

INDIGO-DataCloud

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

ALICE Tier-1/Tier-2 Workshop
Torino, 23/2/2015

Distributed Computing: The Next Steps

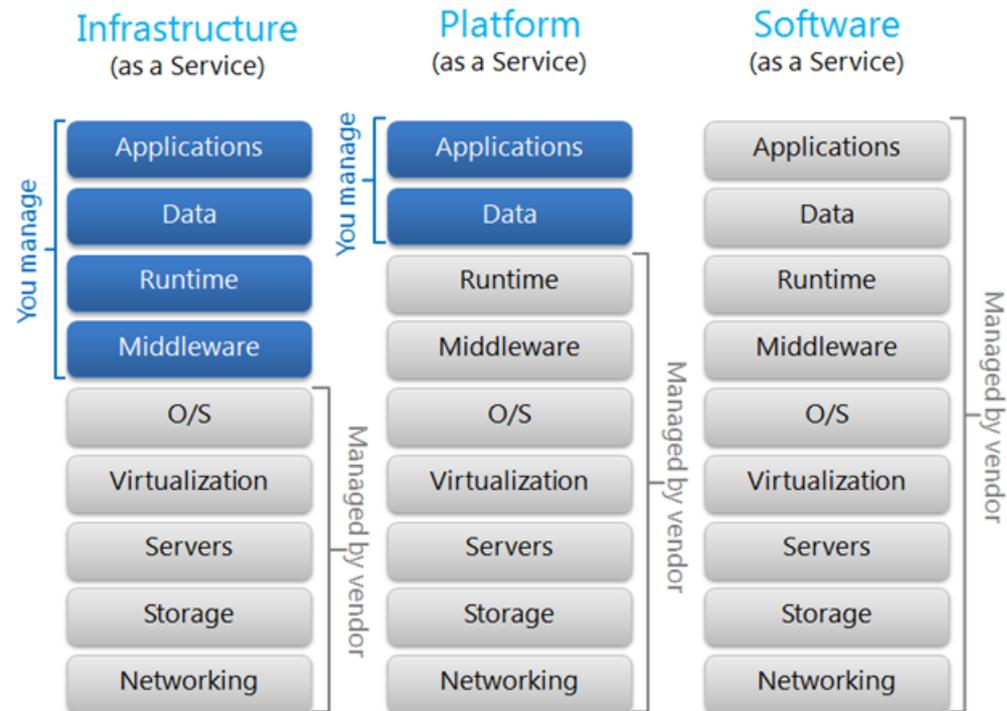
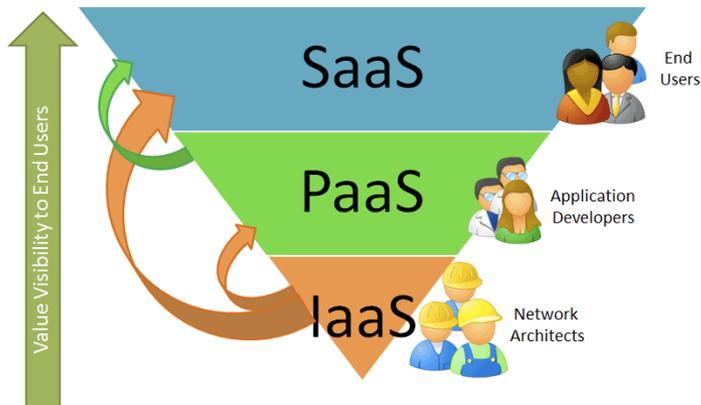
1. Ease of access, for small and big collaborations alike.
2. Software and economic sustainability.
3. Robustness (no single points of failure).
4. Modular, scalable architecture.
5. Open source software, vendor independence.



The 5 Cloud Postulates

1. Self-service, on-demand
2. Access through the network
3. Resource pooling
4. Elasticity (with *infinite resources*)
5. Pay as you go

In the end,
Applications Rule.



The Context of the EINFRA-1-2014 Call

- **EINFRA-1-2014**, “Managing, preserving and computing with big research data”
 - Research and Innovation Action (RIA)
- **What:** “Development and deployment of integrated, secure, permanent, on-demand service-driven, privacy-compliant and sustainable e-infrastructures incorporating advanced computing resources and software.”
- **Why:** “increase the capacity to manage, store and analyze extremely large, heterogeneous and complex datasets[1], including text mining of large corpora.”
- **How:** “[P]rovide services cutting across a wide-range of scientific communities and addressing a diversity of computational requirements, legal constraints and requirements, system and service architectures, formats, types, vocabularies and legacy practices of scientific communities that generate, analyse and use the data.”

[1]: Research data include large datasets collected, developed or generated for/by research, integration of small distributed datasets, as well as data not originally collected for research, which may include environmental, social and humanities data.

EINFRA-1-2014: Items 4-5 (specifically addressed by INDIGO)

- **Item 4:**

- Large scale virtualization of data/compute centre resources to achieve on-demand compute capacities, improve flexibility for data analysis and avoid unnecessary costly large data transfers.

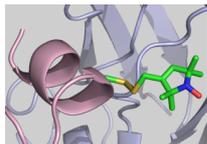
- **Item 5:**

- Development and adoption of a standards-based computing platform (with open software stack) that can be deployed on different hardware and e-infrastructures (such as clouds providing infrastructure-as-a-service (IaaS), HPC, grid infrastructures...) to abstract application development and execution from available (possibly remote) computing systems. This platform should be capable of federating multiple commercial and/or public cloud resources or services and deliver Platform-as-a-Service (PaaS) adapted to the scientific community with a short learning curve. Adequate coordination and interoperability with existing e-infrastructures (including GÉANT, EGI, PRACE and others) is recommended.

Looking for Solutions...

- **Biological and medical science**

- Biological, molecular and medical imaging, life science research applied to medicine, agriculture, bio-industries and society, structural biology.



- **Social sciences, arts and humanities**

- Georeferencing (e.g. of current or historical maps), cultural heritage, smart sensors.



- **Environmental and earth science**

- Biodiversity and ecosystem research, interactions between geosphere, biosphere and hydrosphere, earth system modeling.



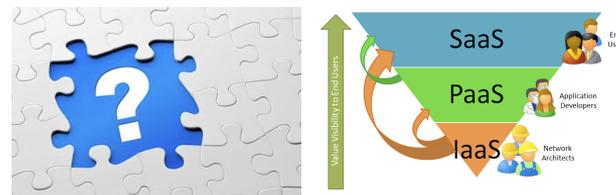
- **Physical sciences**

- Astrophysics, theoretical and experimental research in physics.



What is Missing? (1)

Or, How to...



- ... **orchestrate and federate** Cloud, Grid and HPC [public or private] resources?
- ... **avoid software and vendor lock-in?**
- ... **overcome performance issues** limiting massive adoption of virtualized Cloud resources in large data centers?
- ... **exploit specialized hardware**, such as GPUs or low-latency interconnections?
- ... **manage dynamic and complex workflows** for scientific data analysis?
- ... **combine data** from multiple sources and stored in multiple locations through incompatible technologies?

What is Missing? (2)

Or, How to...



- ... support **federated identities** and provide **privacy** and **distributed authorization** in open Cloud platforms?
- ... provide APIs to exploit the above and **write applications, customizable portals and mobile views**?
- ... **overcome barriers that limit the adoption of PaaS solutions**, such as the use of custom, non-interoperable interfaces, and the limited availability of APIs for technology-independent storage access?
- ... **move beyond static allocation and partitioning** of both storage and computing resources in data centers?
- ... **distribute and deploy applications** in a flexible way?
- ... **exploit** distributed computing and storage resources through **transparent network interconnections**.

Toward a Cloud-based Scientific Data & Computing Platform

- The **INDIGO-DataCloud** project has recently been approved under the *Horizon 2020* EU Framework Programme for Research and Innovation in the EINFRA-1 call. The project is coordinated by INFN.
- INDIGO will build an **open source data & computing platform targeted at scientific communities**, deployable on multiple hardware, provisioned over private or public e-infrastructures.



- The INDIGO Consortium, made of **26 academic and industrial partners from 11 European countries**, will be funded with **11.1 million Euros** in 30 months (expected starting date April-May 2015).

INDIGO: Consortium Partners

Participant no.	Participant organisation name	Participant short name	Country
1 (Coordinator)	Istituto Nazionale di Fisica Nucleare	INFN	Italy
2	Agencia Estatal Consejo Superior De Investigaciones Cientificas	CSIC	Spain
3	Stiftung Deutsches Elektronen-Synchrotron DESY	DESY	Germany
4	Universitat Politecnica De Valencia	UPV	Spain
5	ATOS Spain SA	ATOS	Spain
6	Consorzio Interuniversitario Risonanze Magnetiche di Metallo Proteine	CIRMMMP	Italy
7	Istituto Nazionale Di Astrofisica	INAF	Italy
8	Laboratorio de Instrumentacao e Fisica Experimental de Particulas	LIP	Portugal
9	Karlsruher Institut fuer Technologie	KIT	Germany
10	Universiteit Utrecht	UU	The Netherlands
11	European Organization for Nuclear Research	CERN	Switzerland
12	T-Systems International Gmbh	T-Systems	Germany
13	Centre National de la Recherche Scientifique	CNRS	France
14	Centro Euro-Mediterraneo sui Cambiamenti Climatici	CMCC	Italy
15	Istituto Centrale per il Catalogo Unico delle biblioteche italiane e per le informazioni bibliografiche	ICCU	Italy
16	SANTER REPLY SpA	REPLY	Italy
17	Akademia Gorniczno-Hutnicza Im. Stanislawia Staszica W Krakowie	AGH / AGH-UST	Poland
18	Instytut Chemii Bioorganicznej Polskiej Akademii Nauk	IBCH PAS	Poland
19	Stichting European Grid Initiative	EGL.eu	The Netherlands
20	INDRA Sistemas S.A.	INDRA	Spain
21	Consiglio Nazionale delle Ricerche	CNR	Italy
22	Science and Technology Facilities Council	STFC	United Kingdom
23	CESNET, Zajmove Sdruzeni Pravnickyh Osob	CESNET	Czech Republic
24	Istituto Nazionale di Geofisica e Vulcanologia	INGV	Italy
25	Ruder Bošković Institute	RBI	Croatia
26	Commissariat a l'Energie Atomique et aux energies alternatives	CEA	France

High-level Description

- INDIGO aims at **developing a data/computing platform** targeted at scientific communities, deployable on multiple hardware, and provisioned over hybrid (private or public) e-infrastructures.
- In INDIGO, **key European developers, resource providers, e-infrastructures and scientific communities have joined** to ensure the successful exploitation and sustainability of the project outcome.
- A key strength and benefit of the INDIGO consortium, for both the scientific and industrial sectors, is the potential to **create a new sustainable Cloud competence in Europe for PaaS**, similar to what OpenNebula or OpenStack have done for IaaS.

The Guiding Principles

- **Develop an open framework**, based on open source software, and without initial restrictions on the e-Infrastructure to be accessed (public or commercial, GRID/Cloud/HPC) or its underlying software (OpenStack or OpenNebula). The consortium will take into consideration the medium and long-term exploitation and sustainability in this framework.
- **Exploit existing solutions**, learning, re-using and extending them according to user requirements, and having in mind the expected evolution of technology. INDIGO will also keep a direct contact with e-Infrastructure providers, to assure a successful deployment of the INDIGO Platform and to guarantee an adequate level of support.
- Ensure that the framework offered to final users as well as to developers will have a **low learning curve**, while providing a clear added value. In particular, existing software suites popular in different research communities, and offering already a rich environment for large data processing, like ROOT, OCTAVE/MATLAB, MATHEMATICA or R-STUDIO, will be supported and offered in a transparent way to the final user, while running on powerful remote e-Infrastructure resources.

Relationship among WPs

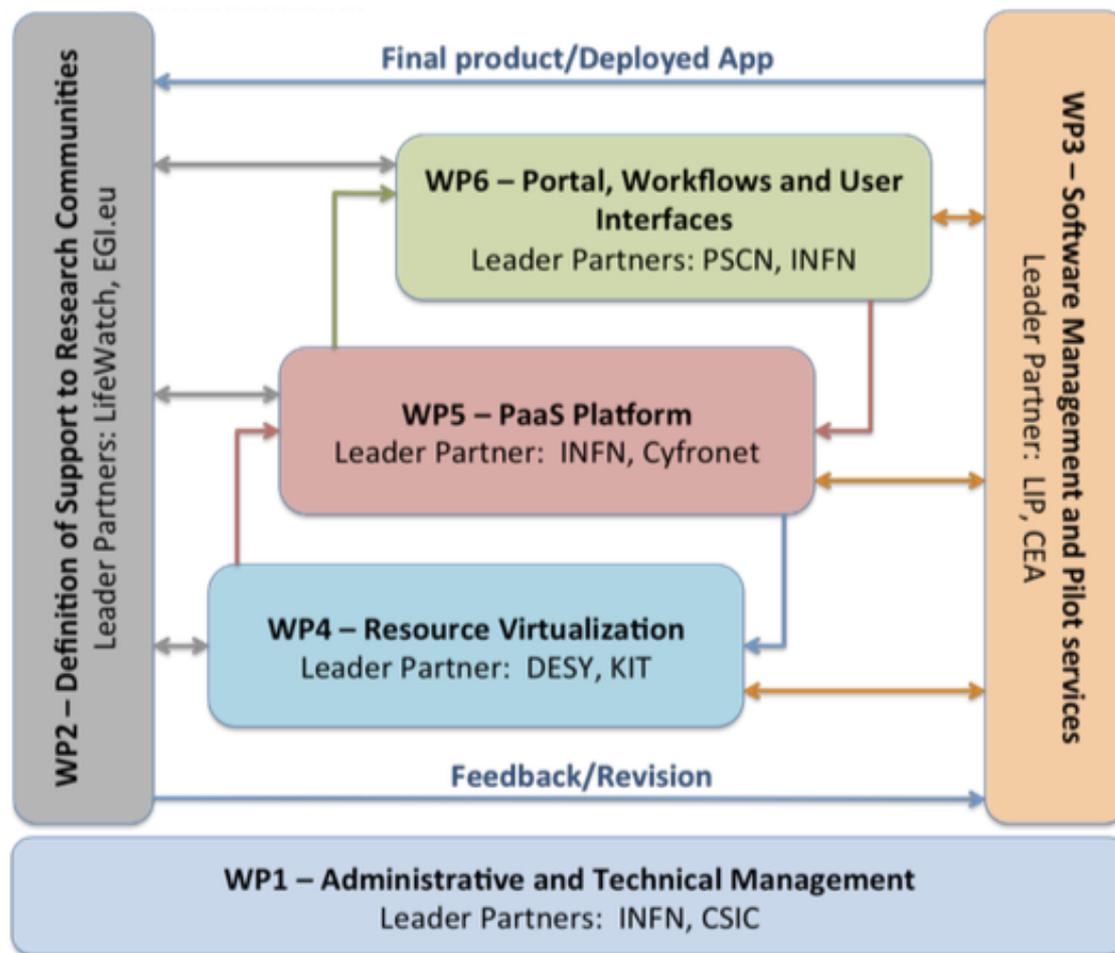


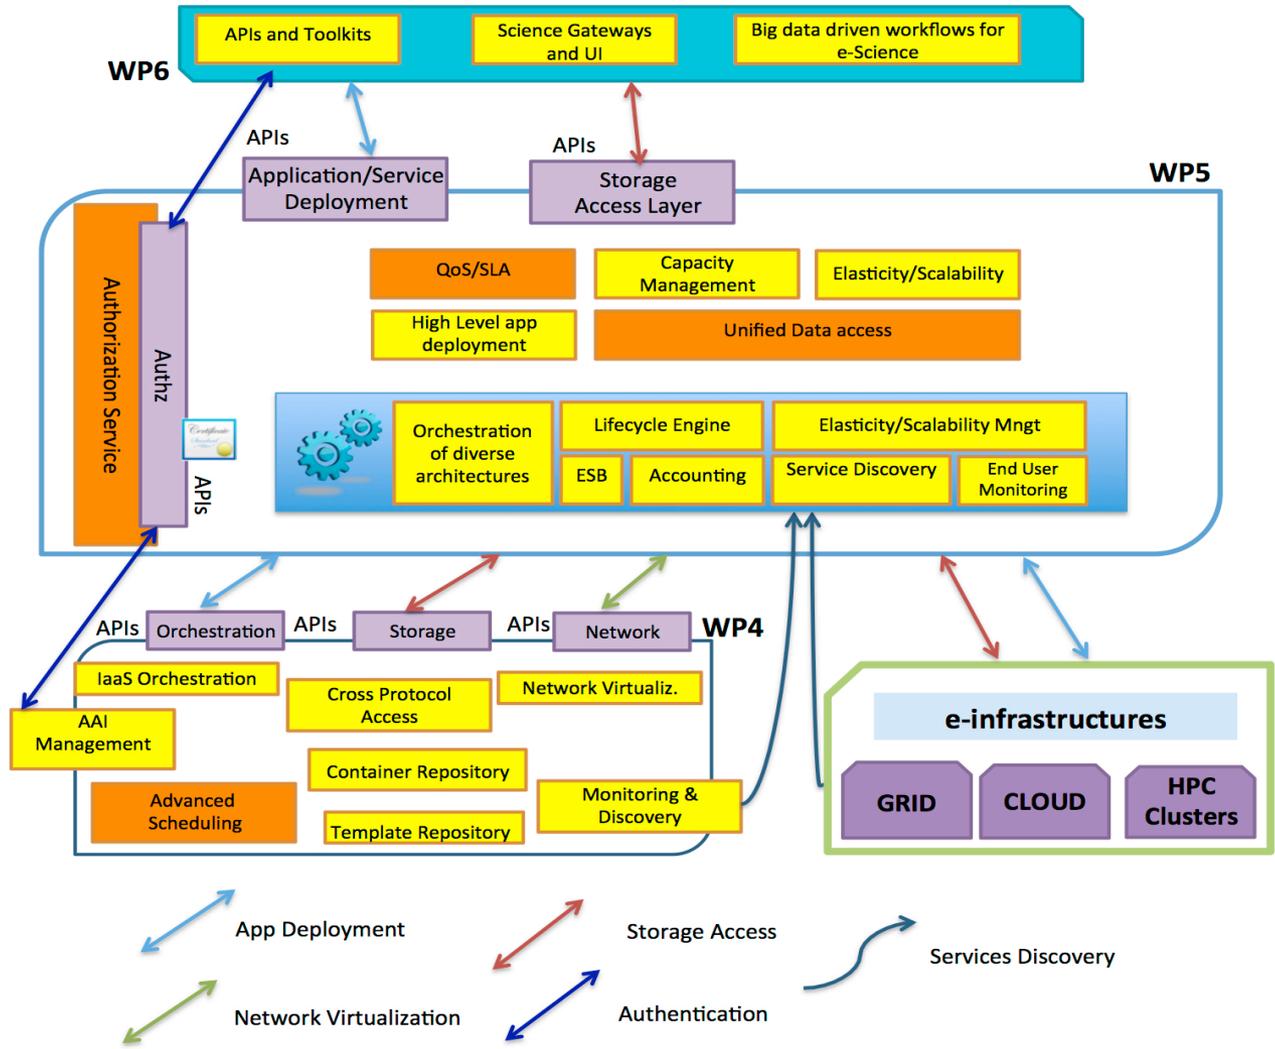
Figure 4: Diagram showing the interrelation among the Work Packages

INDIGO: Global Architecture

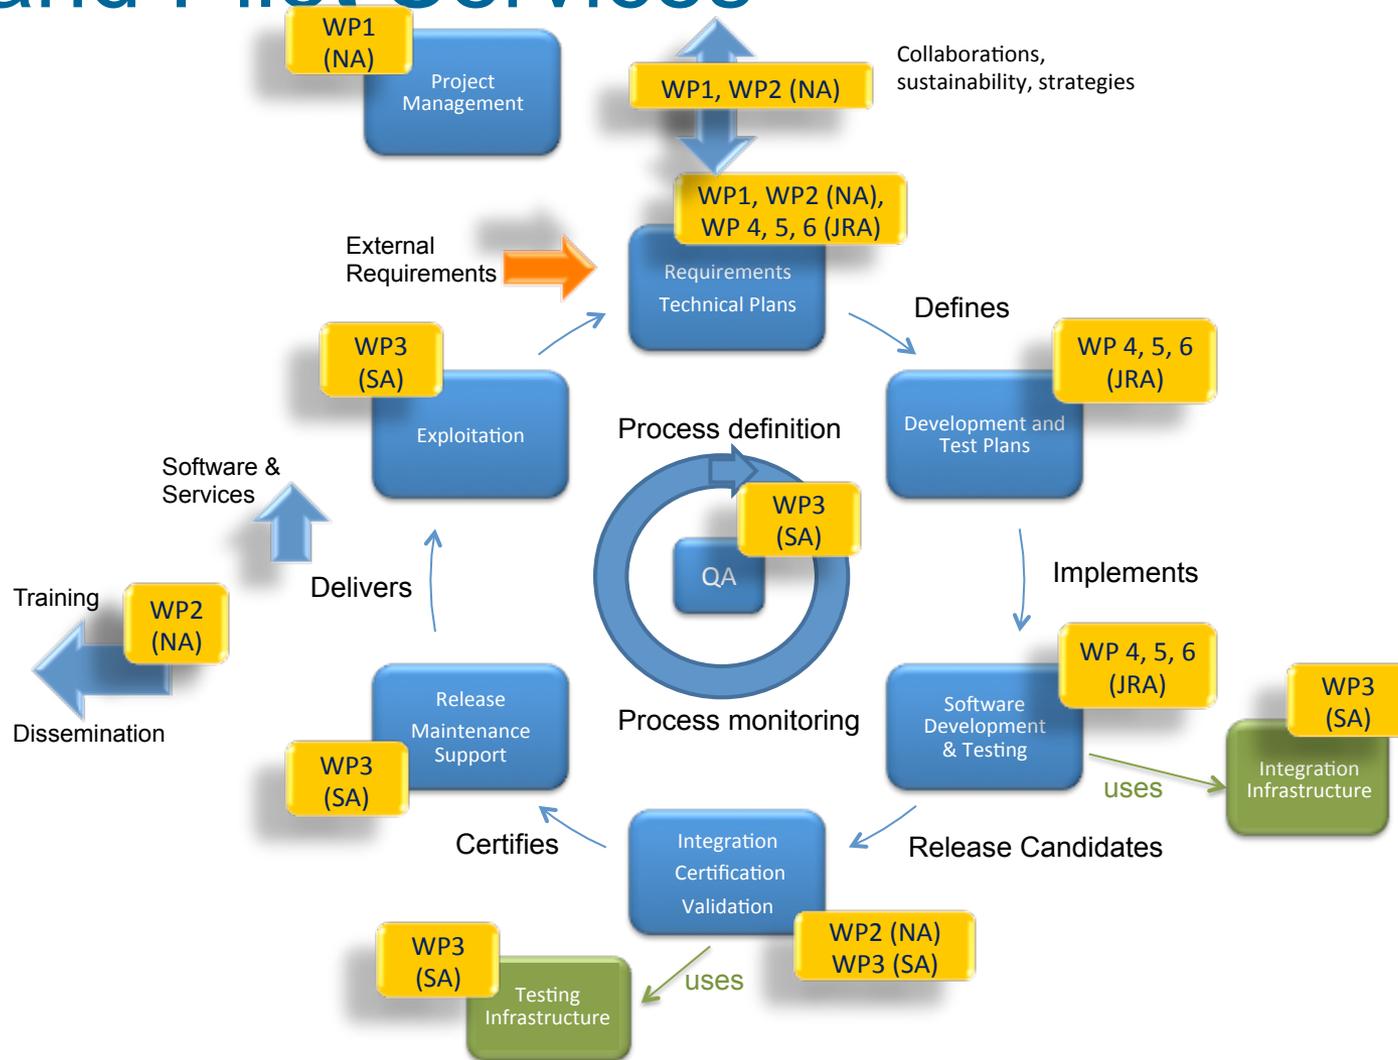
Color codes:

Yellow: implementation based on already available solutions, to be improved or changed.

Orange: Completely new services to be developed.



WP3: Software Management and Pilot Services



WP3: Software Management and Pilot Services

- T3.1 will deal with the **software quality assurance**, compiling and enforcing the necessary quality criteria, indicators, and tests necessary to ensure high quality software components ready for production.
- T3.2 will make the certified software components available as a set of coherent **high quality releases**, supported by an efficient maintenance process.
- T3.3 will provide the **pilot infrastructures** and services for integration and testing supporting the tasks T3.1 and T3.2 activities.
- T3.4 will interface with major production e-infrastructures, collect their feedback, requirements, and will enable a path **towards production exploitation**.

WP4, resource virtualization

Computing	Storage	Network
Providing support for container	Defining interfaces and implementing QoS support for storage systems	Evaluation of available SDN features and operability.
Improving the on-demand compute capabilities through improved orchestration and scheduling	Providing access to the same storage through various standard access protocols.	Using SDN to configure local networks and meet PaaS needs.
		Manage local virtual Networks.
Common Subtasks		
Authentication, Authorization and Identity Management (AAI)		
Service Discovery and Monitoring		

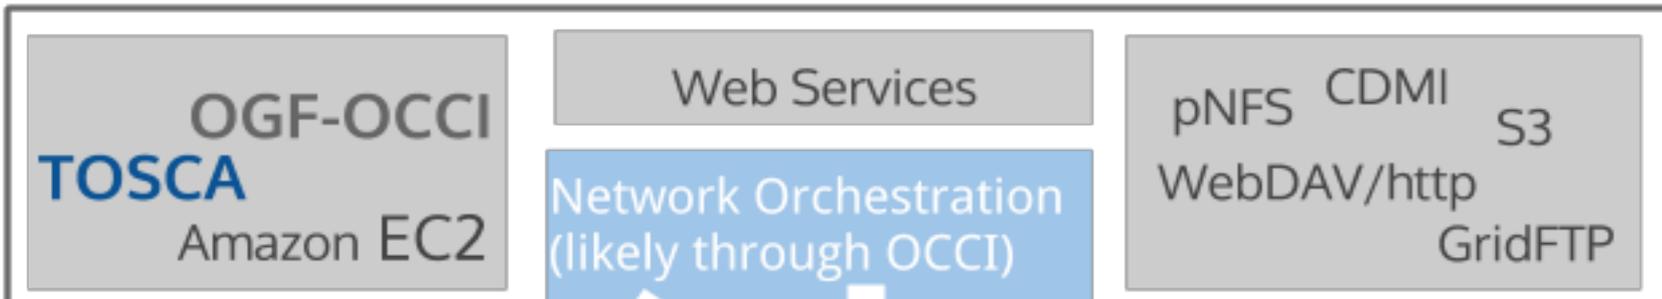
Table 6: Breakdown of WP4 into specific and common tasks

WP4: Resource Virtualization

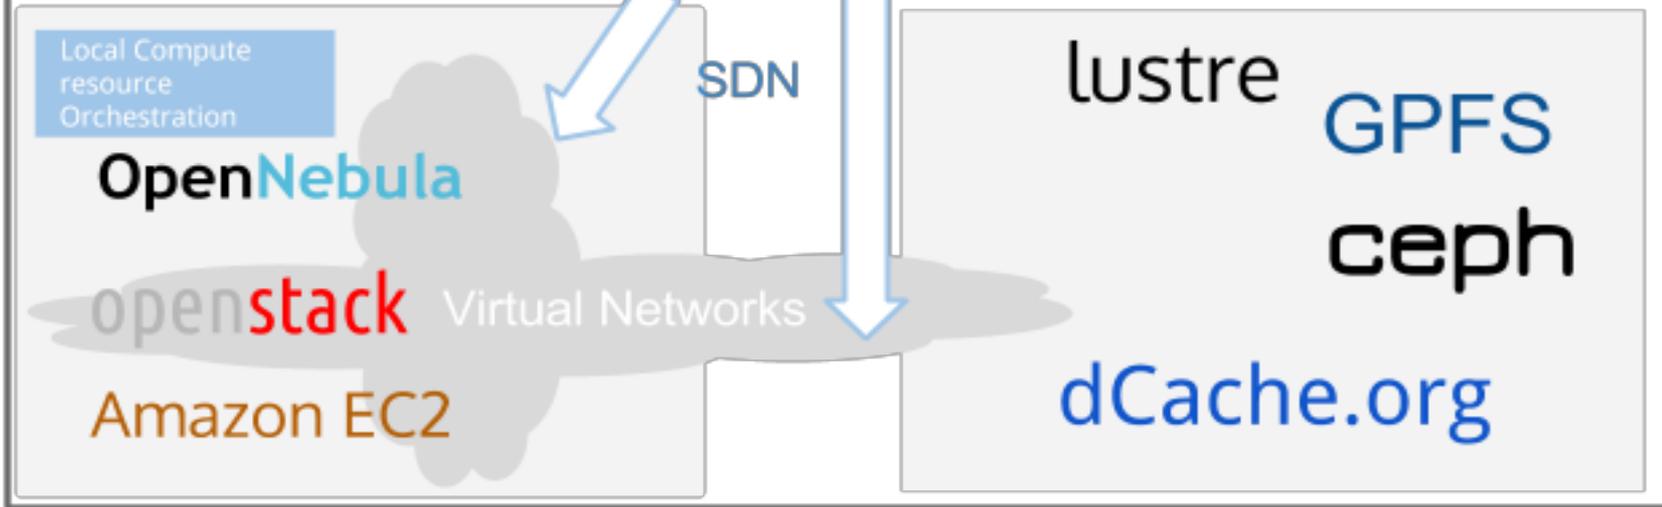
Higher Level Services



Abstraction Layer



Site Local Product Layer



WP4: Resource Virtualization

Common subtask:

- **AAI Management** for the virtualized computing cloud infrastructure
 - Integration of federated AAI technologies into OpenStack, OpenNebula, CEPH, dCache and other supported INDIGO products, allowing users to **access infrastructure** resources using their **home** or guest **IdP account**.
- **Service Discovery** and Monitoring
 - Extension of existing **local-site monitoring services** for all INDIGO products to provide higher-level services in WP5 and WP6 monitoring and accounting information through a query-API.

WP4: Resource Virtualization

Cloud Computing Virtualization

- Providing **support for containers**
 - as a portable and performing platform for the execution and deployment of applications, supporting local site **orchestration features** (e.g. HEAT or OneFlow) simplifying management of the lifecycle of both containers and VMs.
- Improving the on-demand compute capabilities of data-centers by **extending compute orchestration and scheduling**
 - To improve the existing **cloud schedulers** in **OpenStack** and **OpenNebula** to include the support for postponing low priority workloads (by killing, preempting or stopping running containers or VMs) in order to allocate higher priority requests, thus enabling **advanced scheduling policies**, optimizing the usage of the data center and improving its response to the users.

WP4: Resource Virtualization

- **Cloud Storage Virtualization**
- **QoS Support** in storage
 - will enable users to specify **service quality policies** for their data. In collaboration with RDA, we envision standardizing the associated terms and definitions, so users can expect the same quality of service regardless of the underlying implementation.
 - Evaluate/extend **available protocols** (e.g. CDMI, WebDAV, SRM) supporting the defined service levels.
- **Cross Protocol support** for storage solutions
 - Use cases often require storing files with one access protocol and subsequently accessing the **same data with a different protocol**. This requires enabling access to identical data via different protocols.

WP5, PaaS development

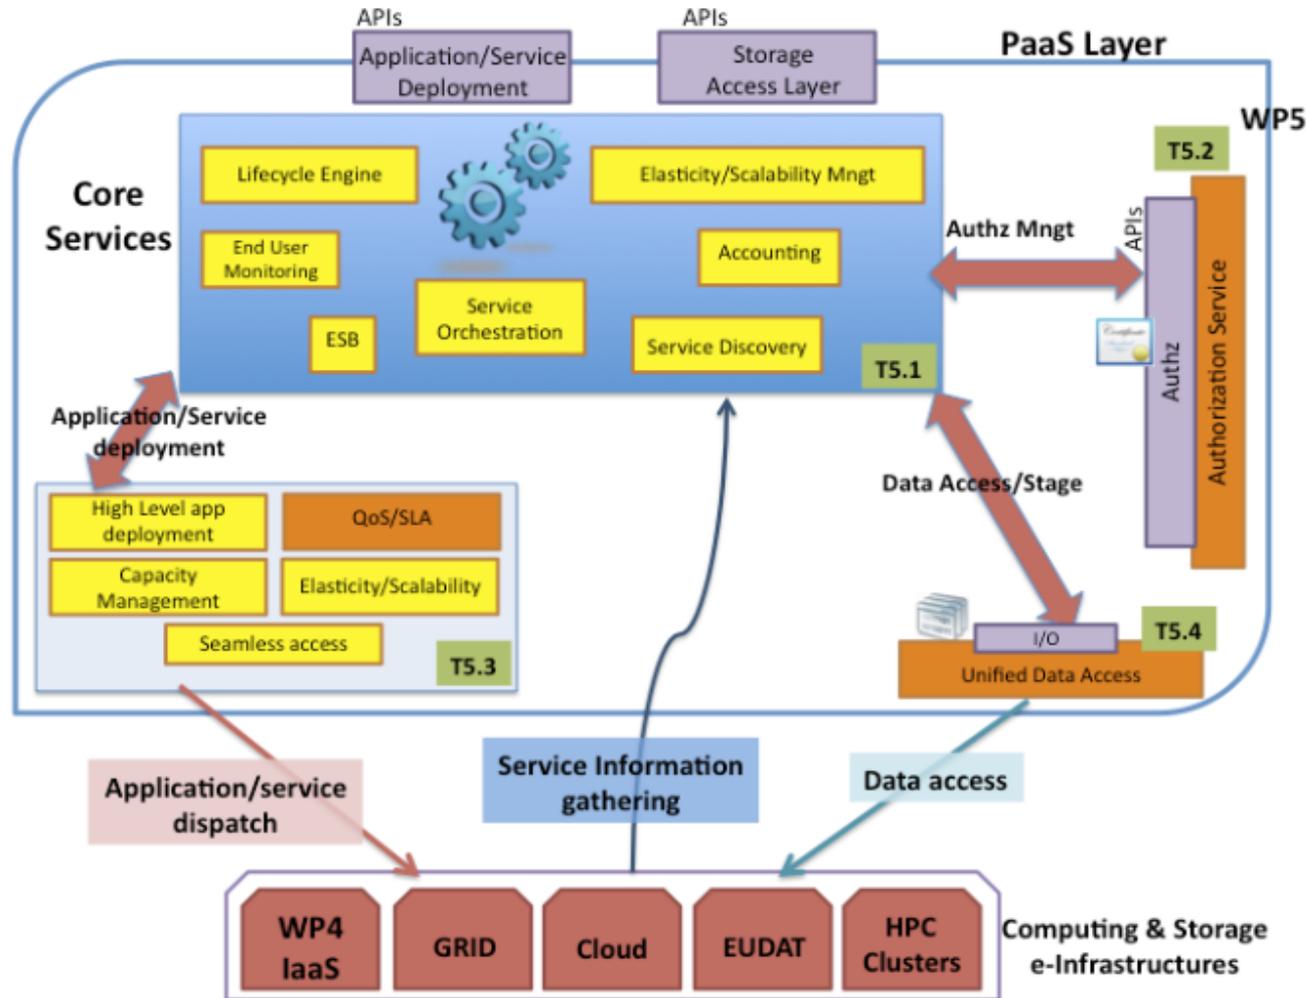


Figure 14: WP5 architecture

WP5 services (1)

- **Core PaaS**
 - An orchestration system that will receive standard descriptions of services or applications (e.g. TOSCA) and their interactions with other services or applications. It is a crucial component conveying interest of both research communities and industrial partners.
- **Authorization services**
 - Services and tools needed to enable a secure composition of services from multiple providers in support of scientific applications.

WP5 services (2)

- **Dispatcher service**
 - Deploy in a transparent way both services and applications in a distributed and heterogeneous environment made by different infrastructures (EGI Grid, EGI Fed Cloud, IaaS Cloud, Helix Nebula, PRACE, local HPC clusters, etc).
- **Unified data access layer**
 - This task will provide users and other INDIGO services a unified and transparent access to distributed storage. As a result, a set of APIs and services will be delivered to face problems of data orchestration, storage access optimization, remote data access. The task will focus on effective and reliable discovery of storage systems capabilities, understanding their data management policies, discovering Retention Time and Access Latency (e.g. out of meta-data or through SRM/CDMI interfaces).

WP6, Science Gateways, Workflows and Toolkits

- **Libraries and toolkits**
 - Build a set of libraries and toolkits on the REST APIs developed by WP5. The aim of these libraries and toolkits is to simplify the development process and speed up the creation of science gateways, desktop and mobile applications. This will be a completely new development with respect to past development of Science Gateway frameworks (such as the Catania Science Gateway).
- **Science gateways and mobile apps**
 - Develop a general-purpose, multi-domain Science Gateway framework, based on the specific requirements of the WP2 communities, exposing and exploiting the PaaS features developed in WP5. This framework, together with mobile applications using the toolkits also developed in WP6, will be easily customized and deployed for the purposes of said communities through the pilot services set up within WP3.
- **Support for big data driven workflows for e-Science**
 - Based on the use cases gathered in WP2, this task will provide dynamic support for scientific workflows management according to a “Workflow as a Service” (WaaS) model. The proper workflow engines will be selected taking into account user needs and requirements. With regard to the different application scenarios, this task will provide workflow services able to seamlessly orchestrate workflows in Cloud, Grid and HPC environments. Input and output data could be stored on Grid, Cloud, local, external storage resources and will be accessed through standards interfaces (e.g. SAGA, OCCI, CDMI).

Timing and Collaborations

- We aim to have the INDIGO kick-off meeting tentatively by April 2015.
 - We would anyhow like to coordinate with other projects approved in the EINFRA-1 call. A concertation of goals and timing might be useful for all projects.
- A few areas where we believe interactions with Alice might especially be fruitful:
 - Exploitation of containers
 - QoS / cross-protocol support in storage systems
 - Testing on heterogeneous Cloud infrastructures (e.g. OpenNebula, OpenStack)
 - Use of dynamically instanced analysis clusters (a' la elastiQ / PoD)
 - Exploration of other areas for collaboration absolutely welcome!

THANKS!

Davide.Salomoni@cnaif.infn.it