

INDIGO Review Bologna, Nov 8, 2016



INDIGO - DataCloud

Welcome to WP4's Virtualized Resources

WP4.1
Computing
Virtualisation

Alvaro Lopez Garcia, CSIC

WP4.2
Storage
Virtualisation

Marcus Hardt, KIT

WP4.3
Network
Virtualisation

Zdenek Sustr, CESNET

Content



INDIGO - DataCloud



- High level objectives of INDIGO WP4.
- High level overview.
- Details on selected highlights.
- Usage and contribution to standards in WP4.
- Summary.

High level WP4 Objective



INDIGO - DataCloud

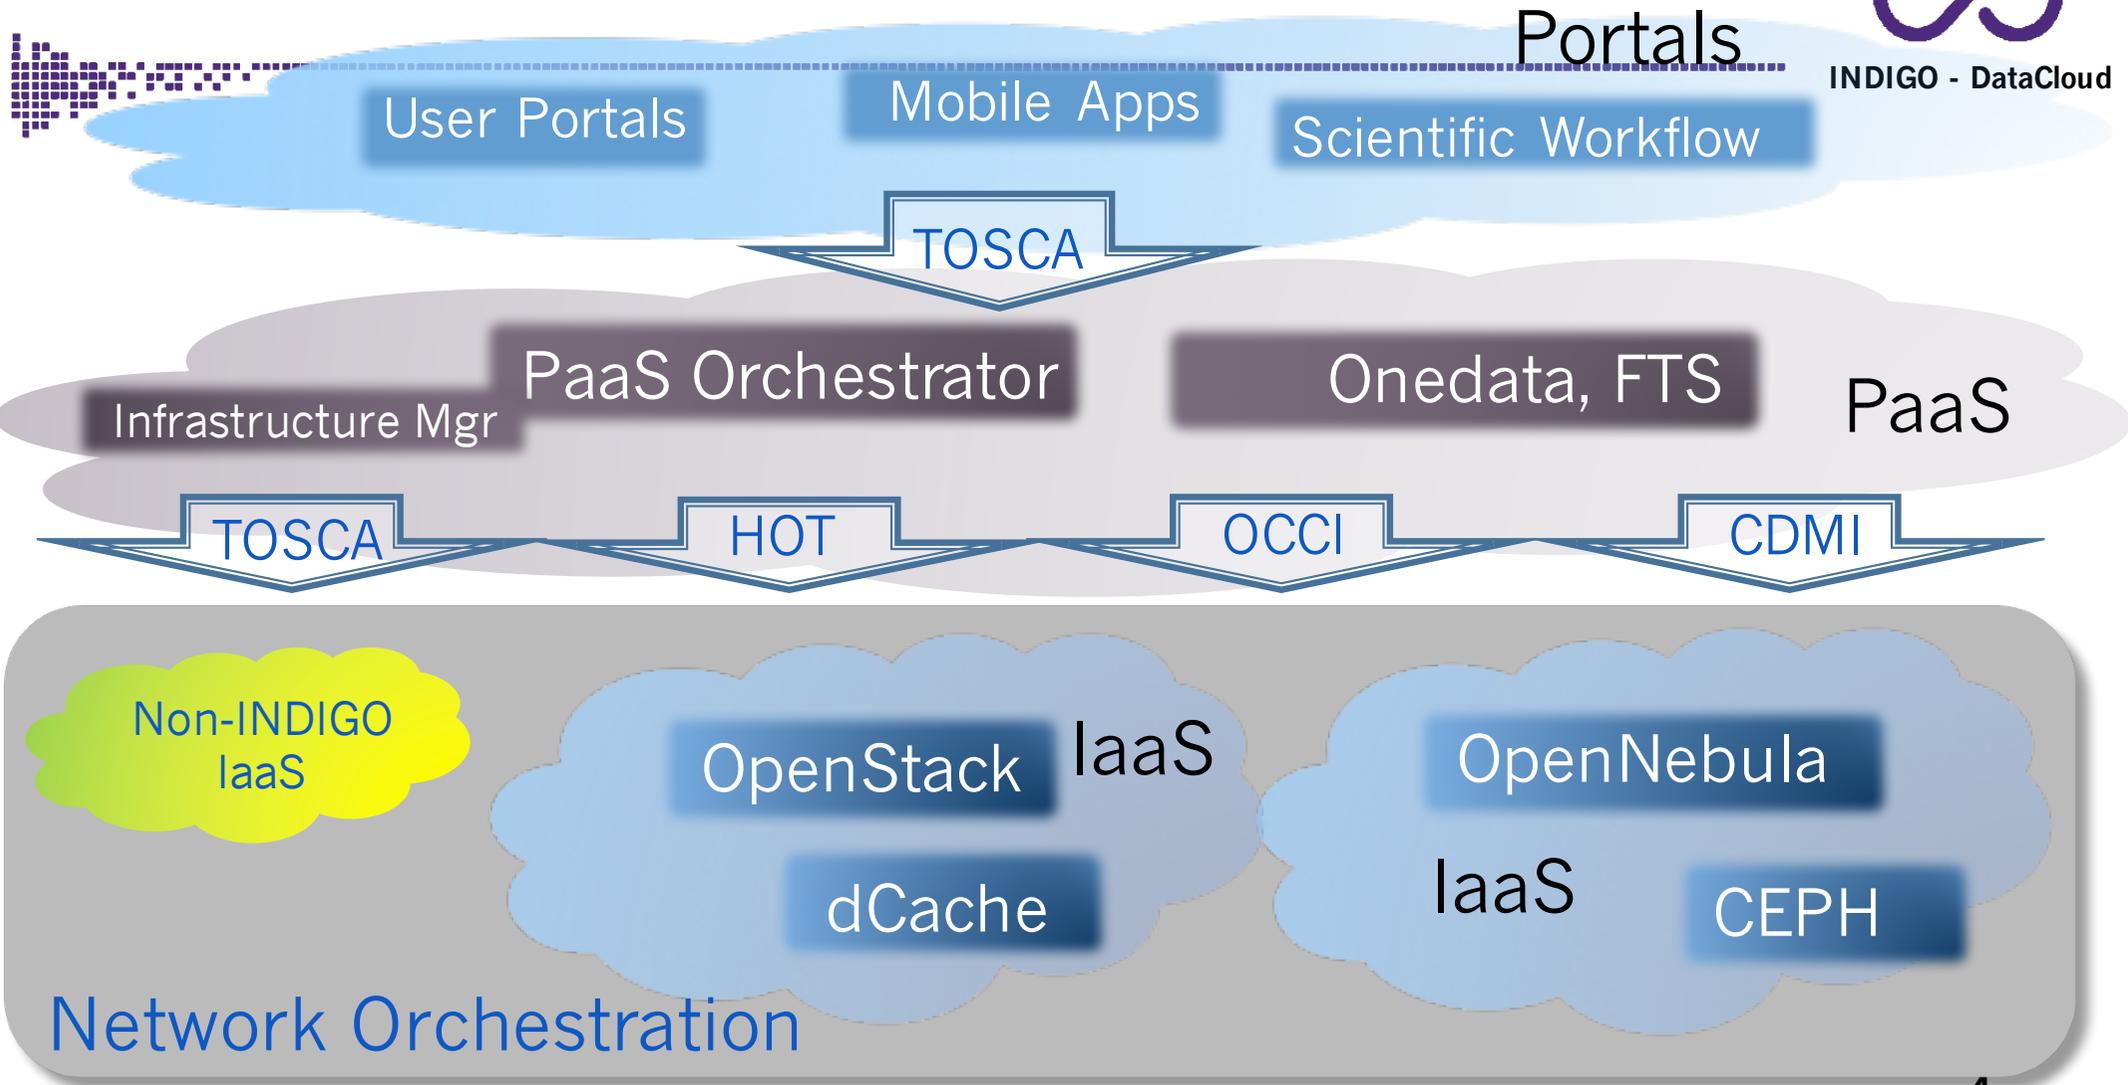


- In WP4 we develop components needed to guarantee that INDIGO solutions can be non-disruptively integrated into the major existing e-infrastructures (Grid, HPC and Cloud).
- In the context of the INDIGO software stack, WP4 components are the building blocks, the PaaS level is build upon.

IaaS in the INDIGO Context



INDIGO - DataCloud





INDIGO - DataCloud

High level overview

High Level Overview



INDIGO - DataCloud



- Providing Support for Containers in addition to the already existing VM support at the IaaS layer
 - Support for selected Cloud Management Frameworks.
 - Using *open-standard* interfaces.
 - With automatic repository synchronization.
- Improving on-demand compute capabilities at the local datacenters
 - Improving cloud scheduling, a mandatory requirement for production systems.
 - Facilitating container execution in HPC and HTC systems.
 - Providing site-level orchestration.

High Level Overview (cont.)



INDIGO - DataCloud

- Providing “Network Isolation” resp. “Multi Tenancy” for Virtual Machines.
 - Protecting user VMs from world access.
 - Protecting the world against maliciously acting virtual machines.
 - However, making VMs within a “Tenant” visible and discoverable.
 - Evaluating vendor support for SDN, as reported in INDIGO Deliverable D4.4.
- Providing a common framework to orchestrate Quality of Service in Storage across various Open Source and industry based storage technologies.
 - Including GPFS, CEPH, dCache, HPSS and StoRM
 - Setting up a European-wide Testbed with different types of storage endpoints.
- Consolidating AAI across the entire IaaS layer.
 - Integration of OpenID Connect when possible.
 - Using Token Translation when necessary.

High Level Overview (cont.)



INDIGO - DataCloud



Supporting Standards

- OCCI support in Compute and Networking for the supported CMFs
- Agreeing on a common vocabulary for QoS in storage across various science communities via RDA.
- CDMI Extension with SNIA to supporting Storage Quality of Service
- TOSCA language as the orchestration standard.
- OpenID Connect for common authentication, allowing *Single Sign On*.
- Standardized *Usage Records*, defined by the Open Grid Forum (OGF), for accounting.



Details on highlights



In INDIGO-1 release.



Development in progress, but not in release yet.



INDIGO teams actively involved in standardization process.

Container Support



INDIGO - DataCloud

- Container support in Cloud Management Frameworks
 - OpenStack (nova-docker)
 - Maintenance and packaging of nova-docker -> bug-fixes and improvement
 - OpenNebula (ONEDocker).
 - New driver development from scratch plus maintenance and packaging
- Container repository (DockerHub) synchronization
 - Generic WebService, responding to events from the central DockerHub service, by automatically synchronizing containers to local repositories.
 - New product developed: java-reposync
- Batch system container support
 - uDocker
 - Running Docker containers in user-space, without Docker at all.
 - BDocker
 - Manager for executing containers in batch-systems.
 - Under development, not yet released.



Improved Scheduling



INDIGO - DataCloud

- Synergy
 - General purpose manager for OpenStack, implementing FairShare scheduling, based on user, group and 'past usage' information.
 - SLURM "Multifactor Algorithm"
 - External product interacting with OpenStack Compute.
- OpenNebula FairShare Scheduler
 - Under development.
 - Will be accepted as an OpenNebula plug-in.
- Partition Director
 - Switches physical compute resources between a *Batch System* and a *Cloud Infrastructure*.
 - Reassigning shares and quotas.
 - Released in INDIGO-1 for the LSF batch system.
 - Next step is to provided the PD for other batch systems, e.g. HTCondor.



Improved Scheduling (cont.)



INDIGO - DataCloud

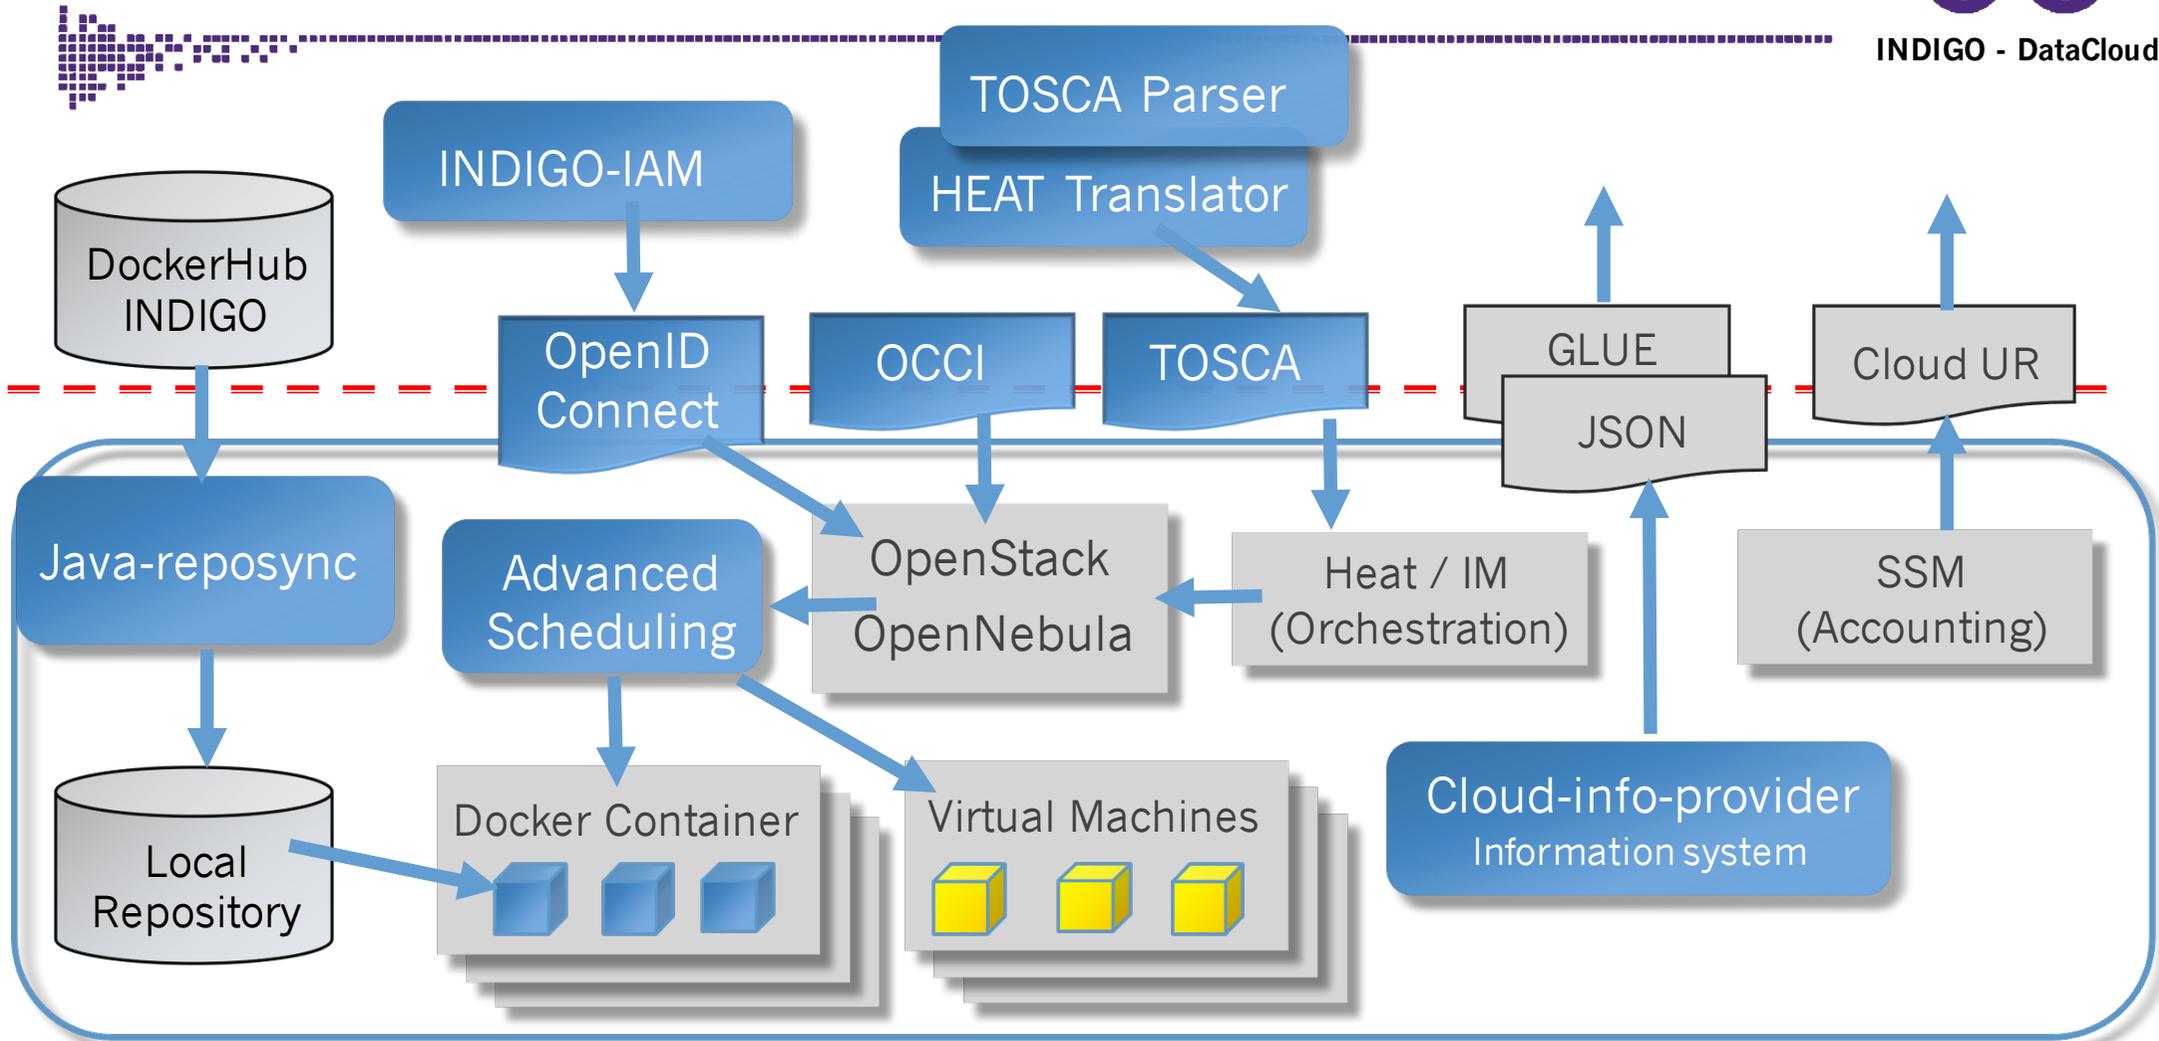


- OPIE: OpenStack *pre-emptible* Instances Extension
 - Pre-emptible instance model: interruptible by higher priority VMs.
 - Technical prerequisite to support spot instance markets.
 - Currently shuts-down VMs, but different actions can be chosen, e.g. suspending and resuming VMs.
 - Implementation as a pluggable scheduler + API extensions.
 - Working on upstream integration.

IaaS Design in Computing



INDIGO - DataCloud



Network Orchestration



INDIGO - DataCloud

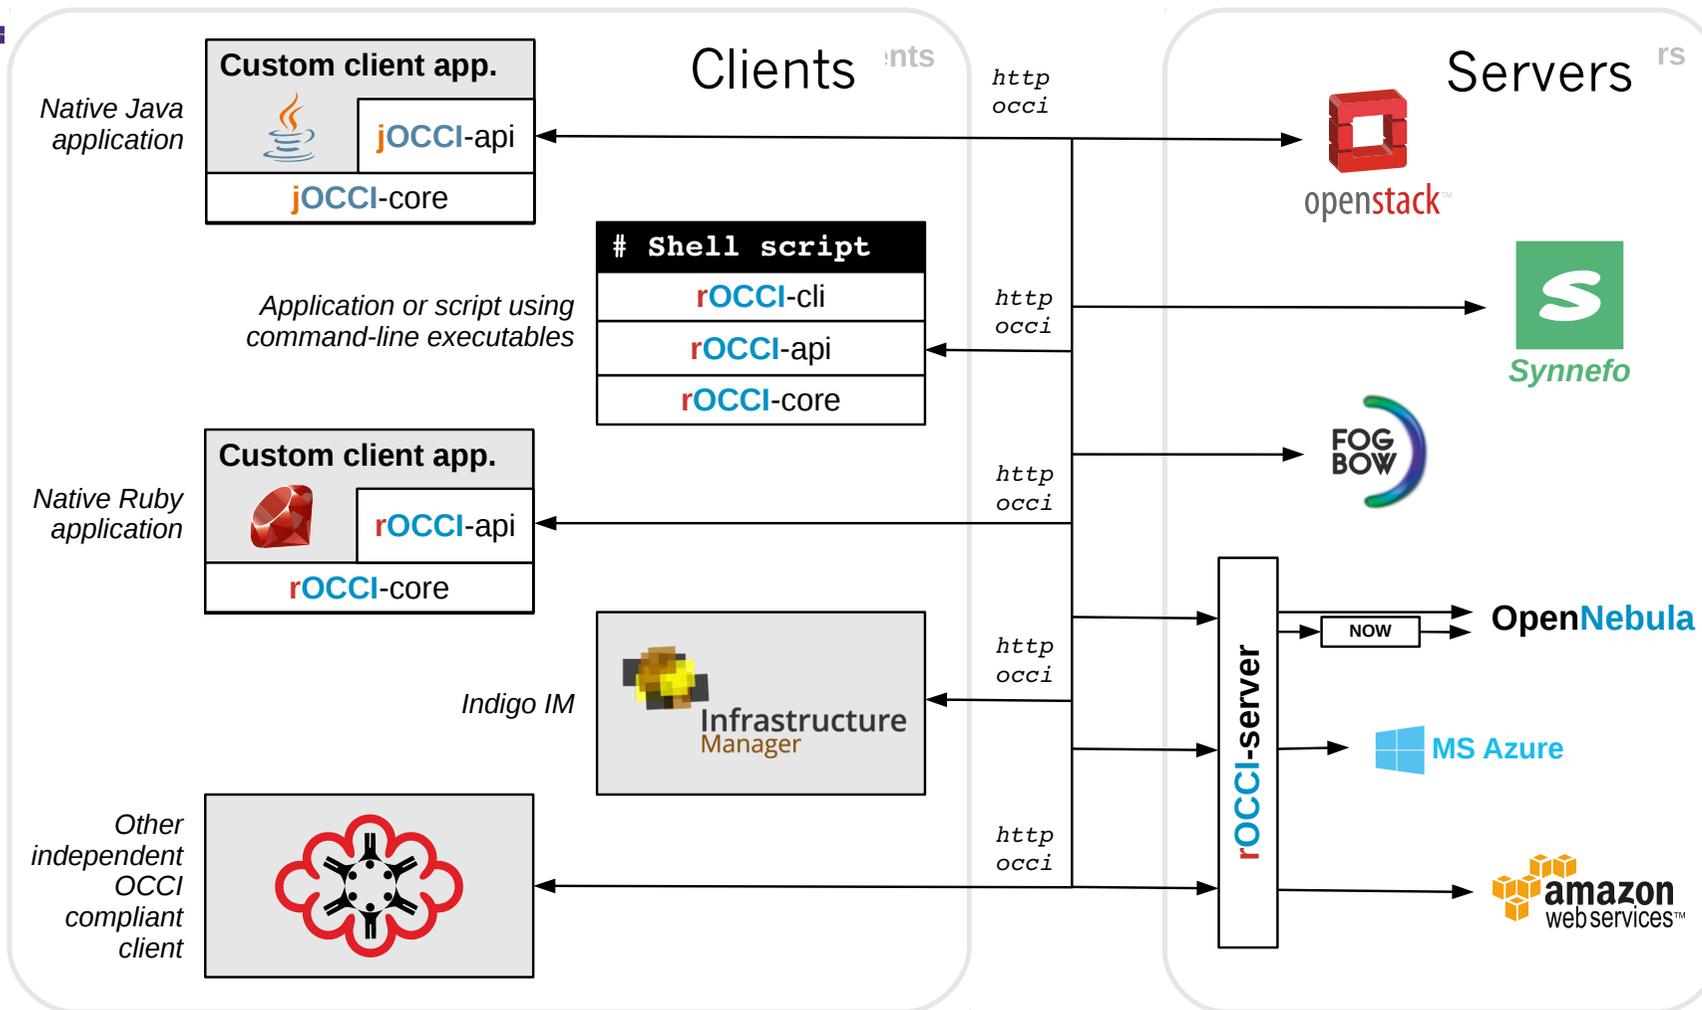


- Networking functions and orchestration for OpenNebula
 - Currently in alpha version.
 - **Network Orchestrator Wrapper**
 - Standalone component to orchestrate VXLAN-based networks in OpenNebula.
 - rOCCI-server backend for NOW to enable network management over OCCI.
 - Provides preliminary consistency checks before translating and forwarding the OCCI request to the local network infrastructure.

OCCI Layout for the network



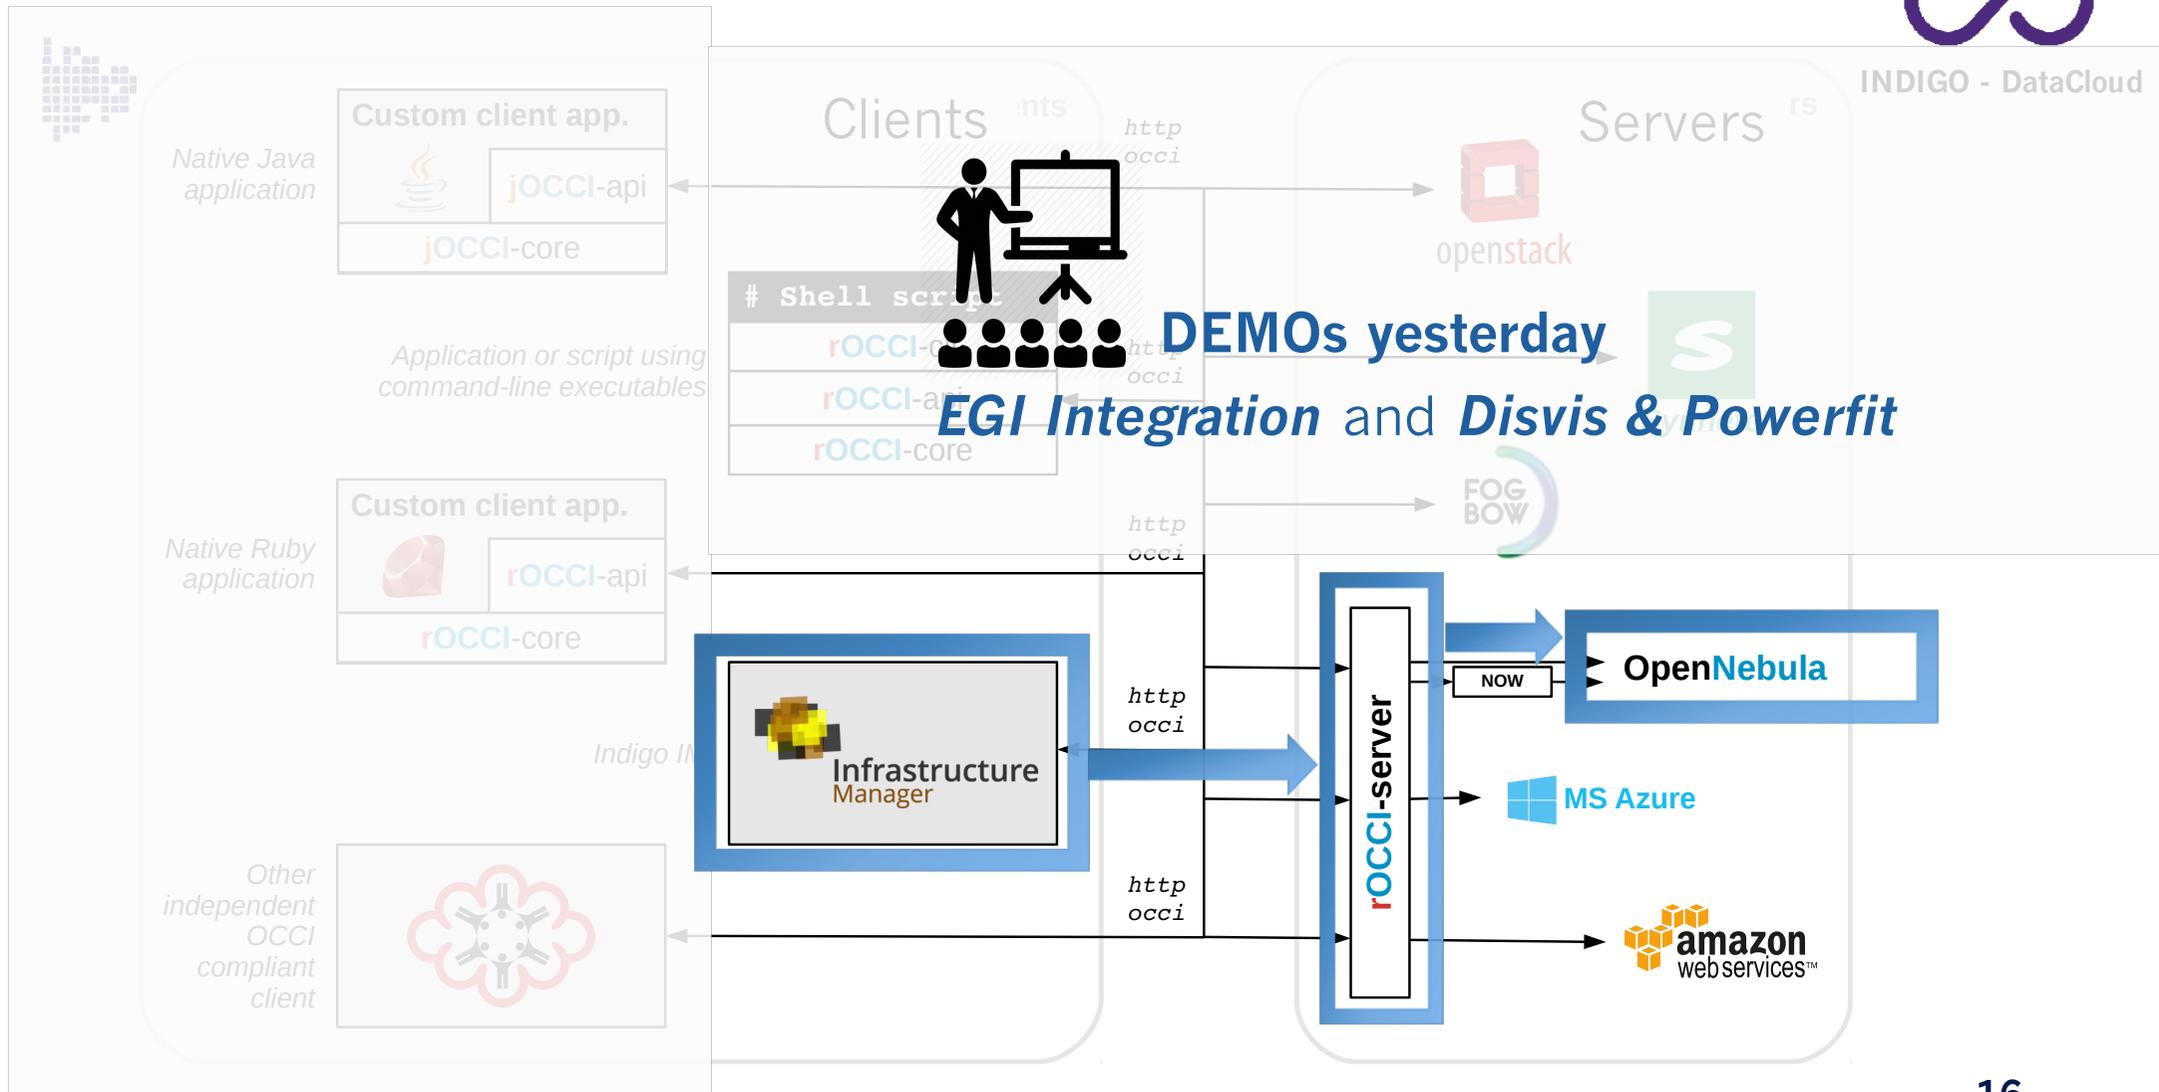
INDIGO - DataCloud



OCCI Layout for the network



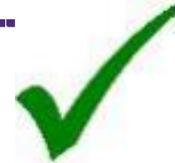
INDIGO - DataCloud



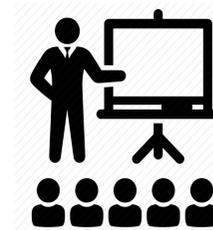
Authentication and Identity Management



iO - DataCloud



- Authentication based on OpenID Connect..
 - OpenID connected matches the INDIGO design requirements to a large extent and provides an industry standard.
 - OIDC is in line with the AARC blueprint architecture.
 - In consequence, INDIGO components, designed to work with the INDIGO IAM system, like the Token Translation service, can authenticate with other, similar, IAM systems, like EGI and Google.
 - OIDC Plug-Fest in preparation.



Demo
yesterday

Authentication and Identity Management



INDIGO - DataCloud

■ OpenID Connect support (Cont.)

■ Direct *OpenID Connect* support:

- OpenStack support through its native KeyStone tool.
 - Refactored Command Line Interface merged and released upstream.
- CDMI Framework System.
- dCache storage technology.

■ Support through the *Token Translation Service, TTS*

- *OpenNebula* : OIDC to Username/Password
- *Grid Systems* : OIDC to X509 Proxies via an Online CA.

Information System



INDIGO - DataCloud



- Cloud Info Provider

- Publishes resource provider information and capabilities.
- Augmented to implement a JSON rendering.
- Changes being merged into EGI's upstream info provider.
- The goal is to provide a single product for both types of rendering.
 - The traditional *LDIF* for Grid and
 - *JSON* for cloud systems.

Cloud Storage Virtualisation



INDIGO - DataCloud



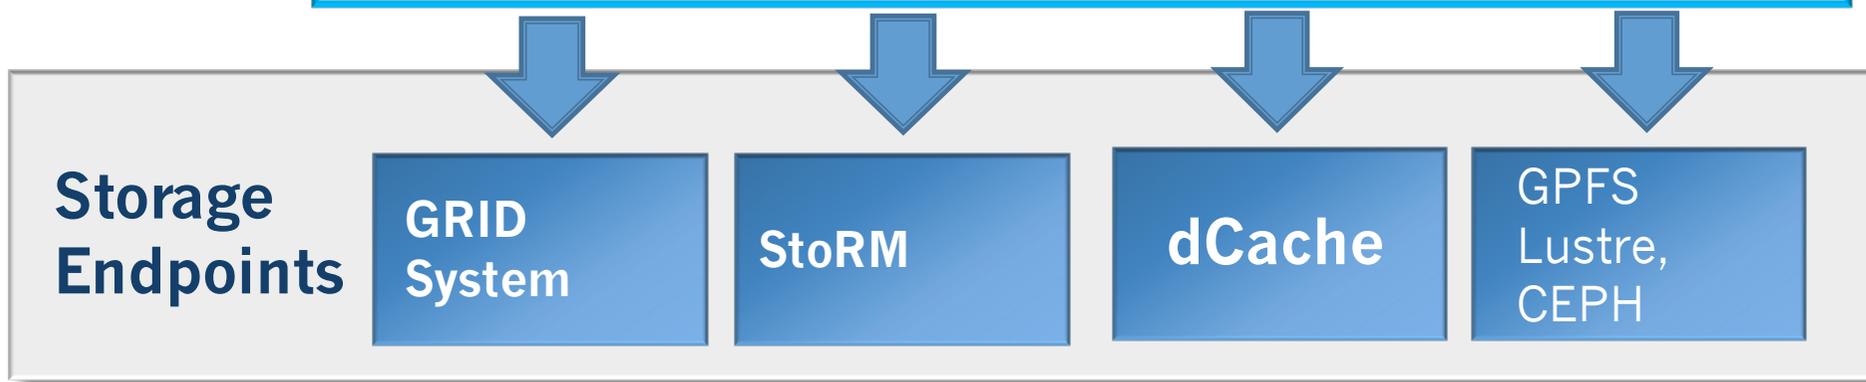
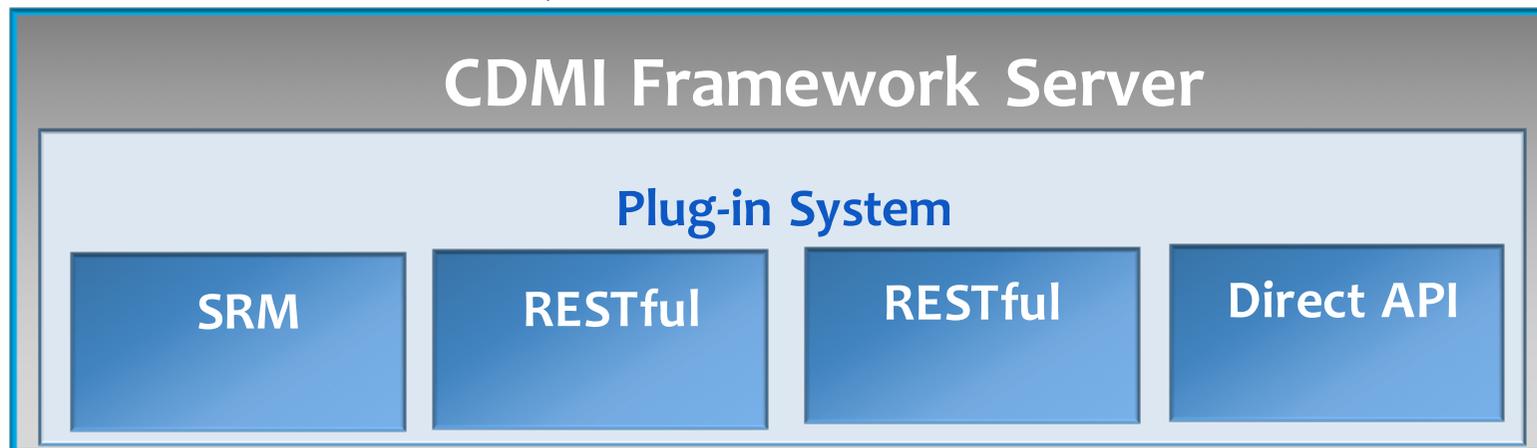
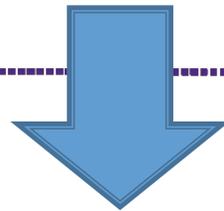
- Implementation of CDMI extensions
 - Information system to provide details about storage quality
 - Initial attributes supported:
 - **Access latency** (estimated time to first byte), **Location** of data (on a country-level granularity for legal purposes), **Number of copies** of data
 - List of **available storage qualities** into which a migration is available
 - Storage systems supported via plugins
 - TSM/StoRM, dCache, CEPH, HPSS
 - Modify storage quality
 - “Bring online” by moving between the available storage classes
 - Established a *European Testbed* with heterogeneous technologies.



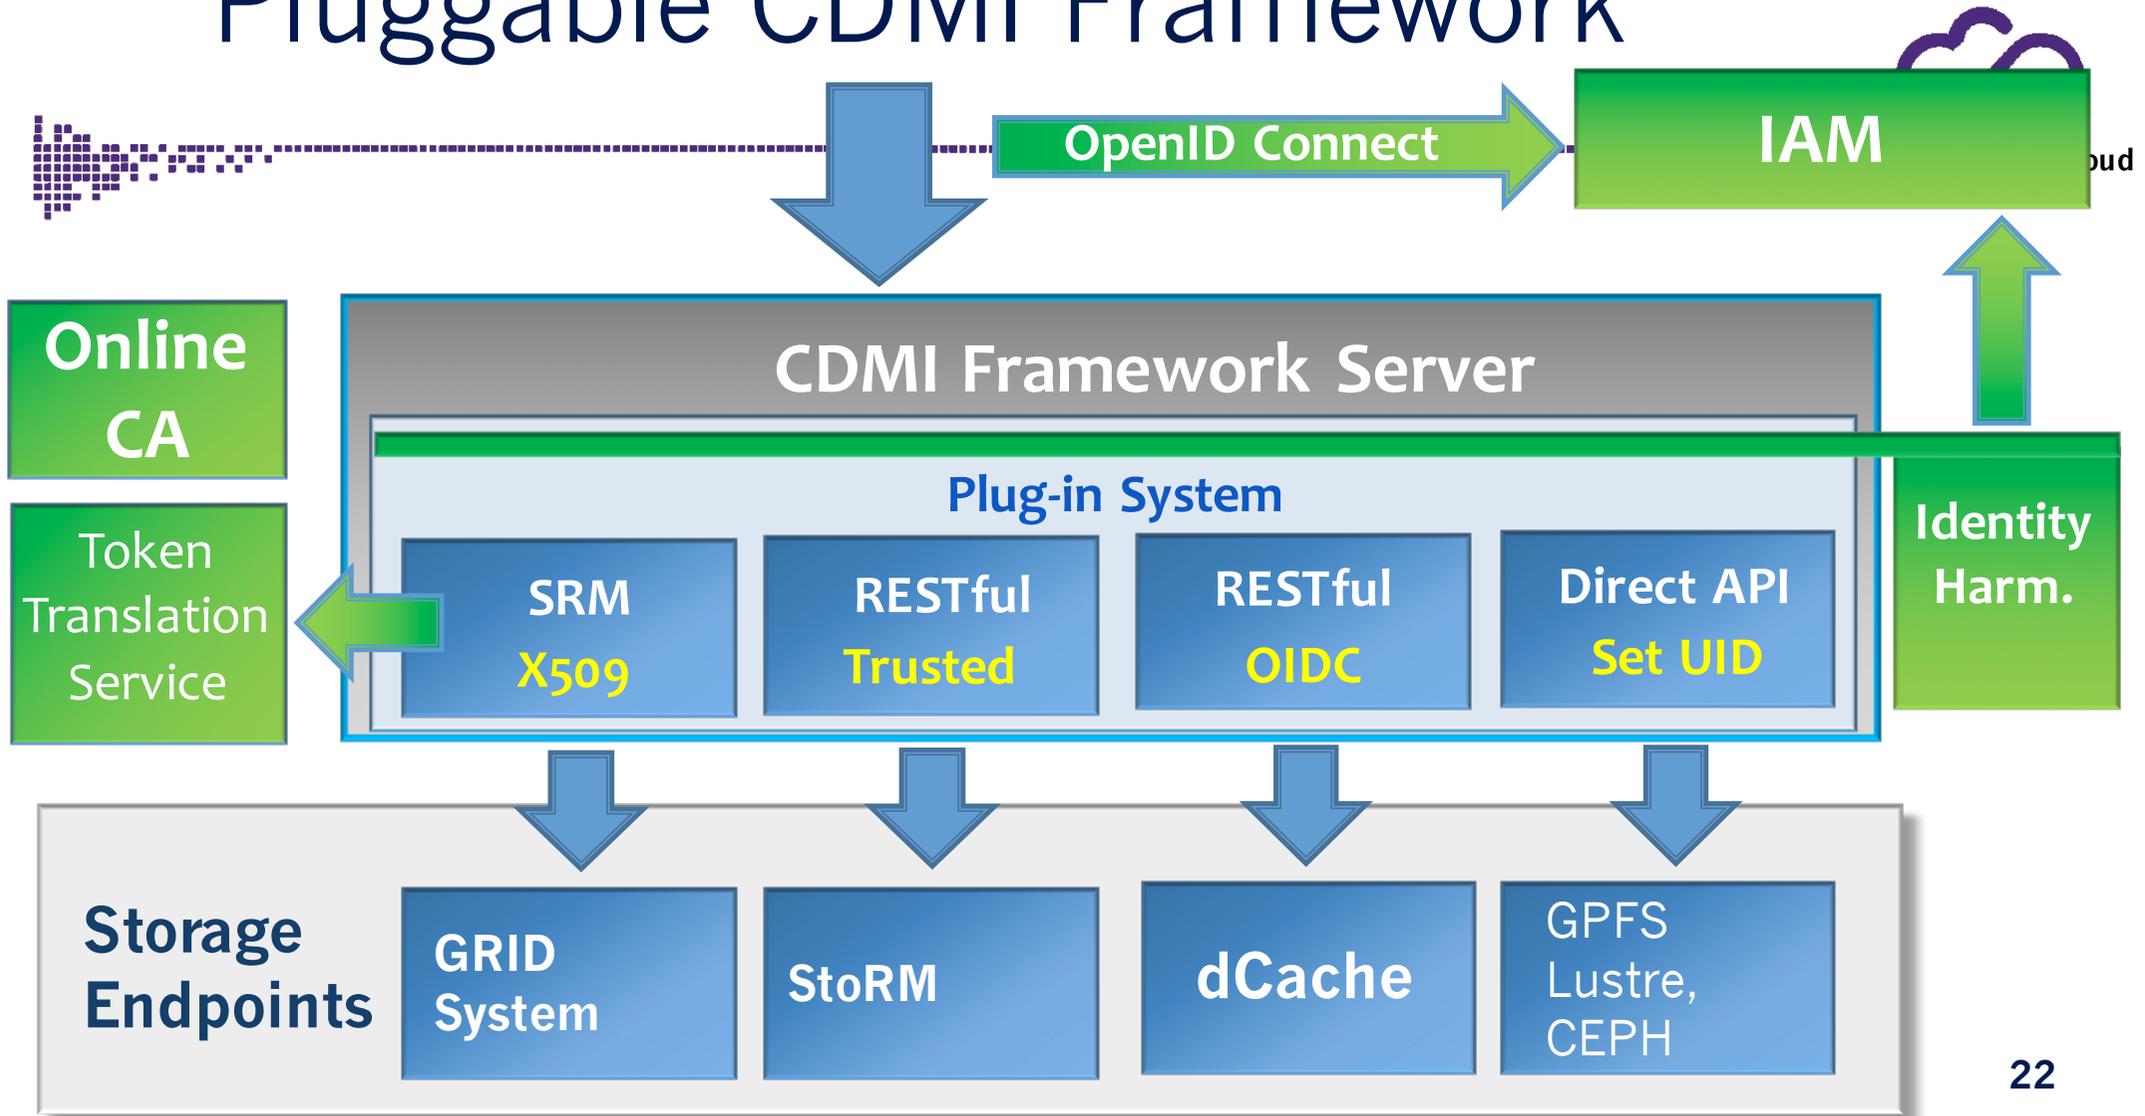
Pluggable CDMI Framework



INDIGO - DataCloud



Pluggable CDMI Framework



CDMI Media Capabilities

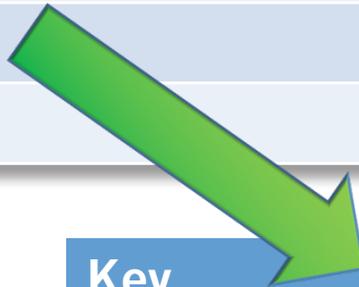


INDIGO - DataCloud

Object	Value
Path	/cdmi_capabilities/container
capabilities	{}
Children	DISK
	TAPE



Key	Value
Data_redundancy	1 Copies
Geographical Placement	"DE"
Latency	600000



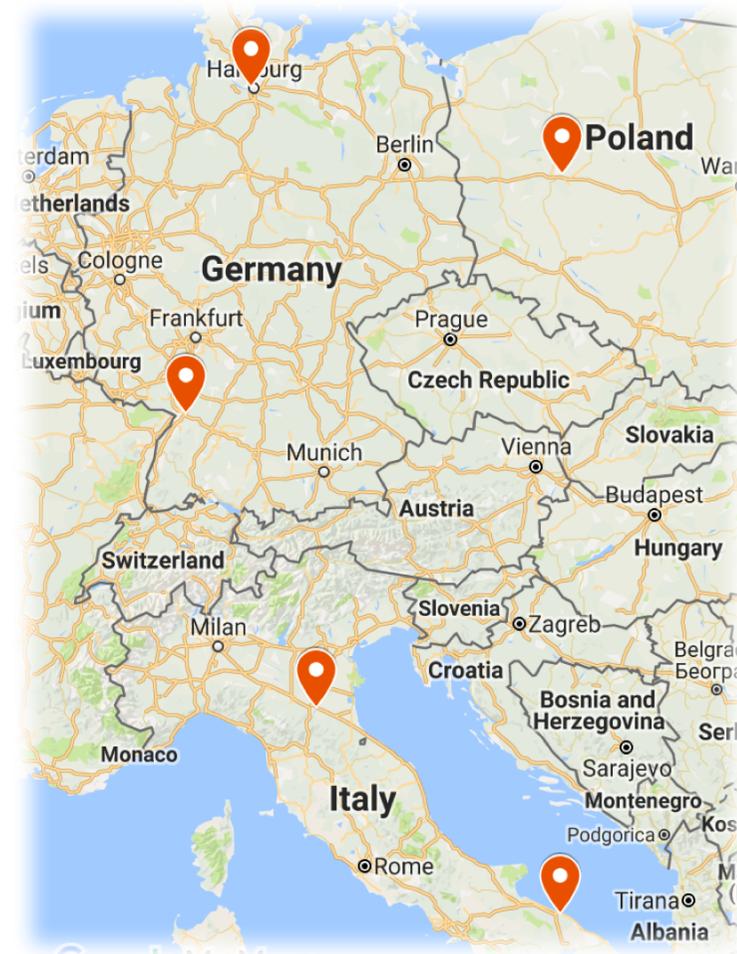
Key	Value
Data_redundancy	3 Copies
Geographical Placement	"DE"
Latency	100

Cloud Storage Virtualisation



INDIGO - DataCloud

- Implementation of CDMI extensions
 - (Including workflows, Staging)
 - Webservice at KIT continuously verifies 5 evaluation endpoints:
 - <http://seemon.data.kit.edu>
 - KIT (GPFS, HPSS)
 - BARI (GPFS, Ask Giacinto)
 - CNAF (GPFS, StoRM)
 - POZNAN (CEPH)
 - DESY (dCache)



Support of standards



INDIGO - DataCloud

- OCCI : IaaS management standard and interface.
 - “ooi” : include OpenStack networking changes.
 - Both rOCCI (ONE) and “ooi” support containers transparently.
 - Currently in Production at EGI Sites.
 - Networking functions in OpenStack OCCI Interface.
 - INDIGO involved in OCCI network definition OGF Working Group:
 - IP address reservation
 - Security Group Controls



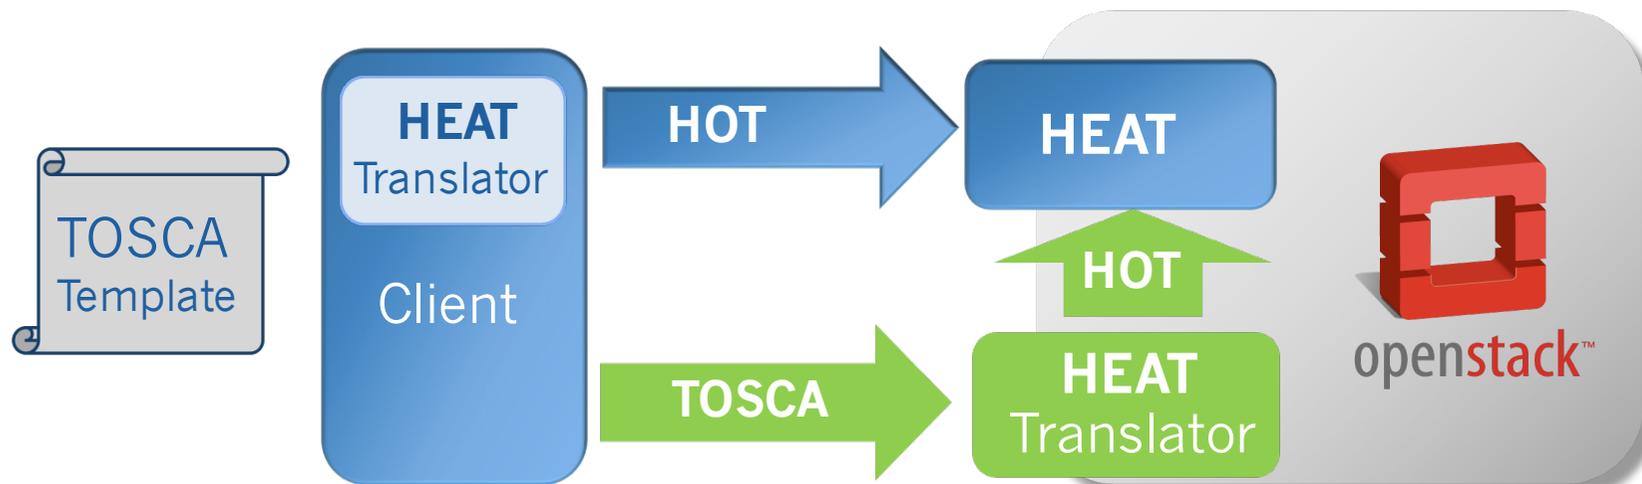
Support of standards (cont.)



INDIGO - DataCloud

■ 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*”.
- Currently in preparation by IBM : “*HEAT Translator as a Service*”



Support of standards (cont.)



- Survey available SDN options to enable spanning inter-site networks.
 - Recommendations and possible approaches for federated environments such as EGI Federated Cloud
 - INDIGO Deliverable D4.4 evaluates options and possible standards.
- Agreement on a common vocabulary for storage QoS across sciences under the umbrella of the Research Data Alliance.
 - INDIGO created a QoS in storage working group.
- Agreement on an extension of the CDMI protocol with the Storage Network Industry Association.
 - Member of the SNIA CDMI extension working group.



Support of standards (cont.)



INDIGO - DataCloud

- OGF Usage Record for accounting.
 - Accounting based on the OGF standard *Usage Record* (UR) resp. the *Cloud Accounting Usage Record*.
- Consistent usage of OpenID Connect authentication.
 - OpenStack via KeyStone.
 - OpenNebula via the INDIGO Token Translation Service.
 - CDMI Framework and dCache : direct OIDC
 -



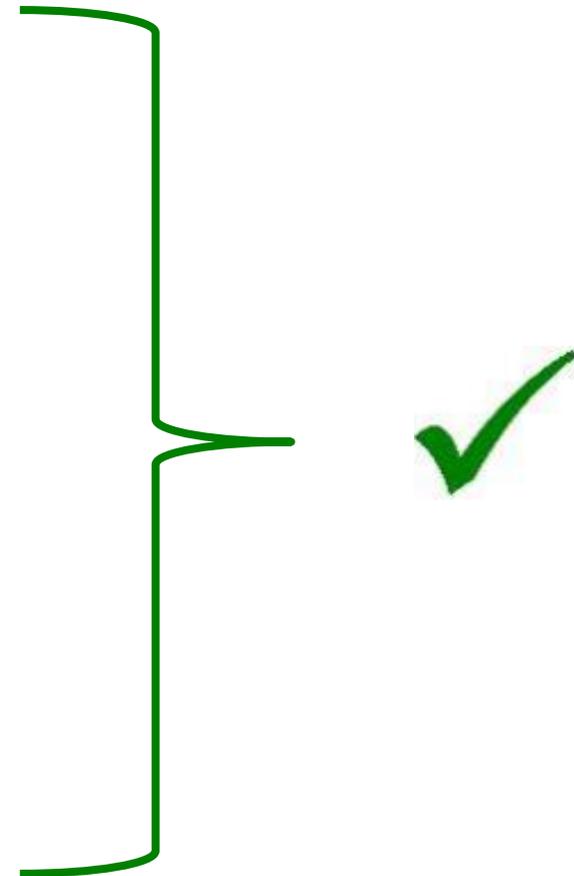
Contribution merged upstream



INDIGO - DataCloud

Essential in the context of sustainability and acceptance.

- OpenStack
 - NovaDocker
 - Heat Translator (On the client side)
 - OOI (OCCI for OpenStack)
 - Cloud info provider has been merged into EGI upstream code.
 - Pre-emptible Instance (in progress)
- OpenNebula
 - OneDock
 - rOCCI
- OpenID Connect
 - OpenID Connect via KeyStone (in the client side, via browser)
 - Lib Cloud authentication.
 - Erlang library
 - Java implementation for dCache.



Summary



INDIGO - DataCloud

- WP4, presenting the IaaS endpoints, is composed of a large variety of heterogeneous sites across Europe, representing a large set of scientific communities.
- In order to provide easy integration into the European infrastructures, WP4 is using standard technologies and protocols wherever possible and contributes to standardization efforts when necessary.
- WP4 puts significant efforts into contributing components to their official upstream repositories:
 - Increases acceptance by local site administrators.
 - Makes those components sustainable beyond INDIGO.

