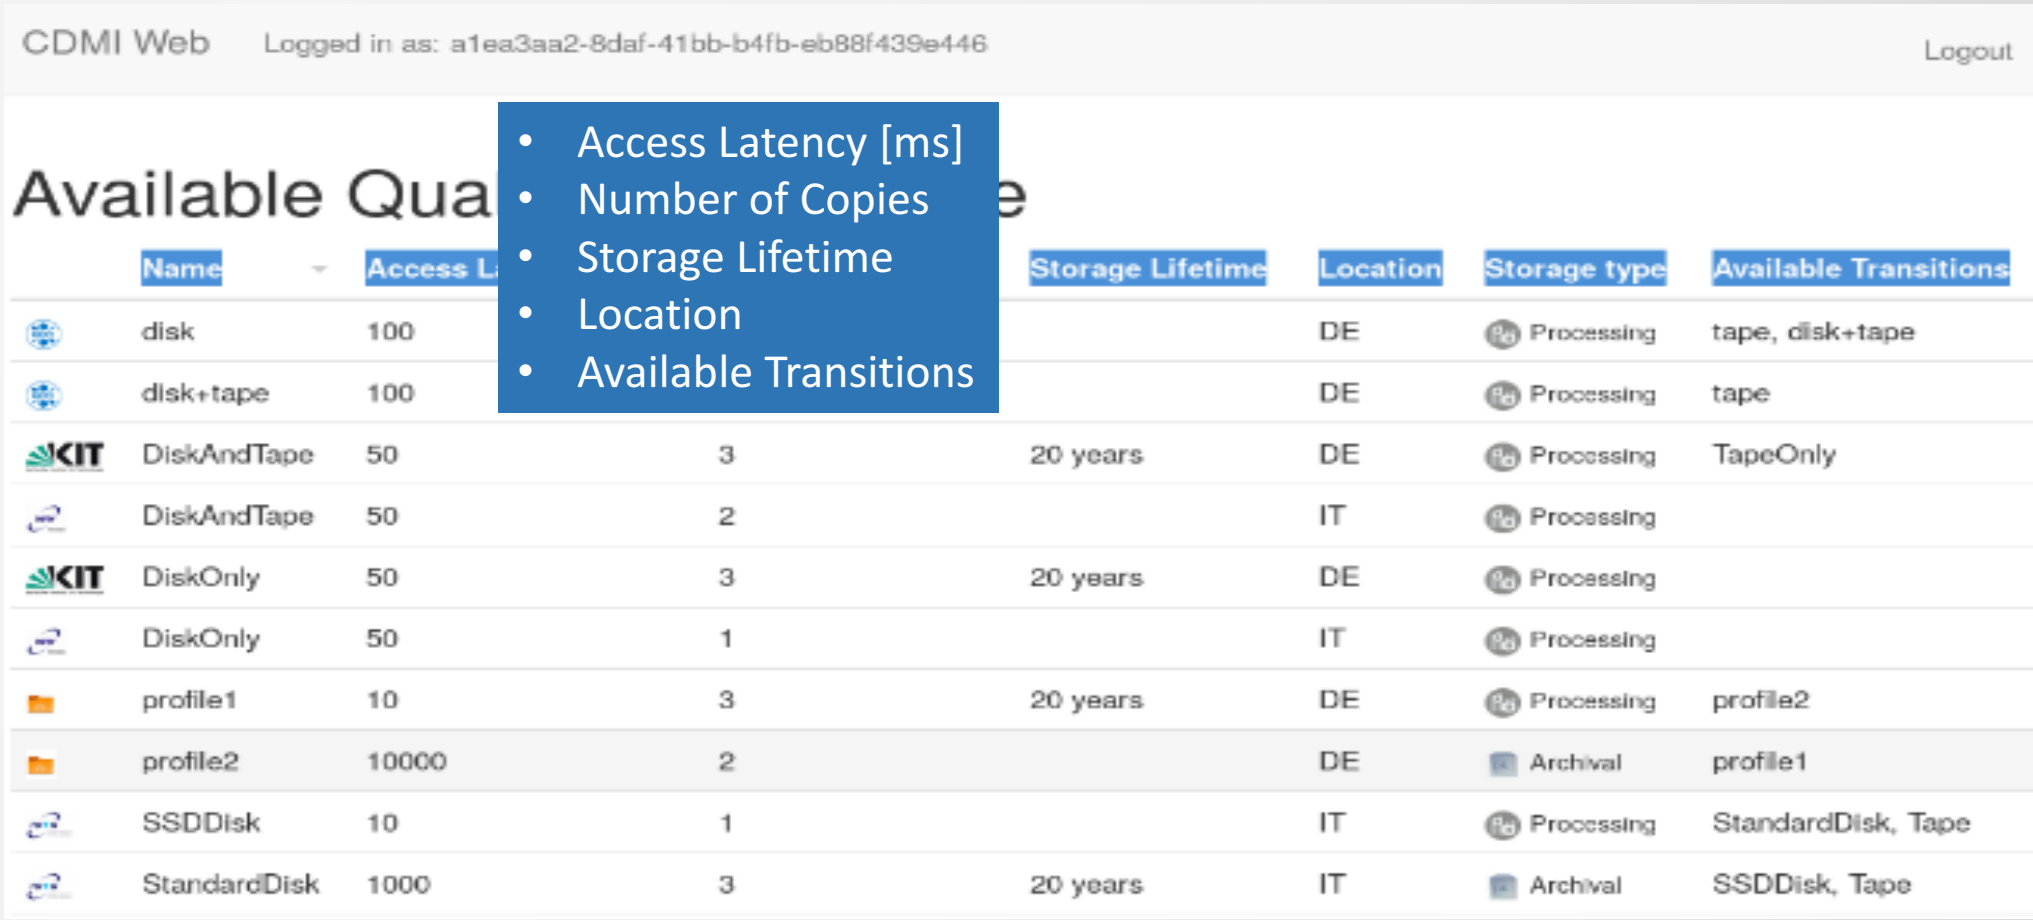

- Access Latency [ms]
- Number of Copies
- Storage Lifetime
- Location
- Available Transitions

Patrick Fuhrmann (DESY) and Marcus Hardt (KIT)

The CDMI Broker Page (at KIT)

INDIGO - DataCloud

https://dcache-qos-01.desy.de:8443/kit/

Create Directory | Upload File

Name	Type	Current QoS	Target QoS	
Picture1.png	File	disk	Select	Delete

https://dcache-qos-01.desy.de:8443/kit/

Create Directory | Upload File

Name	Type	Current QoS	Target QoS	
HelmholtzJRG-2017.docx	File	disk	Select	Delete
Picture1.png	File	disk	Select	Delete

HelmholtzJRG-2017.docx uploaded

Next steps for storage QoS & Data lifecycle: see the eXtreme-DataCloud project (XDC)

Results and Impact

See also “INDIGO-DataCloud:A data and computing platform to facilitate seamless access to e-infrastructures”, https://arxiv.org/abs/1711.01981

Key results

- ✓Categorization of user community requests, identification of common requirements, classification into the Computational, Storage and Infrastructural areas
- ✓Identification of technology gaps
- ✓Definition of the INDIGO Technical Architecture, covering all the three Cloud layers (IaaS, PaaS, SaaS)
- ✓Definition of stringent procedures and quality metrics for software development and packaging
- ✓Guidelines on how applications can be built using a service oriented architecture in the Cloud, using INDIGO solutions
- ✓Two major software releases (MidnightBlue and ElectricIndigo), with several minor updates each.
- ✓Two large distributed testbeds, one for supporting development activities, another for pre-production applications
- ✓Several “generic use case” examples
- ✓16 complex use cases demonstrated in production infrastructures by multiple communities
- ✓Comprehensive success metrics researched, defined and compiled for all use cases
- ✓Extensive documentation and training material

Components and Patches Merged in Upstream Open Source Projects

- OpenStack (https://www.openstack.org)
  - Nova Docker
  - Heat
  - OpenID-Connect for Keystone
  - Pre-emptible instances support (under discussion)
- OpenNebula (http://opennebula.org)
  - OneDock
  - Patch to support FaSS
- Infrastructure Manager (http://www.grycap.upv.es/im/index.php)
- Clues (http://www.grycap.upv.es/clues/eng/index.php)
- Onedata (https://onedata.org)
- Apache Libcloud (https://github.com/apache/libcloud)
- Kepler Workflow Manager (https://kepler-project.org/)
- TOSCA adaptor for JSAGA (http://software.in2p3.fr/jsaga/dev/)
- OCCI implementation for OpenStack (https://github.com/openstack/ooi)
- Extended AWS support for rOCCI in OpenNebula. Python and Java libraries for OCCI support.
- CDMI and QoS extensions for dCache (https://www.dcache.org)
- Workflow interface extensions for Ophidia (http://ophidia.cmcc.it)
- OpenID Connect Java implementation for dCache (https://www.dcache.org)
- MitreID (https://mitreid.org/) and OpenID Connect (http://openid.net/connect/) libraries
- FutureGateway (https://www.catania-science-gateways.it/)

Advancement

- Advancements beyond the state-of-the-art:
  - A **standards-based PaaS layer for both computing and storage**, abstracting the lower (IaaS) layer independently from the actual Cloud technology adopted by IaaS resource providers.
    - → applications run unmodified in private or public clouds
    - → communities now have ways to express their requirements and deploy them as applications over hybrid Cloud infrastructures. This was often simply not possible before INDIGO.
  - **Reductions of the performance impact** brought about by virtualization
    - → transparent adoption of containers, both directly (udocker) or associated with sophisticated and customized managers (Mesos)
  - **Users can access distributed resources** via direct API, custom workflows or flexible Science Gateways, integrating standard and existing authentication technologies
  - **Consistent, extensible authentication and authorization architecture** over distributed infrastructures

Impact beyond the project lifetime

- **Open license** for the INDIGO-DataCloud software components
- **Many INDIGO components or modifications brought upstream** in open source projects
- INDIGO components at Technology Readiness Level (TRL) 8 **included into and supported by the EOSC-hub project** and are part of many Common Services, Thematic Services and Competence Centers
- **Two direct INDIGO-DataCloud follow-on projects**, eXtreme-DataCloud and DEEP-HybridDataCloud, just started in November 2017
- We have set up **strong connections** with standard bodies, such as **RDA and SNIA**
- INDIGO leads a **EC-funded Common Dissemination Booster (CDB) project group**, defining a common portfolio of results that will be disseminated to end-users, with a specific eye on exploitation opportunities
- An evaluation about bringing the **evolution of INDIGO components in the Horizon2020 WP 2018-2020** is ongoing

INDIGO in the EOSC-hub Thematic Services

THEMATIC SERVICE	DESCRIPTION	INDIGO SOLUTIONS specified in the proposal
ECAS	Climate Analytics Service (ECAS) provided by ENES	Ophidia, KEPLER, FutureGateway
DARIAH SG	DARIAH science gateway tailored for the digital arts and humanities communities	OneDock, OpenStack Nova Docker, FutureGateway, Onedata
OPENCoastS	OpenCoastS: On-demand Operational Coastal Circulation Forecast Service	INDIGO udocker Infrastructure Manager Orchestration (TOSCA, HEAT)
WeNMR	Structural biology services DISVIS, POWERFIT, HADDOCK, GROMACS, AMPS-NMR, CS-ROSETTA, UNIO, FANTEN	IAM, PaaS Orchestrator, Infrastructure Manager FutureGateway, Onedata
DODAS	Dynamic On Demand Analysis Service	IAM, TTS, PaaS Orchestrator, Orchent, IM, TOSCA, Onedata, FutureGateway
LifeWatch	PAIRQURS, Citizen science services, GBIF, Digital Knowledge preservation framework, remote monitoring and smart sensing.	IAM, PaaS Orchestrator, Infrastructure Manager FutureGateway, Onedata
EO Pillar	Earth observation services coordinated by ESA. The tools are: MEA, EPOSAR, Sentinel playground, Datacube analytic service, Geohazards exploitation platform,OSS-X Sentinel service	Onedata. Analysis of INDIGO cloud software add-ons for OpenStack.
CMI	The Component MetaData Infrastructure Including the Virtual Language Observatory and the Virtual Collection Registry, provided by CLARIN	Providing interoperable metadata for (digital) humanities between both CLARIN and DARIAH.
GEOSS	GEO DAB (Discovery and Access Broker) GEOSS portal	-

Thank you

https://www.indigo-datacloud.eu

Better Software for Better Science.

@indigodatacloud www.indigo-datacloud.eu https://www.facebook.com/indigodatacloud/

INDIGO - DataCloud

BACKUP Slides

INDIGO & EOSC **in production**: >= TRL8*

- Several **INDIGO solutions and activities** are in the new **EOSC-hub proposal**, submitted by a consortium of 74 partners **under the coordination of EGI, EUDAT and INDIGO-DataCloud**
- With **INDIGO components** such as Identity and Access Management, Token Translation, Virtual filesystems (Onedata), Advanced IaaS Services, the Infrastructure Manager, the INDIGO PaaS and its orchestrator, web front-end services, user-level containers
- INDIGO was also given the **overall technical coordination of the project** and will contribute to **training, support, external liaison, stakeholder engagement, policy**.

*: **TRL** = **T**echnology **R**eadiness **L**evel. TRL 8 means production-ready at scale

INDIGO & EOSC in evolution: < TRL8

- **Novel features evolving INDIGO components** are a key part of two proposals to the **EINFRA-21-2017 call (eXtreme-DataCloud and DEEP Hybrid DataCloud)**, both recently approved and started in November 2017:
  - Intelligent dataset distribution and data lifecycle management
  - Smart caching
  - Orchestrating Computing Workflows based on policy driven or adaptive data movements
  - Flexible metadata management for big data sets
  - Access to bare-metal resources on the Cloud
  - PaaS-Level access to HPC resources
  - Extensions to the INDIGO Orchestrator for hybrid IaaS deployments and scale out to 3rd party clouds
  - Extensions to the INDIGO Virtual Router Appliance
  - Real-time, streaming-based data ingestion and processing

eXtreme-DataCloud (XDC)

Participant No	Participant organisation name	Participant short name	Country
1 (Coord.)	Istituto Nazionale di Fisica Nucleare	INFN	Italy
2	Universidad de Cantabria	UC	Spain
3	Stiftung Deutsches Elektronen-Synchrotron	DESY	Germany
4	European Organization for Nuclear Research	CERN	Switzerland
5	Centre National de la Recherche Scientifique	CNRS	France
6	Akademia Gorniczo-Hutnicza Im. Stanislawa Staszica W Krakowie	AGH-UST	Poland
7	European Clinical Research Infrastructure Network	ECRIN	France
8	Stichting European Grid Initiative	EGI Foundation	The Netherlands

- From the DoA:
  - "The eXtreme DataCloud (XDC) project will develop **scalable technologies for federating storage resources and managing data in highly distributed computing environments.**"
  - "The XDC roots lie in the experience of the INDIGO-DataCloud project. It catches concepts and methodologies developed in that project, in particular those connected with Storage Quality of Service and data lifecycle management."

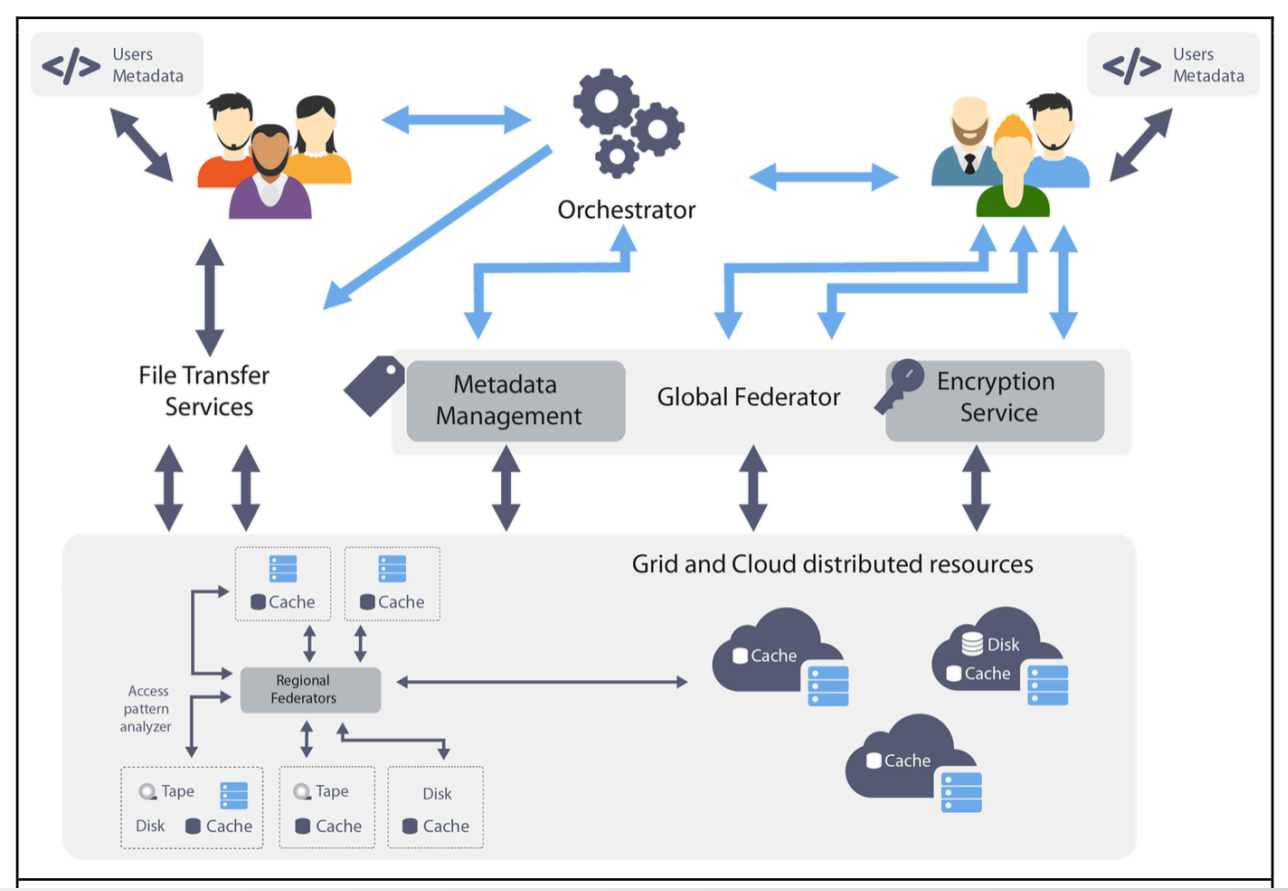

DEEP-Hybrid DataCloud

- From the DoA:
  - “DEEP Hybrid DataCloud project [will] support **intensive computing techniques** that require specialized HPC hardware, like GPUs or low latency interconnects, to explore very large datasets. A **Hybrid Cloud approach** enables the access to such resources that are not easily reachable by the researchers at the scale needed in the current EU e-infrastructure.”
  - “The project will evolve to TRL8 existing services and technologies at TRL6+, including relevant contributions to the EOSC by the INDIGO-DataCloud H2020 project, that the project will enrich with new functionalities.”

Participant No	Participant organisation name	Participant short name	Country
1 (coord.)	Agencia Estatal Consejo Superior de Investigaciones Científicas	CSIC	Spain
2	Laboratório de Instrumentação e Física Experimental de partículas	LIP	Portugal
3	Istituto Nazionale di Fisica Nucleare	INFN	Italy
4	Instytut Chemii Bioorganicznej Polskiej Akademii Nauk	PSNC	Poland
5	Karlsruhe Institute of Technology	KIT	Germany
6	Universitat Politècnica de València	UPV	Spain
7	CESNET, zájmové sdružení právnických osob	CESNET	Czech Republic
8	Ustav Informatiky, Slovenska Akademia Vied	IISAS	Slovakia
9	Atos Spain S.L.	ATOS	Spain
10	Helmholtz Zentrum Muenchen, Deutsches Forschungszentrum fuer Gesundheit und Umwelt (GmbH)	HMGU	Germany

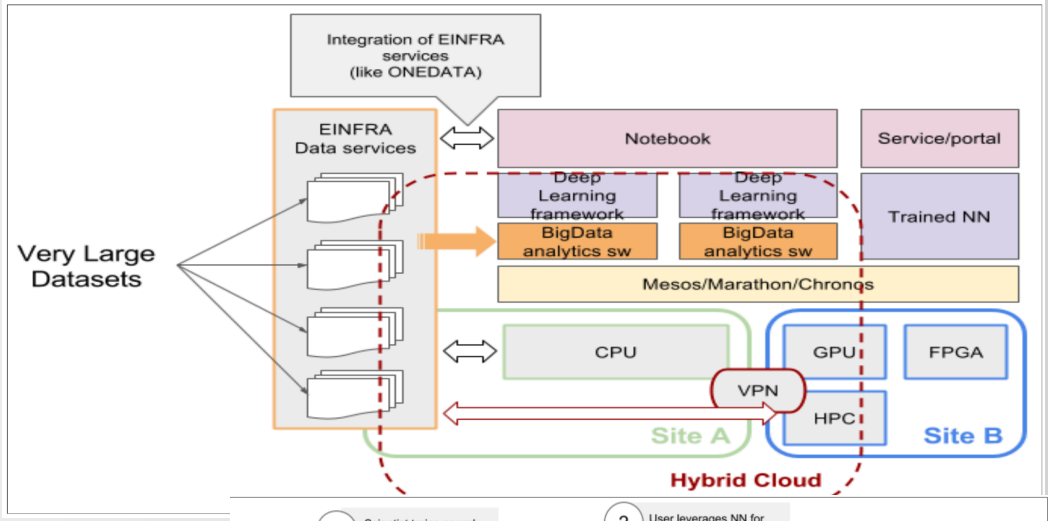

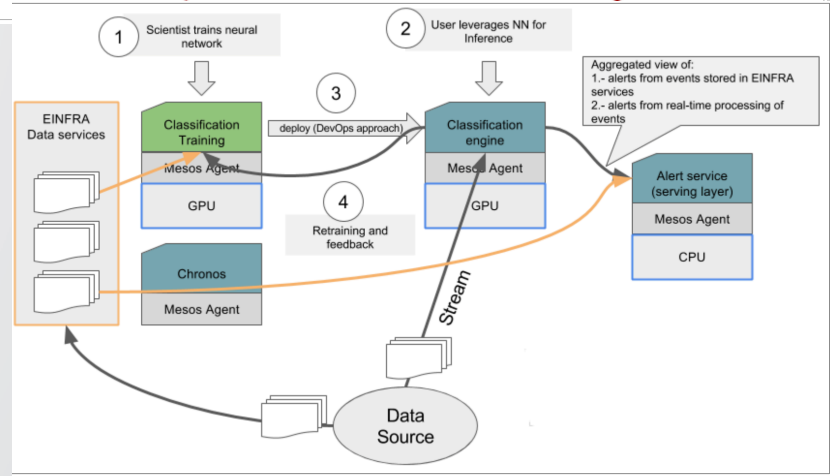

The *Common Dissemination Booster*

- INDIGO-DataCloud is the leading project in a group of 3 projects (INDIGO-DataCloud, eXtreme-DataCloud, DEEP-HybridDataCloud), which together **submitted a proposal (approved) to the Common Dissemination Booster** (CDB) for all 5 CDB services.
- This will help to **establish branding in EU projects at the EOSC level**, and will contribute to the exploitation of project results **at both research and business levels**

Continuing the INDIGO brand

- A proposal for an "**INDIGO Software Stack Providers Collaboration Agreement**" is being discussed by the INDIGO PMB and will be submitted soon to the INDIGO-DataCloud Collaboration Board
- In this agreement, we propose that **partners join in contributing, extending or maintaining INDIGO-DataCloud software components** beyond the boundaries of – for example– Horizon 2020 projects.
- The **intended benefits** are:
  - Support of software components as a part of a well established and known Service Catalogue and Architecture
  - Avail of well-established dissemination channels (e.g. CDB)
  - Find common exploitation channels to reach users with new services
  - Stronger influence on larger initiatives (such as those linked to the EOSC)
  - Opportunities for new common proposals

udocker Execution Modes

https://github.com/indigo-dc/udocker/

3.25. Setup

```
udocker setup [--execmode=XY] [--force] CONTAINER-ID|CONTAINER-NAME
```

Choose an execution mode to define how a given container will be executed. Enables selection of an execution engine and related execution modes. Without --execmode=XY, setup will print the current execution mode for the given container.

Options:

- `--execmode=XY` choose an execution mode
- `--force` force the selection of the execution mode, can be used to force the change of an execution mode when it fails namely if it is transferred to a remote host while in one of the Fn modes.

Mode	Engine	Description	Changes container
P1	PRoot	accelerated mode using seccomp	No
P2	PRoot	seccomp accelerated mode disabled	No
F1	Fakechroot	exec with direct loader invocation	symbolic links
F2	Fakechroot	F1 plus modified loader	F1 + ld.so
F3	Fakechroot	fix ELF headers in binaries	F2 + ELF headers
F4	Fakechroot	F3 plus enables new executables and libs	same as F3
R1	runC	rootless user mode namespaces	resolv, passwd
S1	Singularity	uses singularity if available in the host	passwd

The default execution mode is P1.

Udocker inner workings

Features

- udocker pre-compiled code and the containers are download to ${UDOCKER_DIR}, by default ${HOME}/.udocker.
- Docker layered FS is UnionFS based. Images are pulled by downloading the corresponding layers and metadata (docker Hub REST API).
- udocker implements parsing of docker container and of a subset of metadata.
- Different execution engines: PTRACE, LD_PRELOAD, runC, Singularity.
- Tested with GPGPU and MPI aware applications.

PTRACE engine

- Implements through PRoot.
- PRoot uses PTRACE to change the pathnames dynamically and to execute the binary transparently inside the container (P2 mode).
- Patches have been written to make SECCOMP works with PTRACE (P1 mode).

Mode	Description	Changes container
P1	PRoot+SECCOMP	No
P2	PRoot	No

Udocker inner workings

Features

- udocker pre-compiled code and the containers are download to ${UDOCKER_DIR}, by default ${HOME}/.udocker.
- Docker layered FS is UnionFS based. Images are pulled by downloading the corresponding layers and metadata (docker Hub REST API).
- udocker implements parsing of docker container and of a subset of metadata.
- Different execution engines: PTRACE, LD_PRELOAD, runC, Singularity.
- Tested with GPGPU and MPI aware applications.

LD_PRELOAD engine

- Based on the Fakechroot library.
- Implemented several workarounds to address Fakechroot shortcomings and to avoid letting the containerized application load system libraries.
- Modified version of PatchELF to perform the modifications of the binaries.

Mode	Description	Changes container
F1	exec w/ direct loader	symlinks
F2	F1 + mod. loader	F1+ld.so
F3	ELF header mod.	F2+ELF headers
F4	F3 + new execs and libs	as F3

Udocker inner workings

Features

- udocker pre-compiled code and the containers are download to ${UDOCKER_DIR}, by default ${HOME}/.udocker.
- Docker layered FS is UnionFS based. Images are pulled by downloading the corresponding layers and metadata (docker Hub REST API).
- udocker implements parsing of docker container and of a subset of metadata.
- Different execution engines: PTRACE, LD_PRELOAD, runC, Singularity.
- Tested with GPGPU and MPI aware applications.

RunC engine

- Support for unprivileged User Namespace and rootless container using RunC.
- udocker performs the translation between docker metadata and cli args and the OCI specs to run the container in unprivileged mode.

Singularity

- Support running singularity containers.

Really Docker containers via Singularity, if the latter is installed on the host

Mode	Description	Changes container
R1	rootless usermod namesp.	resolv,passwd
S1	singularity	passwd

Udocker inner workings

Features

- udocker pre-compiled code and the containers are download to ${UDOCKER_DIR}, by default ${HOME}/.udocker.
- Docker layered FS is UnionFS based. Images are pulled by downloading the corresponding layers and metadata (docker Hub REST API).
- udocker implements parsing of docker container and of a subset of metadata.
- Different execution engines: PTRACE, LD_PRELOAD, runC, Singularity.
- Tested with GPGPU and MPI aware applications.

RunC engine

- Support for unprivileged User Namespace and rootless container using RunC.
- udocker performs the translation between docker metadata and cli args and the OCI specs to run the container in unprivileged mode.

Singularity

- Support running singularity containers.

Really Docker containers via Singularity, if the latter is installed on the host

Mode	Description	Changes container
R1	rootless usermod namesp.	resolv,passwd
S1	singularity	passwd