Stato di INDIGO-DataCloud

Davide Salomoni, INFN-CNAF
(davide.salomoni@cnaf.infn.it)

Riunione CCR
Presidenza, 16/9/2015

Work Packages in INDIGO

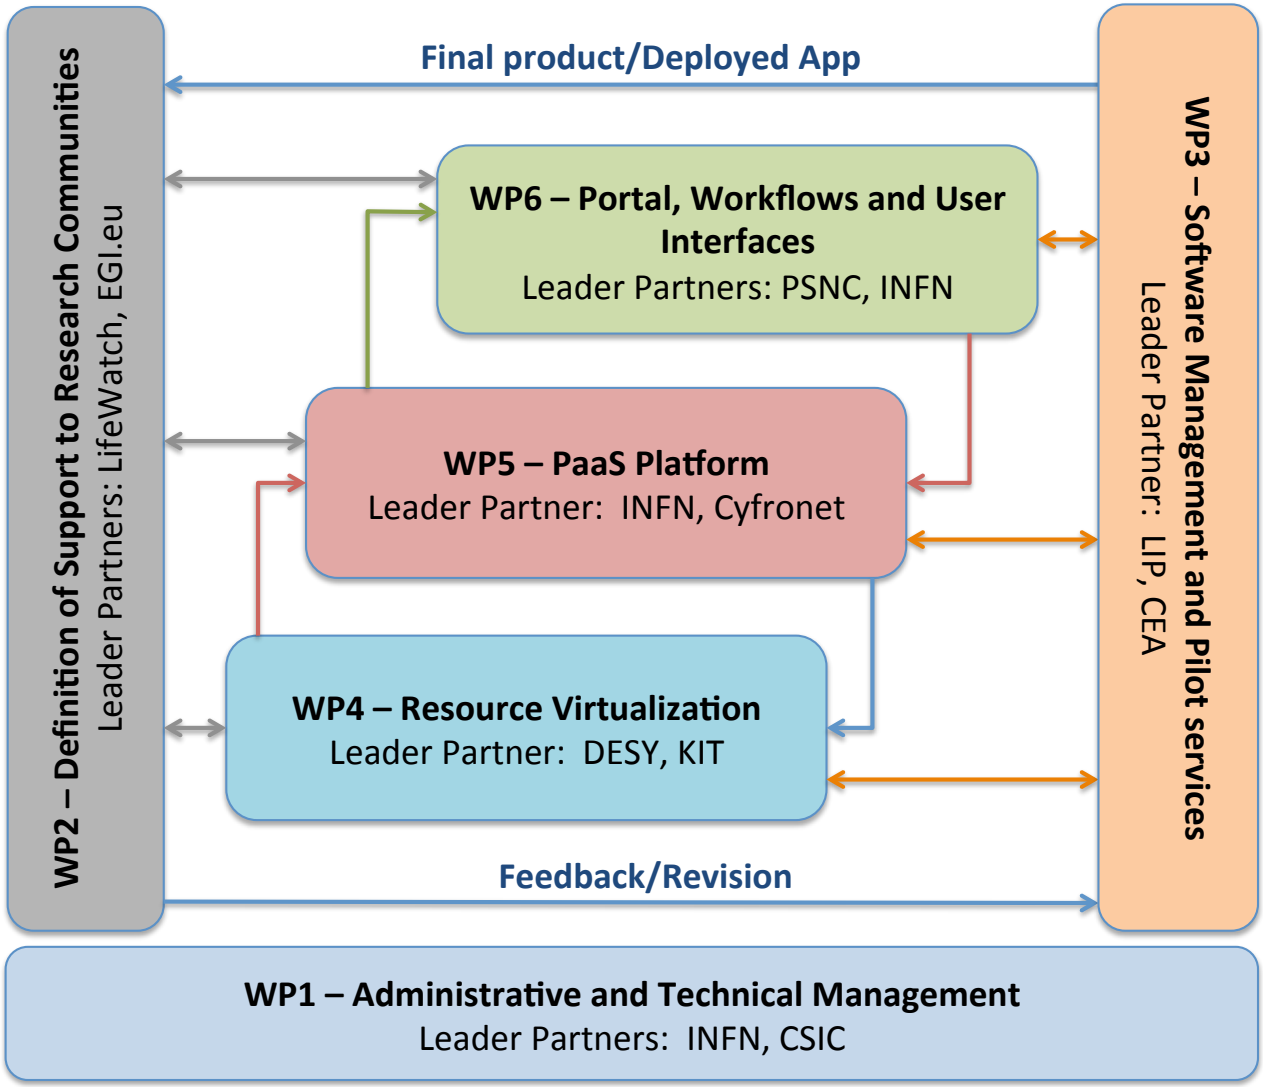

Cosa è successo dal report del 26/5 in CCR (1)

- Sono stati pubblicati 7 deliverables
  - Human resources allocation (confidential)
  - Quality plan
  - Initial requirements from research communities
  - Consolidation of dissemination plan, including project website and promotion
  - Risk analysis and risk response
  - Initial plan for WP3 (Software management and pilot services)
  - State of the art for the technologies at the base of PaaS architecture

Cosa è successo dal report del 26/5 in CCR (2)

- Altri 5 deliverables verranno pubblicati entro settembre 2015
  - Communication plan (confidential)
  - General architecture
  - Software design document for WP4 (compute, storage, network virtualization)
  - Detailed work plan of WP4 activities for the duration of the project
  - Design document and work plan for the PaaS architecture (WP5)
  - Software architecture and work plan for WP6 (user interfaces)
- Trovate tutti i deliverable pubblici su https://www.indigo-datacloud.eu/pages/components/deliverables.html

Cosa è successo dal report del 26/5 in CCR (3)

- Architectural meeting del 7-8/7/2015 a Valencia
  - “Fondamentale per definire le interazioni inter-WP e l’architettura di dettaglio delle componenti di INDIGO” (dal report di maggio in CCR)
- Definizione dell’architettura di base di AuthZ (cross-WP)
- Prossimi eventi con presenza INDIGO:
  - RDA Europe, 22-24/9, Parigi
  - ISC BigCloud, 28-30/9, Francoforte
  - HEPiX, 12-16/10, BNL
  - H2020 ICT2015, 20-22/10, Lisbona
  - OpenStack Summit, 27-30/10, Tokyo (TBC)
  - e-Concertation Meeting, 9-10/11, Bruxelles
  - EGI Community Forum, 10-13/11, Bari
  - SuperComputing 2015, 15-20/11, Austin TX
  - Liaison prevista con progetto H2020 HNSciCloud
- Internal face-to-face meetings
  - 6-7/10, Bologna (WP4)
  - Fine ottobre, Roma – TBC (WP5)

La domanda nasce spontanea…

"Ma questi fanno qualcosa
oltre a scrivere pezzi di carta
e ad andare in giro per il mondo?"

Requirements from research communities (WP2)

"The proposal is oriented to support the use of different e-infrastructures by a wide-range of scientific communities, and aims to address a wide range of challenging requirements posed by leading-edge research activities conducted by those communities." (dal testo del proposal di INDIGO)

- Raccolti casi d'uso da 11 diverse comunità scientifiche
  - LifeWatch, EuroBioImaging, INSTRUCT, LBT, CTA, WeNMR, ENES, eCulture Science Gateway, ELIXIR, EMSO, DARIAH
- Da una lista di circa **100** requirements diversi necessari a supportare questi casi d'uso è stata derivata una lista molto più corta, divisa nelle 3 categorie di Computational requirements, Requirements linked to storage, Requirements on infrastructures, con un ranking associato (mandatory / convenient / optional)
- Questo lavoro è stato essenziale per la definizione della architettura dei WP di sviluppo

Simplified impact table

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

#REQ	Description	Type	Rank	Proposed Improvement
CO#1	**Deployment of Interface SaaS**	Computing / PaaS	M	A mechanism to facilitate the deployment of a customised Haddock portal and backend in system in a panoply of infrastructures with minimal intervention.
CO#2	**Deployment of Customized computing back-ends as batch queues**	Computing / PaaS	M	Each instance may have an independent software configuration, potentially incompatible with other projects or specially tailored without side-effects.
CO#3	**Deployment of user-specific software**	Computing / PaaS	M	Manual installation may be cumbersome for large-scale application involving many computing resources or when requesting users to update VMIs. This should be automated.
CO#4	**Automatic elasticity of computing batch queues**	Computing / PaaS	M	When moving to the cloud, users should be provided with the exact number and size of resources they need. Overprovisioning will produce an undesirable cost or inability to serve other requests. On the other side, underprovisioning will lower QoS.
CO#5	**Terminal access to the resources.**	Computing / PaaS service	M	This feature must be linked to the AAI
CO#6	**Privileged access**	Computing / PaaS service	M	This feature must be linked to the AAI
CO#7	**Execution of workflows**	Computing / PaaS	M	Processing done on the cloud where the outputs of the processing are stored. Orchestration of complex pipelines.
CO#8	**Provenance information**	Computing / PaaS Service	C	Very important for revision of papers and project proposals.
CO#9	**Cloud bursting**	Computing /	M	Supplementing the computing capacity with special

Common use cases: excerpt

7.2 Generic Solution B: A Data Analysis Service

This generic solution example is based on the need of a user community that has a coordinated set of data repositories and software services to access, process and inspect them. Processing is interactive, requiring accessing a console deployed on data's premises. The application consists on a console / Scientific Gateway that interacts with the data:

- Examples include "R", Python, Ophidia
- It can be a complementary scenario from the previous one.
- It can expose programmatic services.

Requested by the Case Studies from INSTRUCT, INAF-LBT, CMCC-ENES, LifeWatch-Algae–Bloom, EMSO–MOIST.

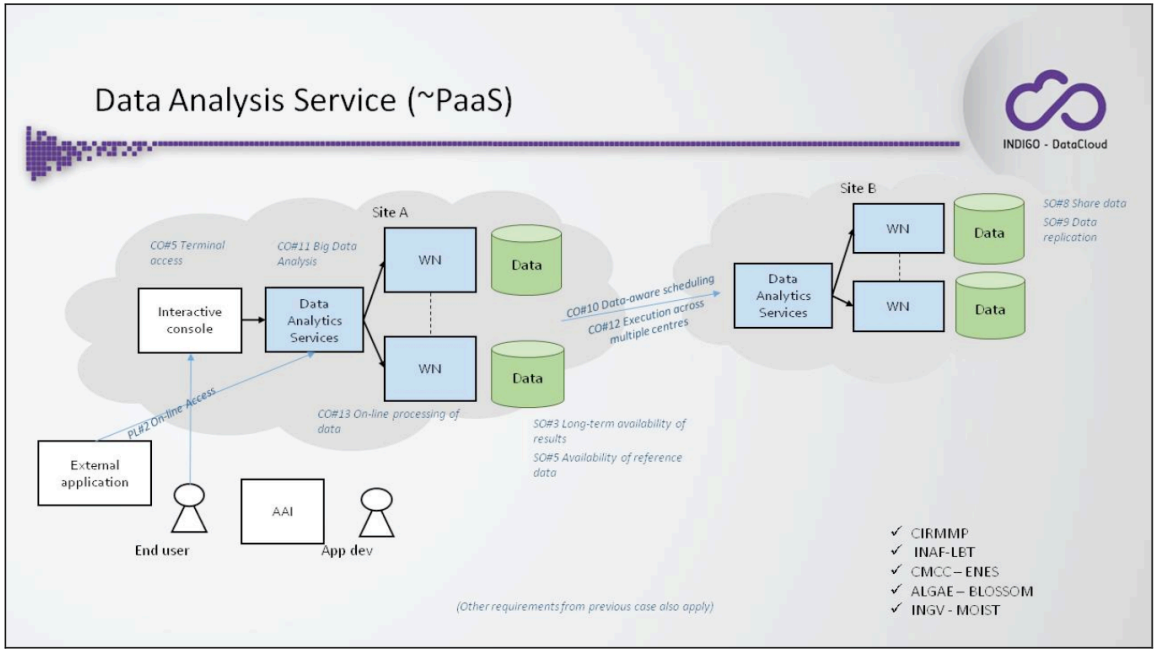

Figure 6: An example of potential generic solution: Data Analysis Service

Requirements: Analisi pratica

Software management and pilot services (WP3)

INDIGO-DataCloud è un progetto che produce software: da qui l'importanza della definizione chiara della produzione, gestione e testing del software stesso

- Abbiamo definito i processi di software lifecycle, i tool per la Quality Assurance (QA) del software, le metriche, i ruoli, le responsabilità e la metodologia di misura per QA, le procedure di software release, maintenance e di supporto, la definizione dei testbed di integrazione e di testing

Software and pilot processes

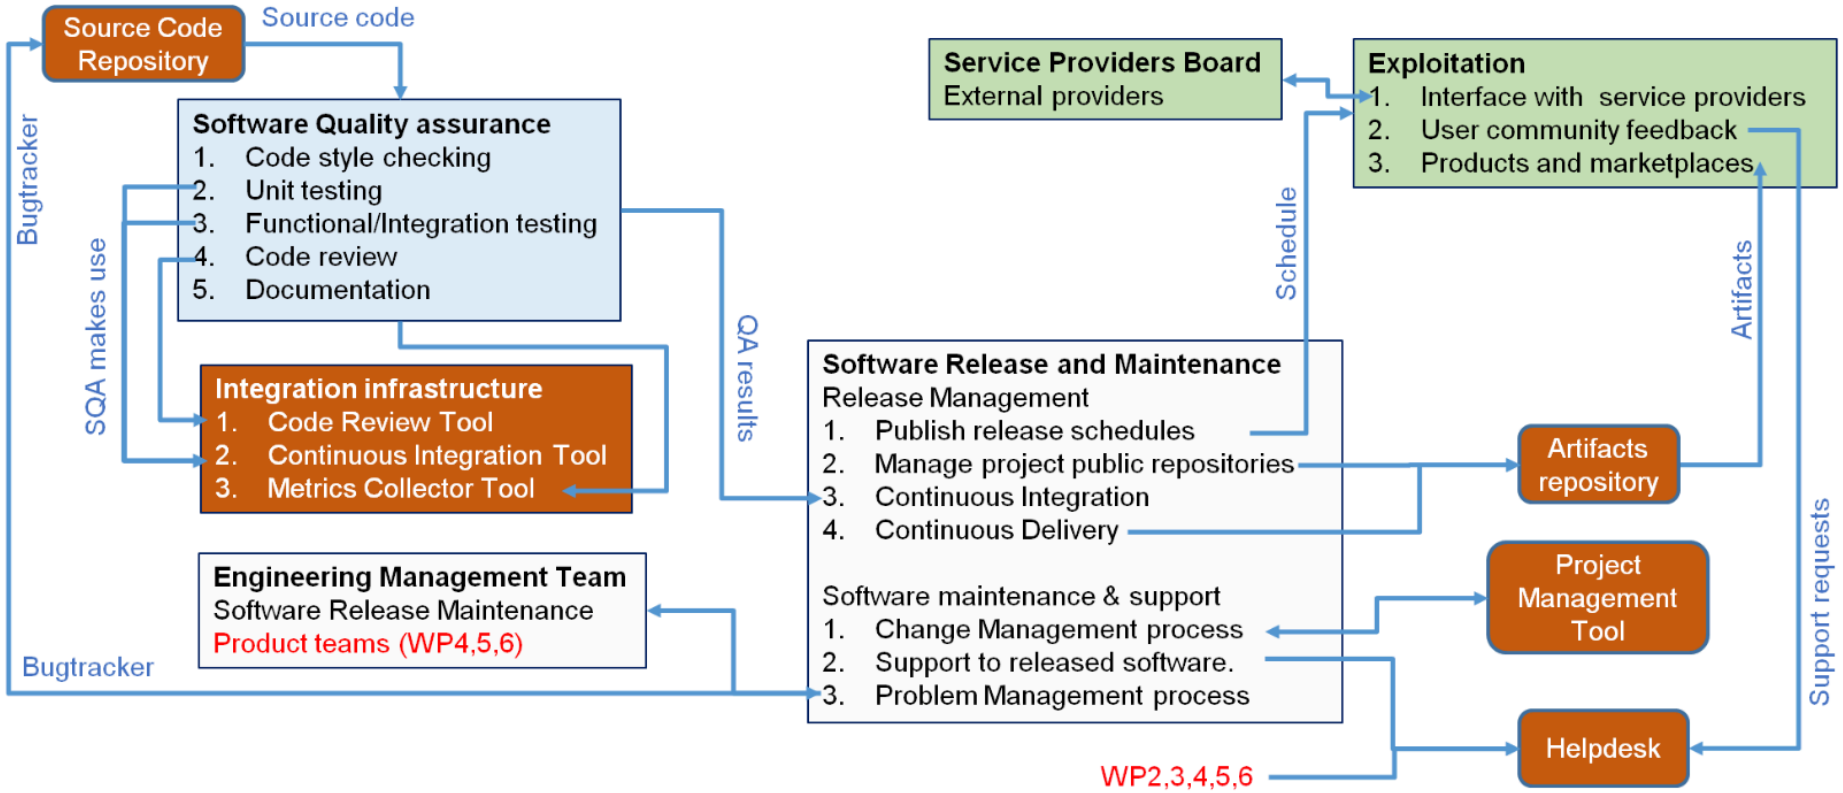

Figure 3: Overview of the process defined by the WP3 initial plan

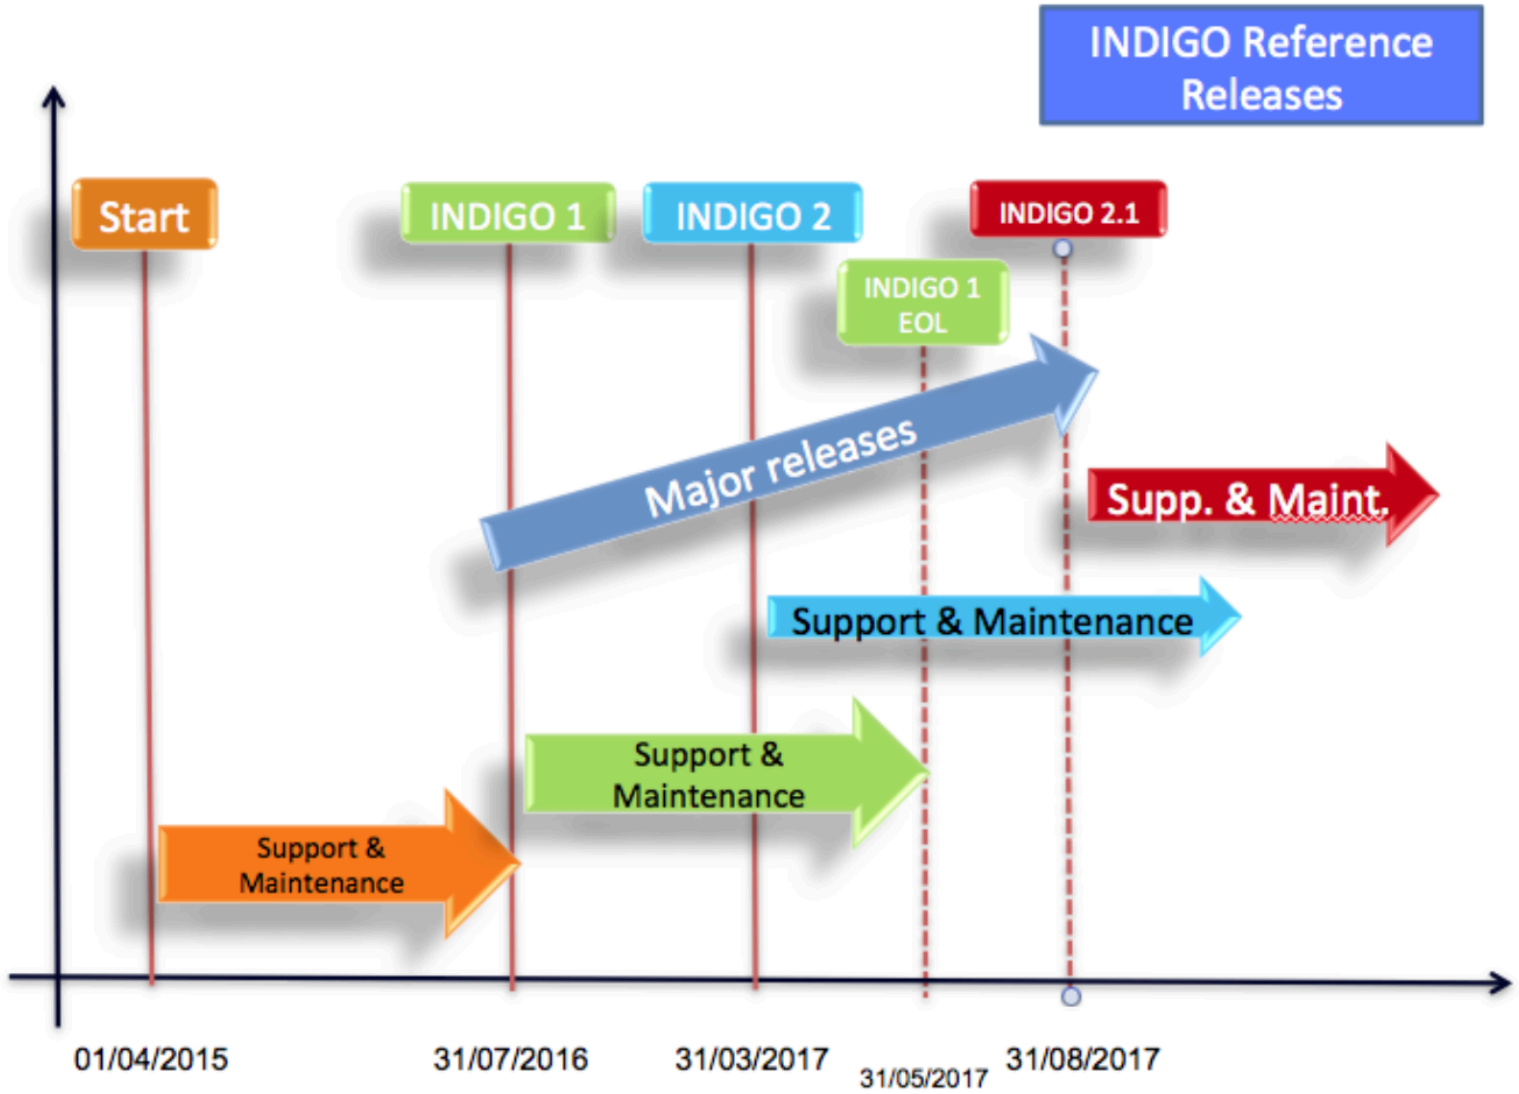

Figure 8: INDIGO-DataCloud Release Timeline

SQA and integration services

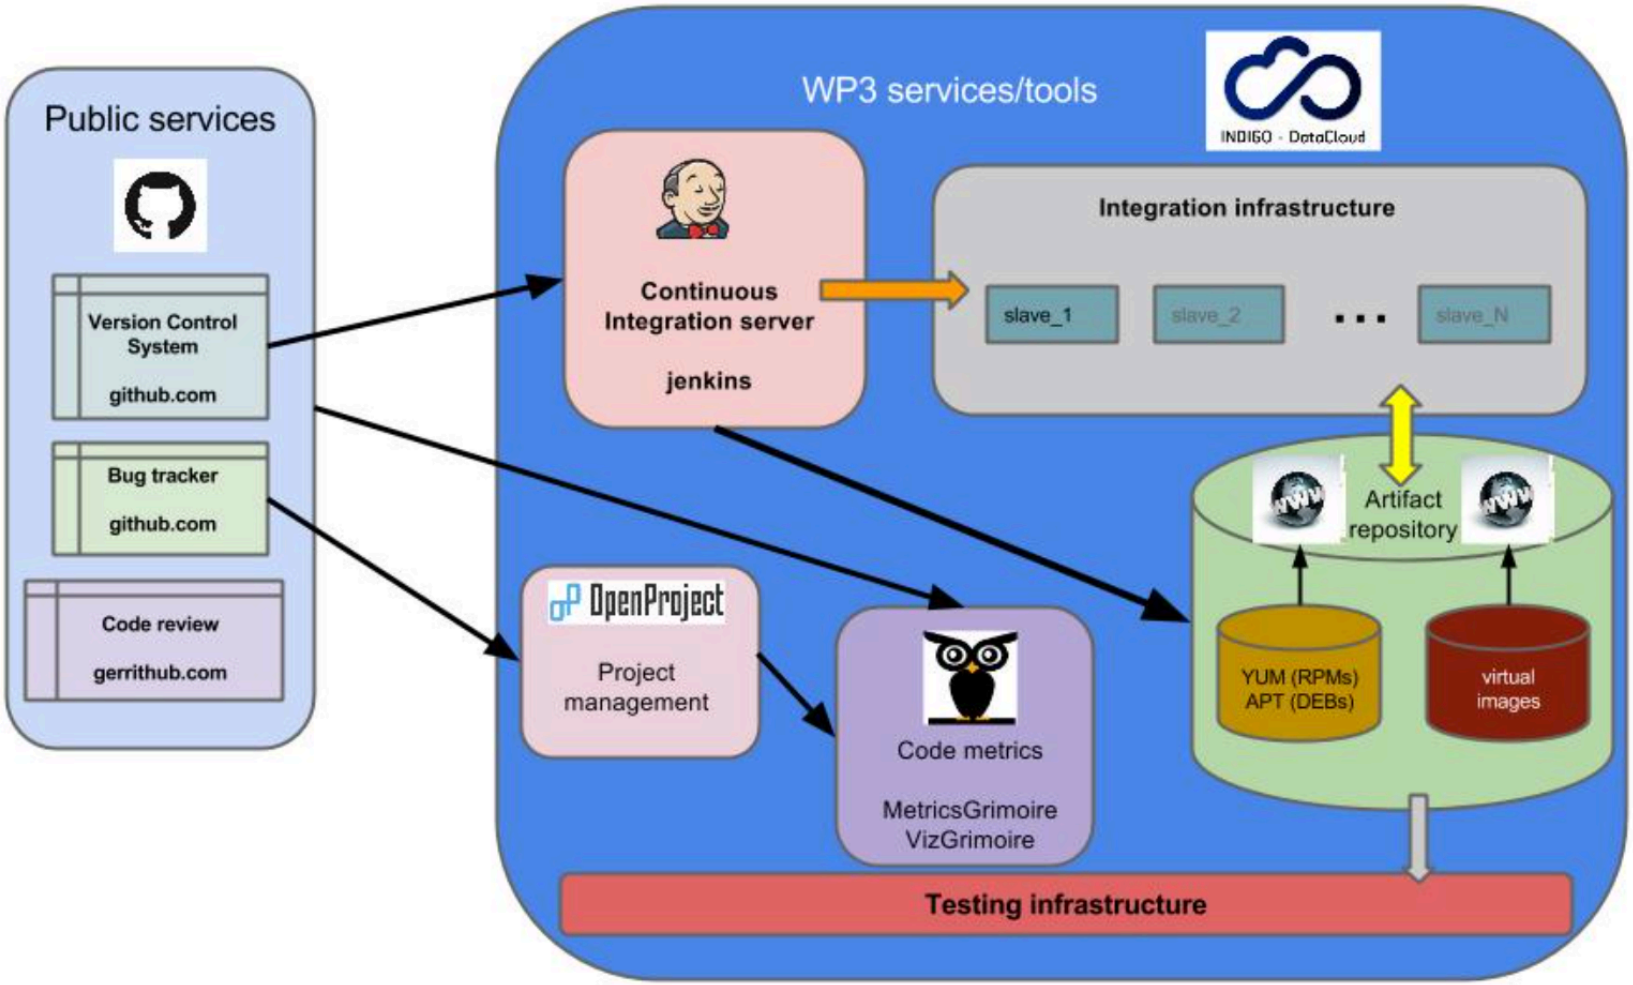

Figure 12: Overall architecture of the services provided by WP3.3 for the SQA and Continuous Integration process.

AuthZ architecture

I componenti sviluppati da INDIGO necessitano di una piattaforma comune per la gestione delle autorizzazioni: è stata creata una AAI task force cross-WP per la definizione dell'architettura di AuthZ (AuthN viene mutuata da tecnologie esistenti, in particolare SAML, OpenID Connect, X.509)

- AAI in INDIGO: fornitura di servizi per l'armonizzazione e la decorazione di informazioni di identità provenienti da IdP
  - Provide unique user identifier linked to various user identities
  - Enrich with additional information (e.g. group membership)
  - Support *constrained delegation* by design
    - i.e., the ability to delegate only a subset of privileges to an agent acting on a user's behalf
  - Integrate well with HTTP RESTful services
- Integrate with AARC (KIT – an INDIGO partner – is also directly involved with AARC)

INDIGO AAI from 10.000 ft

Identity information linked to a computational activity in the INDIGO AAI in maintained in a **session** living at one INDIGO Identity Service (called here **Login Service**) that is trusted by INDIGO services

In order to create a session, a user authenticates through one of the supported mechanisms (SAML, X.509, OIDC) at the Login Service

As the outcome of the creation of a session, the user receives an **INDIGO-token** linked to that session

- The INDIGO-token is a bearer token, and grants access to the identity information stored in the session living at the Login Service

The INDIGO-token is then included in calls to INDIGO services

INDIGO services use the iNDIGO-token to obtain user identity information related to specific computational activities from the Login Service

Delegation

The INDIGO-token is a bearer token

- any agent presenting a valid token gets access to the related session at the INDIGO Login Service

Delegation is implemented handing out the INDIGO-token to an agent or a service acting on the user's behalf

INDIGO-token can be **constrained** before delegation by the agents, so that their validity can be attenuated

- e.g., to limit the validity of the token only for the next 10 minutes and/or to limit its validity only if presented from a specific IP address

Constrained delegation along the chain of agents is possible since the INDIGO-token is a *Macaroon* token

- See http://research.google.com/pubs/pub41892.html for details

WP4, scelte architetturali

Draft

- Cloud Middleware Frameworks: OpenStack and OpenNebula.
  - OCCI implementations available.
  - HEAT (OpenStack) and Infrastructure Manager (OpenNebula) + TOSCA to provide orchestration
  - Initial container support in Stackforge for OpenStack; ATOS, UPV and INFN-Torino will develop container support for ONE.

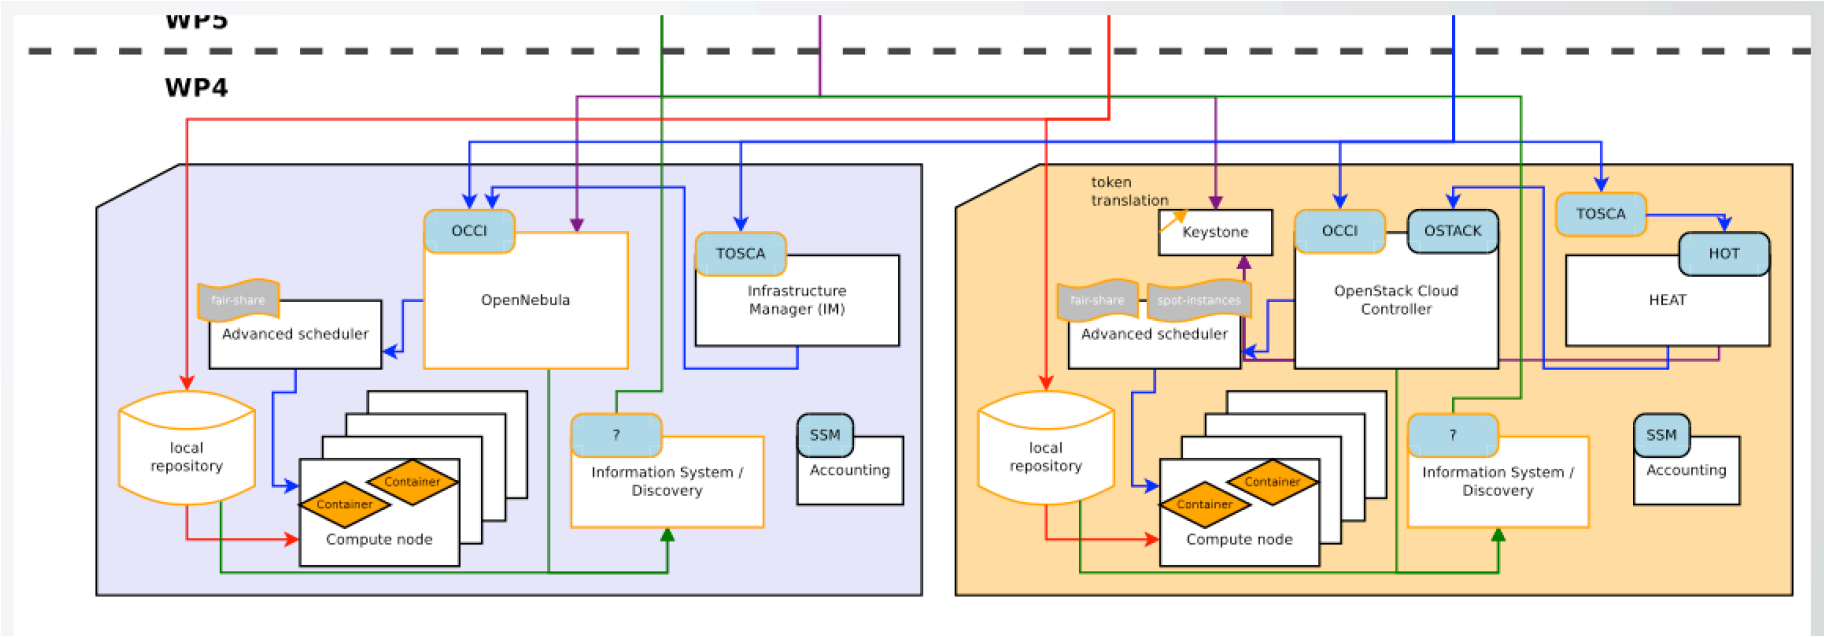

Compute virtualization

- Support for containers (will initially leverage Docker)
  - Local cointainers catalogs, pulling containers from the INDIGO repository
  - Allow execution in batch systems; PoC accessing InfiniBand and GPUs
- Improve cloud scheduling with fairshare scheduling and support for spot instances (can be terminated by a higher priority task)

Storage virtualization

- Storage Management Control Protocol: decision was to go for CDMI, as follows:
  - Agree on terminology for storage QoS and DLC with RDA (we'll be at RDA-Europe next week also for this reason)
- Agree on ways to implement those findings in CDMI with SNIA
- Implement CDMI for QoS and DLC
  - Full Implementation for dCache (in contact with DPM people)
  - Agreement on CDMI front-end for different back-ends (GPFS, StoRM, TSM, HPSS)

PaaS platform development (WP5)

- PaaS core architecture and implementation
- Security and Authorization
- High-level geographical application/service deployment
- Unified Data Access

- Use standards as much as possible to ensure interoperability

On standards

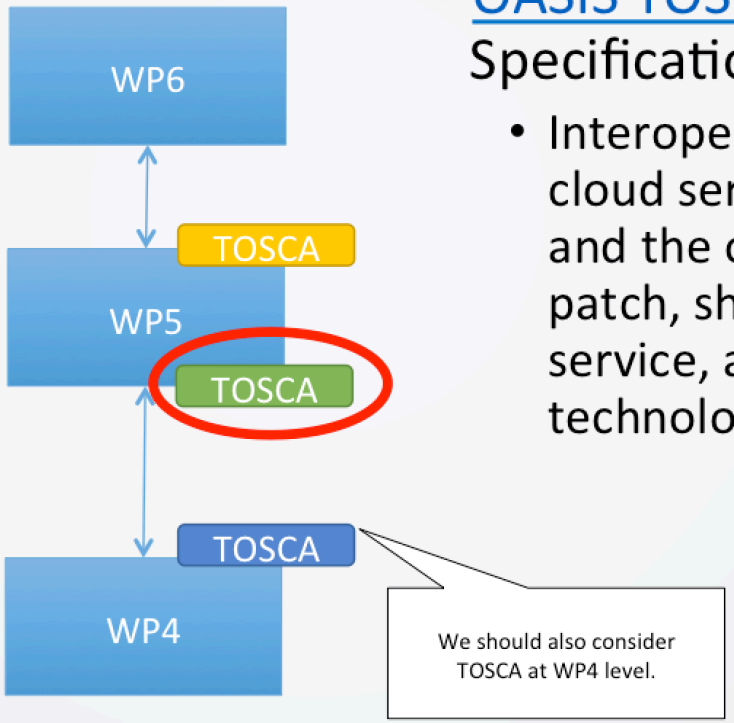

- OASIS TOSCA (Topology and Orchestration Specification for Cloud Applications) 1.0 (11/2013)
  - Interoperable description of application and infrastructure cloud services, the relationships between parts of the service, and the operational behavior of these services (e.g., deploy, patch, shutdown) independent of the supplier creating the service, and any particular cloud provider or hosting technology.

100+ participants from 40+ companies:

Capgemini, ca technologies, IBM, CISCO, EMC², WSO2, SAP, CITRIX, NetApp, redhat., pwc, software AG

Source:

Interfaces between Layers of INDIGO Architecture

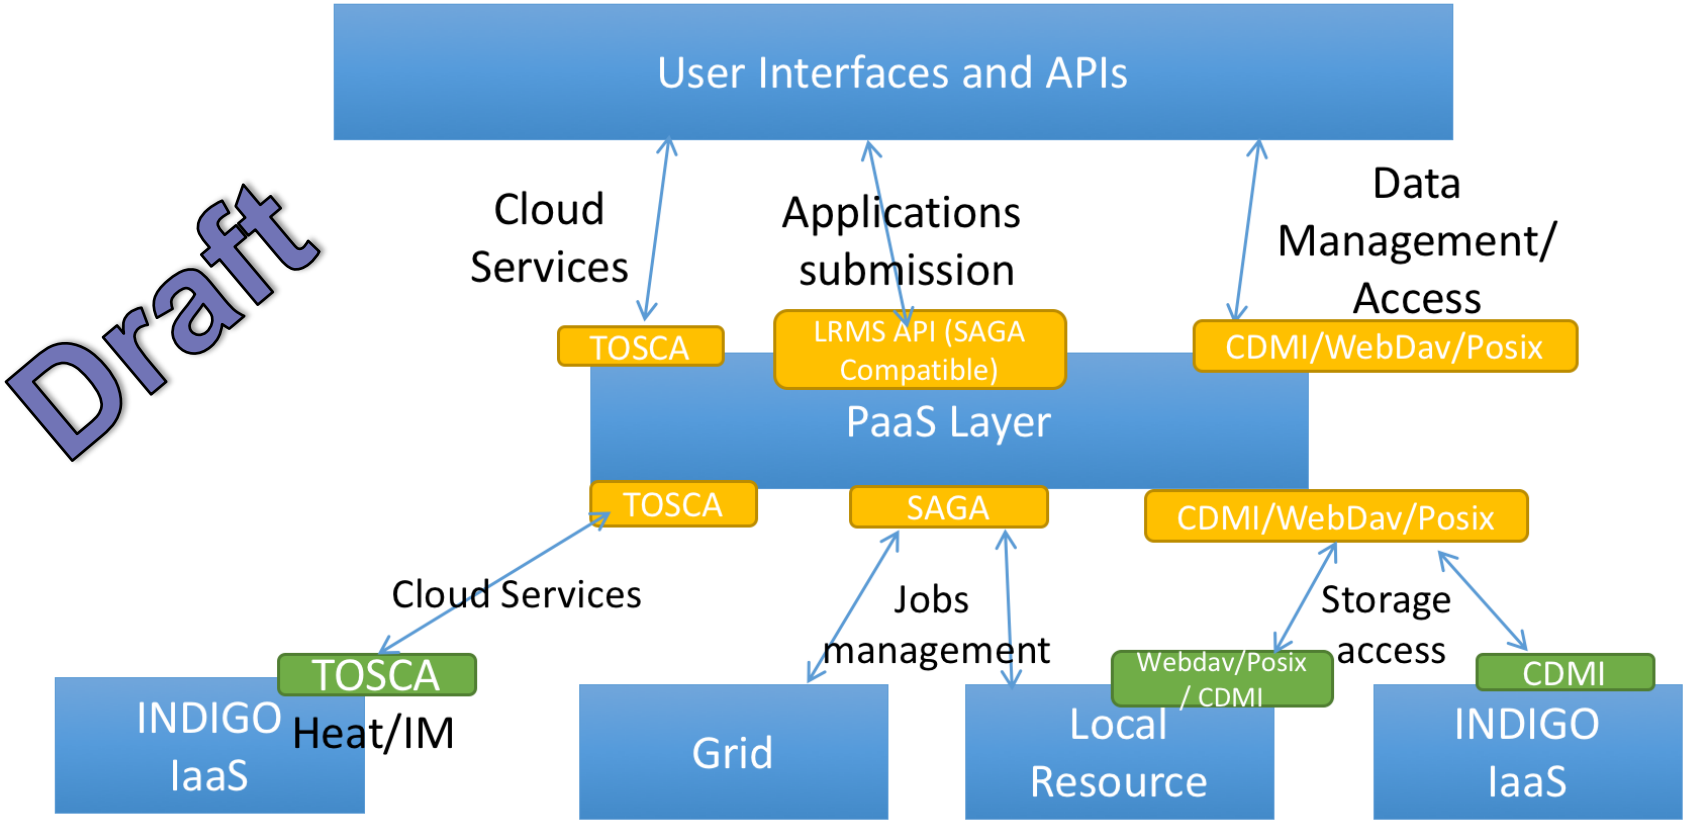

INFN

Overall architecture

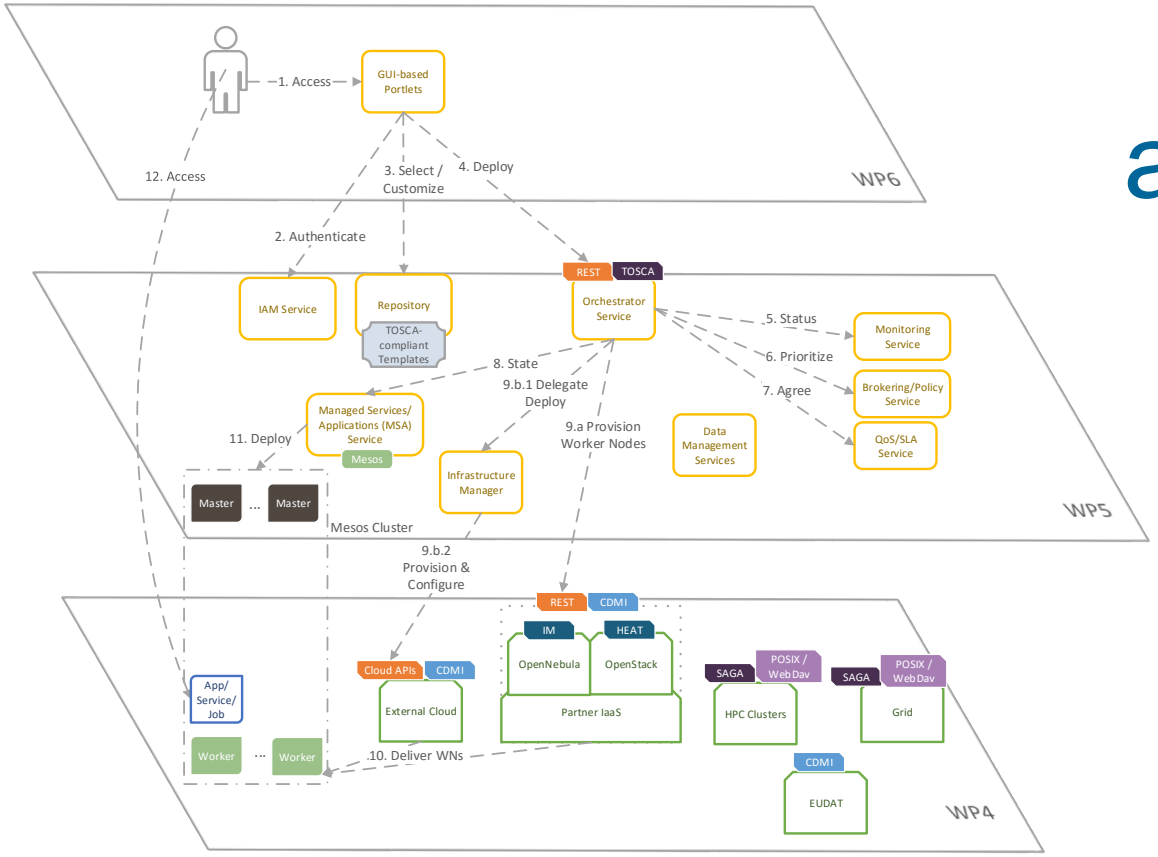

Draft

Portals and user interfaces (WP6)

"Guarantee a simple and effective end user experience, both for software developers and for researchers running the applications."

- Develop Toolkits (libraries) to allow the platform usage at the level of the Scientific Gateways, desktop and mobile applications
- Develop the Open Source Mobile Application Toolkit for iOS, Android and WindowsPhone platforms that will be the base for development of the Mobile Apps.
- Provide User Friendly front-ends proving the usability of the INDIGO PaaS:
  - Provide a Science Gateway and customized examples for selected user communities/ scenarios, that will make use of the proposed Toolkits, including Data Analytics Gateways for e-Science
  - Develop example cross platform native Mobile Apps for selected use cases, based on the Mobile App Toolkit
- Manage the execution of complex workflows using PaaS layers
  - Support for both interactive and batch parallel data analytics workflows
  - Provide the dynamic scientific workflows services in a Workflows as a Service model
  - Provide workflow interfaces extensions for distributed and parallel data analytics on large volume of scientific, multidimensional data).

WP6 overall architecture

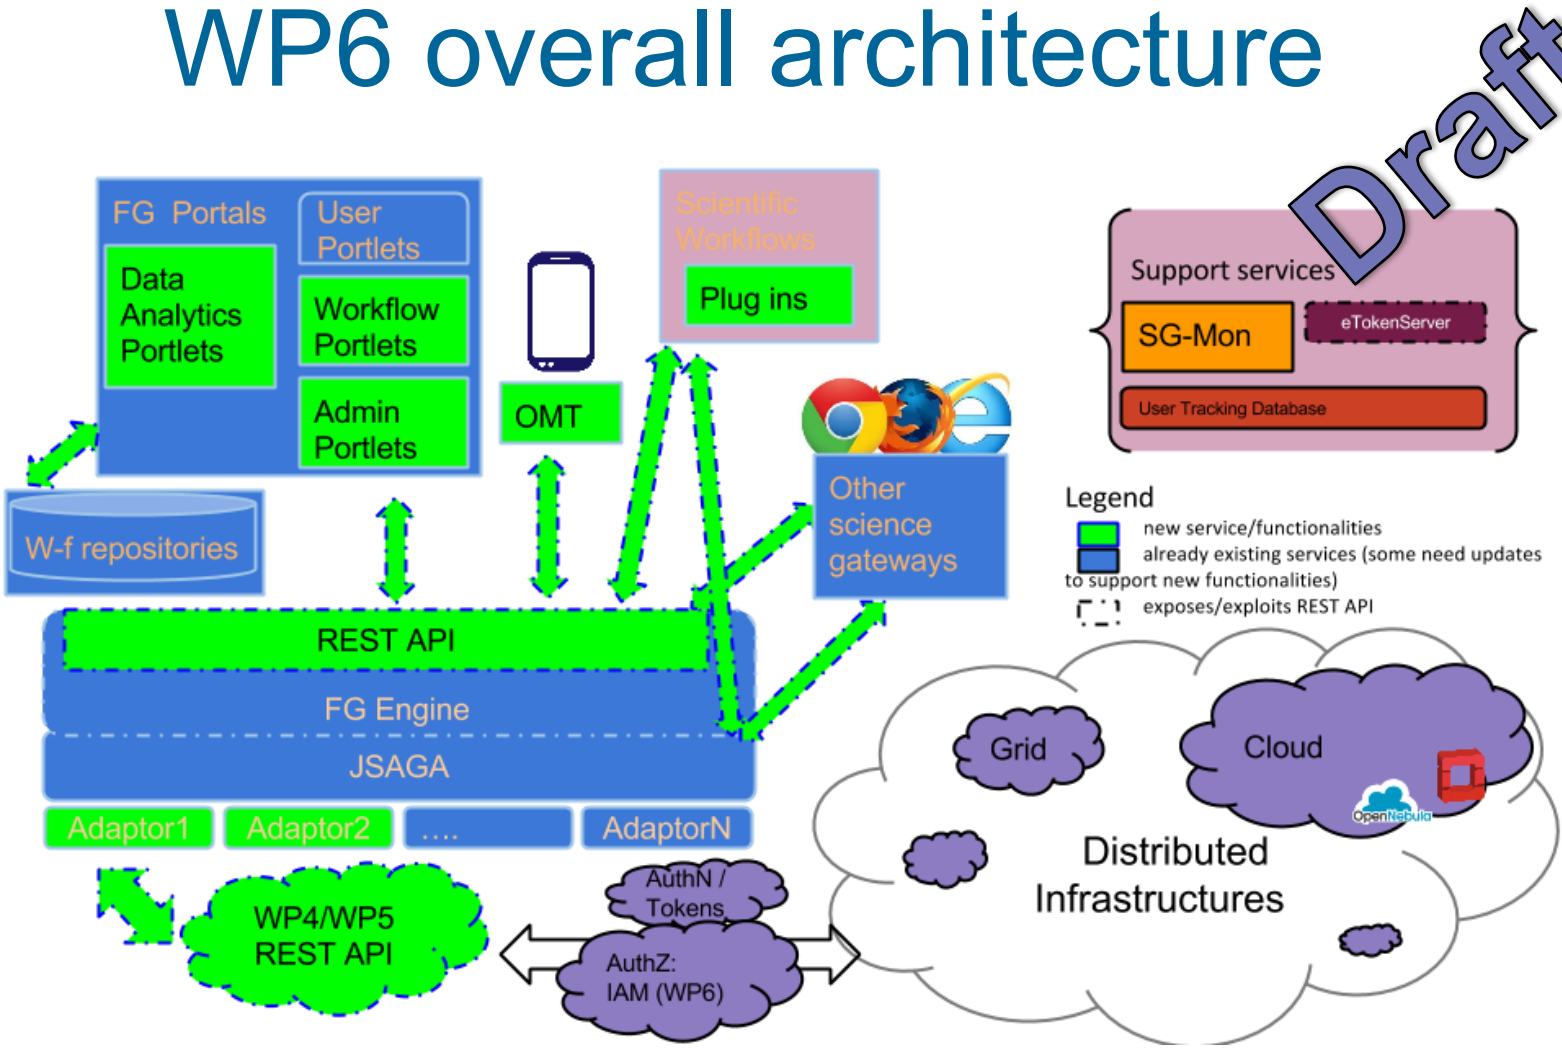

A preliminary demo based on an HPC use case (provided by CMCC) is planned at the EGI CF in November

Relazione con CCR (vedi anche presentazioni a riunione CCR, marzo e maggio 2015)

- INDIGO è un progetto che nasce e cresce su esperienze direttamente connesse ad attività discusse e gestite in CCR. Tra queste, in particolare, il gruppo Cloud INFN, ma anche progetti legati ad esempio ad autenticazione, security o gestione della rete → vedi anche presentazione di Stefano Stalio su proposta di architettura per INFN-CC (Corporate Cloud)
- Ringraziamo il management INFN, il gruppo C3S e la CCR per il loro fondamentale supporto a INDIGO
- E rinnoviamo l'invito a contribuire a INDIGO a chiunque sia interessato al progetto, anche solo leggendo i deliverable e fornendo feedback, casi d'uso, suggerimenti vari, etc.

GRAZIE!

Davide.Salomoni@cnaf.infn.it