Istituto Nazionale di Fisica Nucleare

WNoDeS: un servizio per la gestione di infrastrutture condivise Cloud e Grid

Peter Solagna - INFN

Davide Salomoni - INFN

Alessandro Italiano - INFN

The INFN WNoDeS Project

✓INFN (Italian National Institute of Nuclear Physics) conducts theoretical and experimental research in the fields of subnuclear, nuclear, and astroparticle physics.

✓For the past 10 years INFN has been working in several European projects with the aim of establishing a European Grid Infrastructure to serve scientific needs. Recently, Cloud Computing has emerged as a complementary pattern to Grid Computing.

✓The Worker Nodes on Demand Service (WNoDeS) is an INFN-developed architecture that makes it possible to dynamically allocate virtual resources out of a common resource pool.

✓WNoDeS is currently running in production mode at the INFN National Computing Center (CNAF) supporting, among others, LHC experiments. There are ongoing work to deploy it in the LNL computing farm.

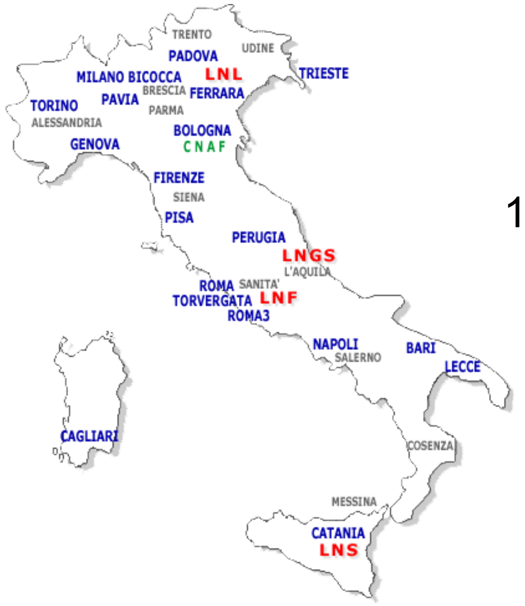

INFN:
19 sites, 4 National Laboratories,
1 National Computing Center

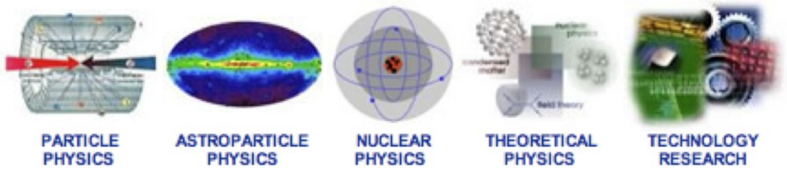

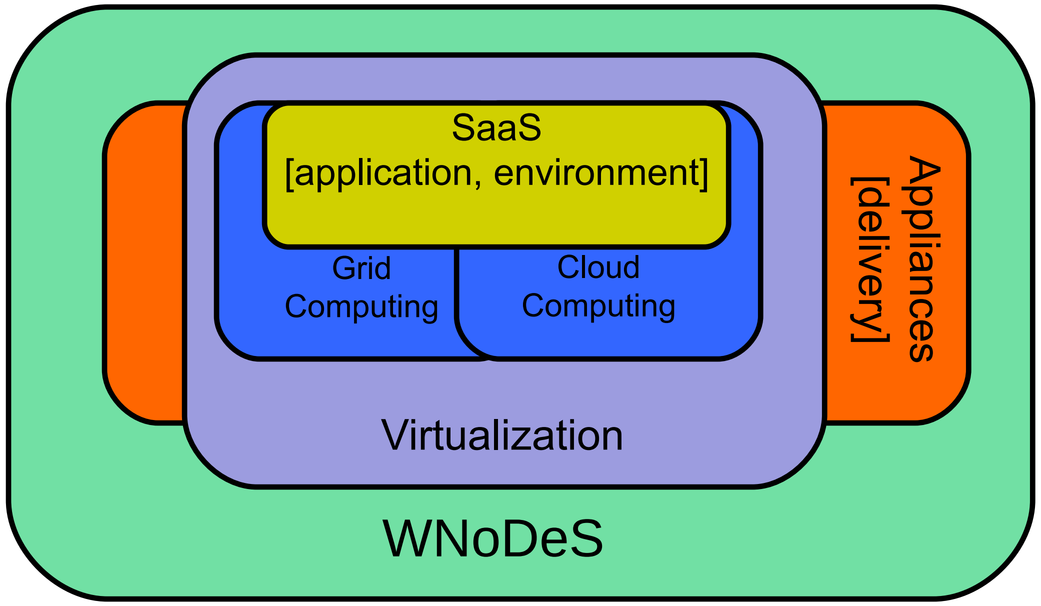

SaaS
[application, environment]

Grid Computing

Cloud Computing

Appliances
[delivery]

Virtualization

WNoDeS

What WNoDeS is

WNoDeS (Worker Nodes on Demand Service) is a virtualization framework, its main characteristics are:

- Full integration with existing computing resource scheduling, policing, monitoring and accounting workflows.
- On-demand virtual resource provisioning and VLAN support to dynamically isolate Virtual Machines depending on service type / customer requests.
- Exploit resource usage through sharing and virtualization
- Provide multiple interfaces to access resources
- Expand the use of existing (grid, cloud) infrastructures
- Support novel user requirements
- Attract new customers, either inside or outside our community of reference

Everything as a Service, where Everything may be hardware,software, data, platform, infrastructure – you name it.

The WNoDeS Architecture

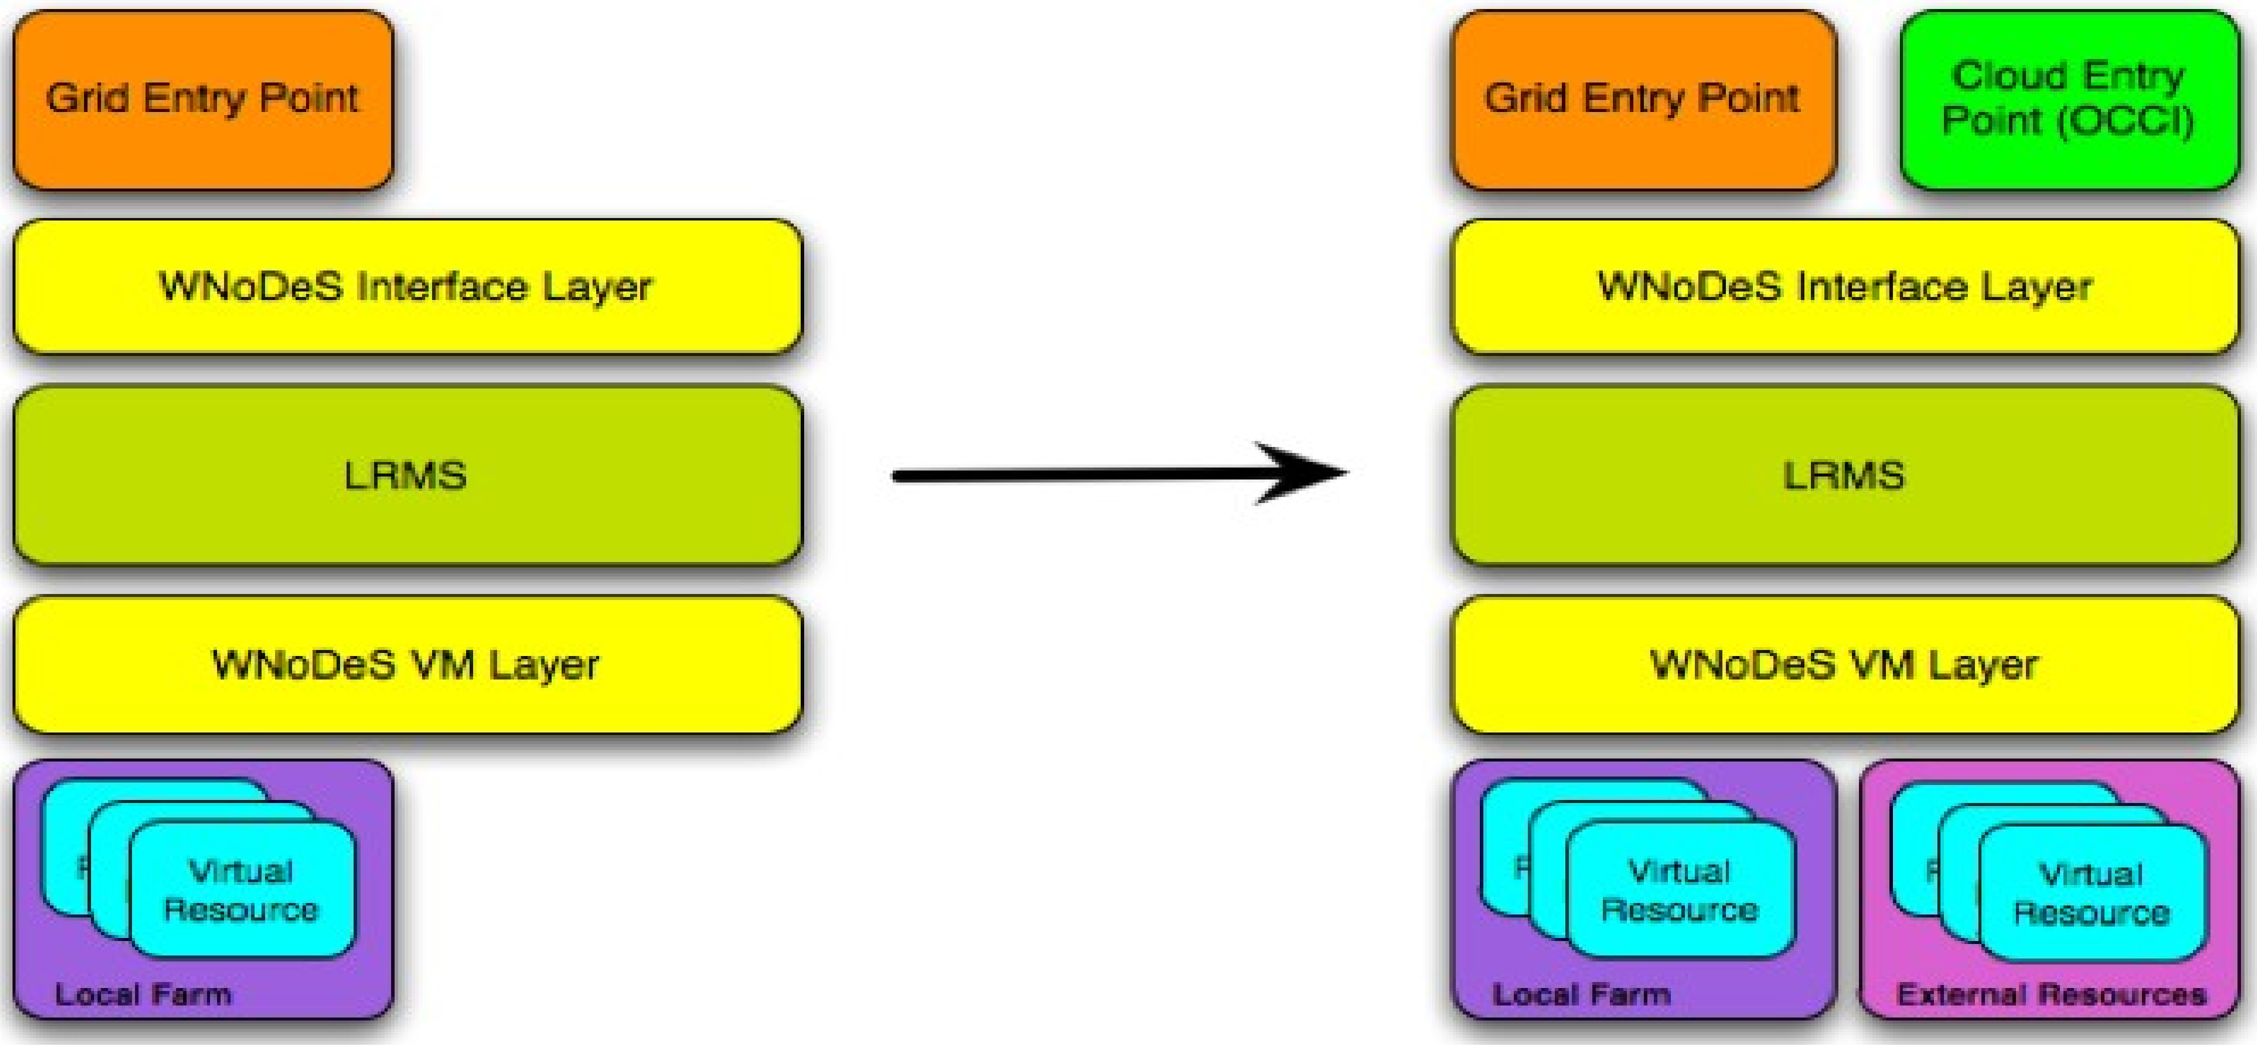

- Production deployement at the INFN Tier-1 in Bologna, Italy, with 2000 running VMs
- Virtualization layer based on Linux KVM
- Scheduling mechanism based on LRMS (Local Resource Management System) software – currently Platform LSF

- Grid interface based on the gLite middleware
- Cloud interface based on OGF OCCI

WNoDeS interfaces

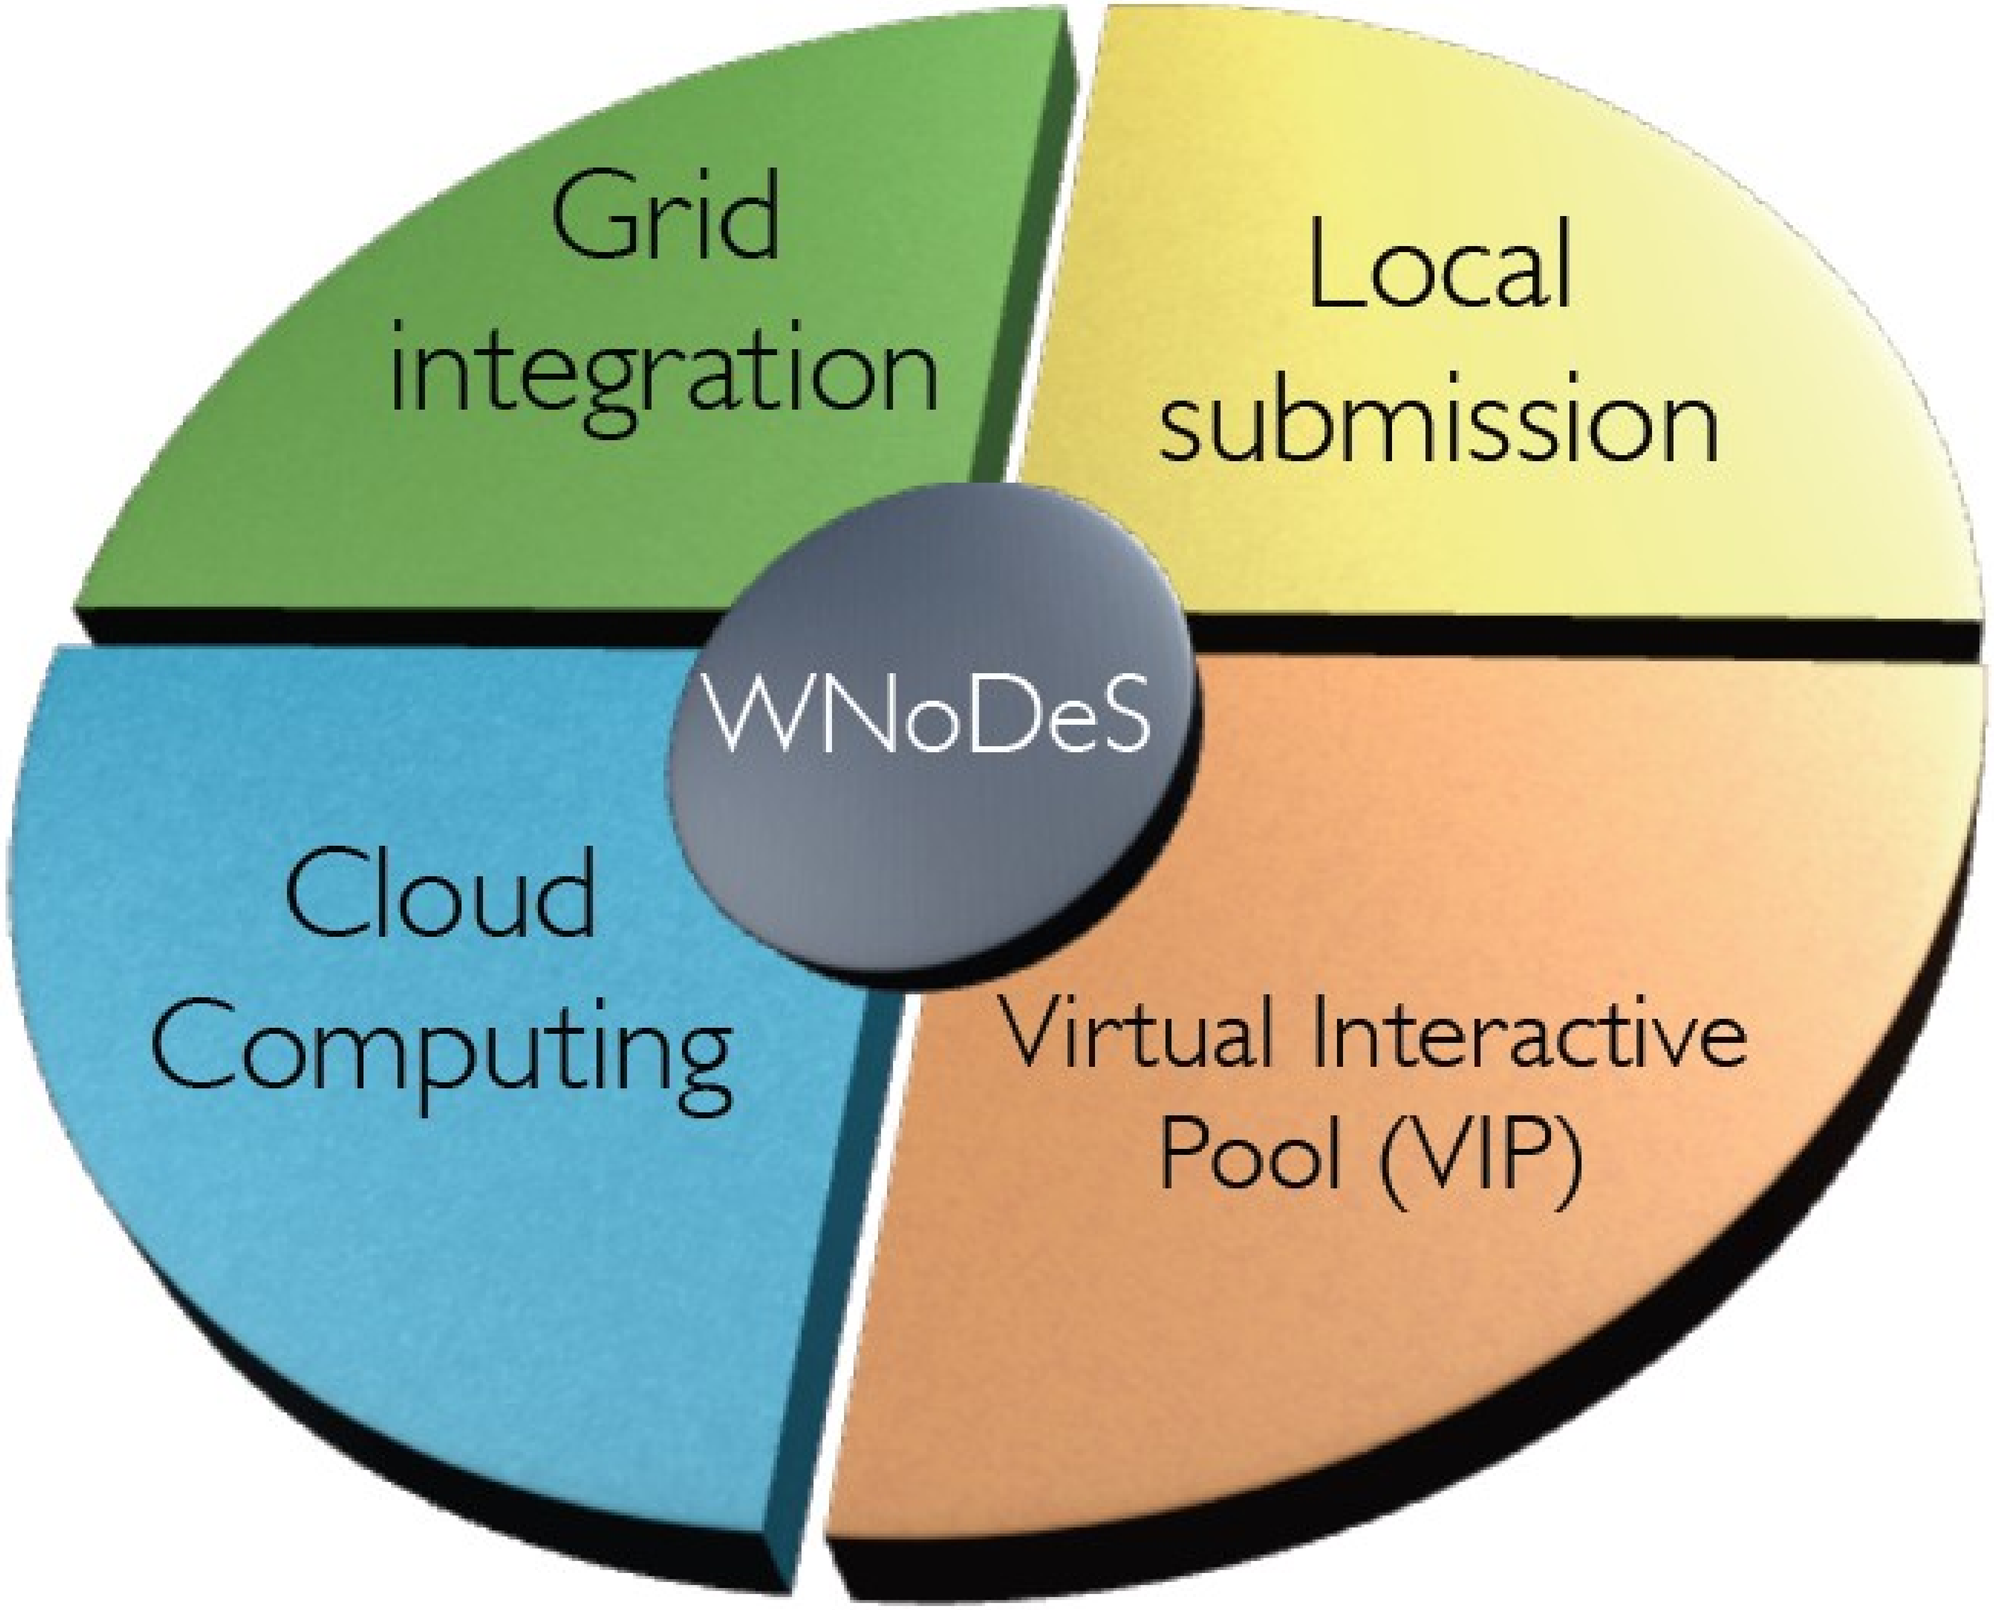

WNoDeS virtualization and VM instantiation layer

A software layer handles communication between the LRMS, the KVM VMM, the bait, and the VWN.

The bait is a special virtual machine running on each computing resource, it exposes to the LRMS the capabilities of the machine it is running on.

These layer implement mechanisms:

- to accept jobs from the LRMS on the bait;
- to trigger creation, destruction, or suspension of a KVM VM via the KVM VMM;
- and to keep job state information between the VWN and the bait.

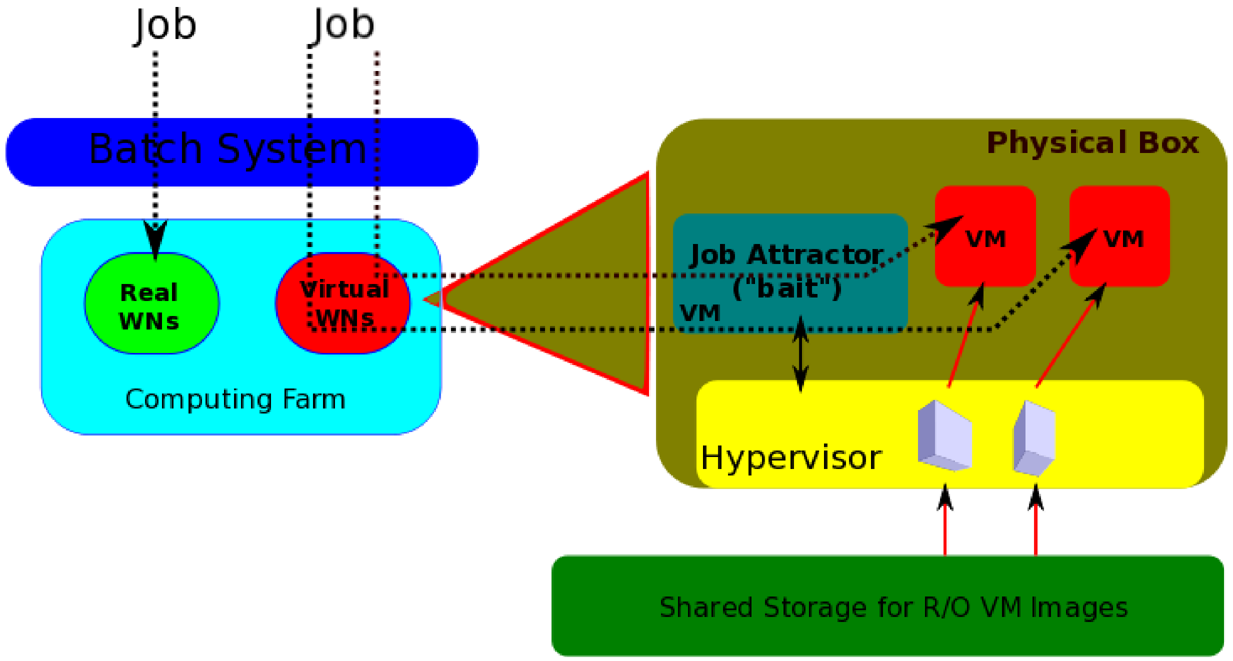

The WNoDeS Cloud interface

We chose the Open Cloud Computing Interface (OCCI) as the main API for the WNoDeS Cloud interface for a set of reasons:

- OCCI is an OGF open standard

- The use of open standards facilitates interoperability between different distributed computing infrastructures

- The OCCI standard is implemented by other Cloud frameworks like OpenNebula.

- It is simple and easy to implement.

WNoDeS integrates Grid and Cloud technologies under the same architecture

- It is not necessary for either users or resource providers to drastically change their existing workflows.

The WNoDeS OCCI Implementation

```
curl -i -H "Accept: occi/text" -X POST \
-H "Category: img_ETICS_sl53; label='eticsVm'" \
--cert x509.pem \
https://fenrir.cnaf.infn.it:8443/grid-cloud/ws/compute
```

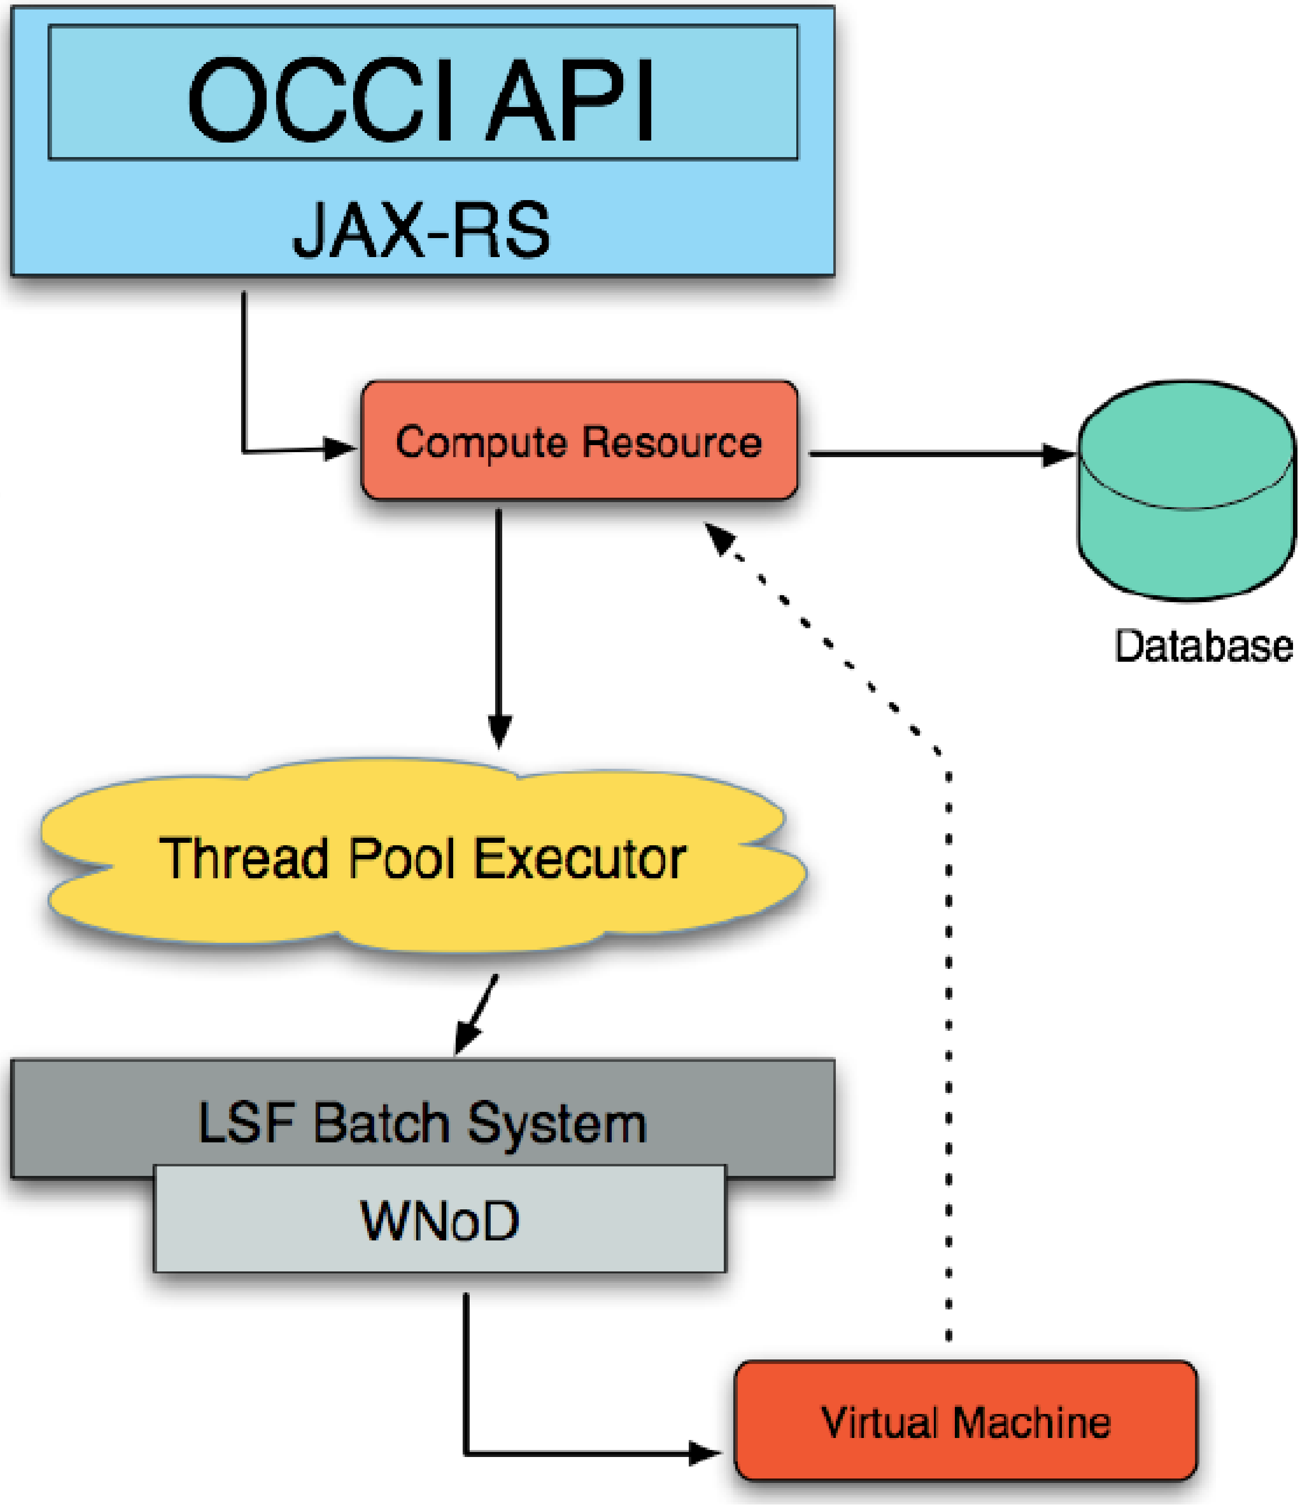

- The Cloud interface is exposed as a RESTful Web Service, implemented in Java, and running on Tomcat.

- The interface can be accessed by the REST protocol (e.g. via command-line, running a curl command to the Web Service endpoint) or by the web interface shown in the next slide.

- The Cloud layer creates computing resources when requested; this resource is then instantiated as a Virtual Machine (VM).

- As soon as the VM is up and running, the requester can retrieve its hostname and other information by querying the computing resource that was just created, through the cloud interface.

The WNoDeS Cloud Web Interface

- Cloud requests to WNoDeS may also be performed through a web application, accessible with a browser with an X.509 certificate.
- This allows to easily instantiate new virtual machines specifying their attributes, like CPU speed or number of cores.
- Existing Grid users may also specify their organization choosing one of the supported VOs.
- Another view permits to monitor the instantiated virtual machines, showing their status and attributes. Only the machines belonging to the user are shown. From this view a user can also delete his virtual machines.

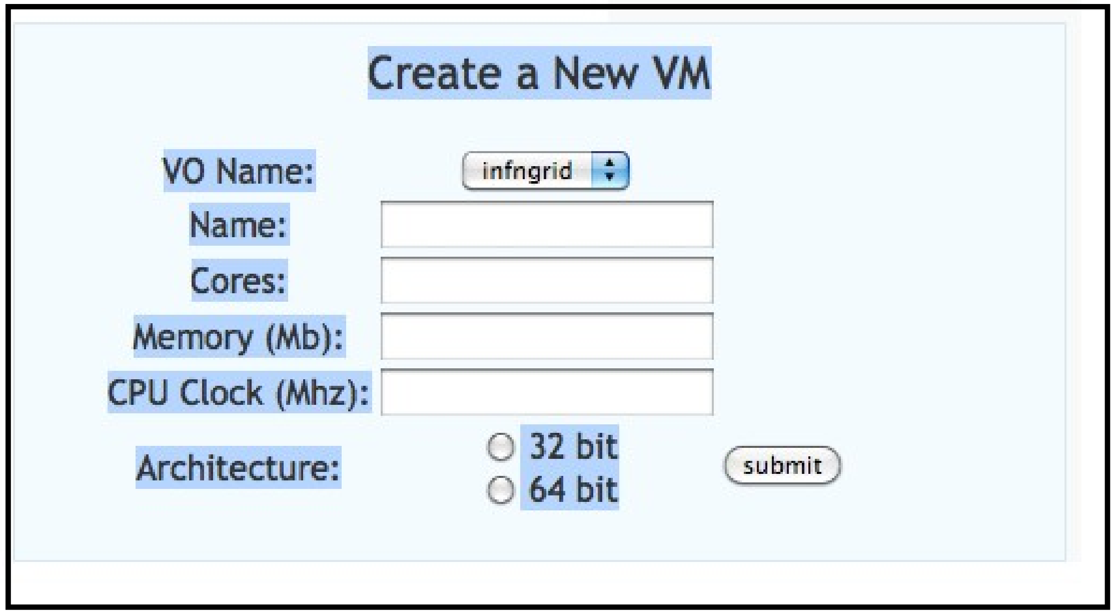

Create a New VM

VO Name: infngrid

Name:

Cores:

Memory (Mb):

CPU Clock (Mhz):

Architecture: ○ 32 bit ○ 64 bit

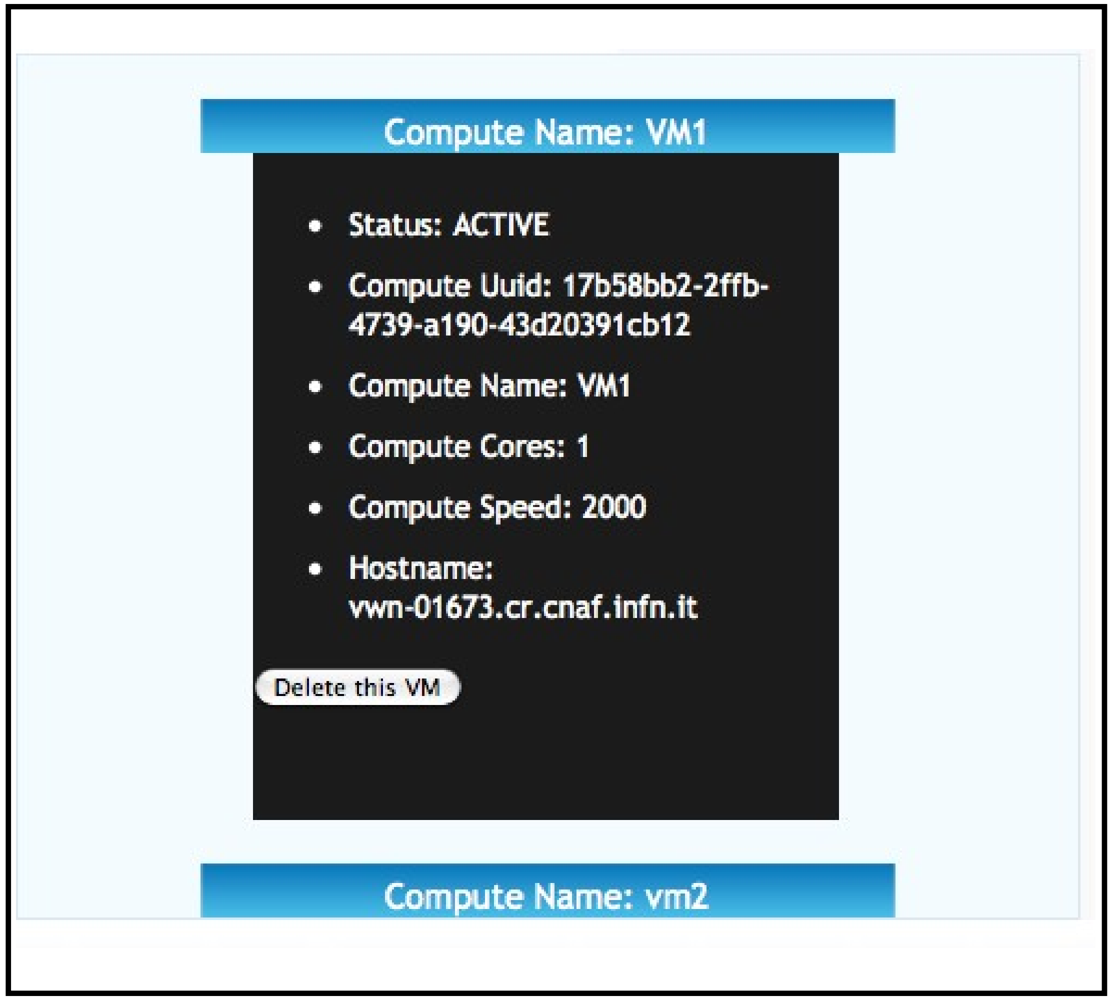

Compute Name: VM1

- Status: ACTIVE
- Compute Uuid: 17b58bb2-2ffb-4739-a190-43d20391cb12
- Compute Name: VM1
- Compute Cores: 1
- Compute Speed: 2000
- Hostname: vwn-01673.cr.cnaf.infn.it

Compute Name: vm2

Virtual Interactive Pools

User often need: interactive access, local access to the batch system, user interfaces to access distributed infrastructure.

- Classic way: statically allocate machines for interactive access, with static configuration.
- WNoDeS way: dynamically allocate resources to different tasks regardless of the needed configuration.

- Joint project of INFN and Universita` di Bologna.

- Instantiate virtual machines (VM) with the desired profile on demand, including VO specific customizations.

  - Throttle bandwidth to avoid abuse by individual users

- Open an interactive shell on that VM.

-Clients for VIP layer available on bastion front-end hosts or installable on personal workstation.

An authentication Gateway

- If X.509 certificates are already used, pass-through.
- If Kerberos or Shibboleth are used, kCA and SLCS solve the issue.
- If username/password is used, then an IdP needs to be setup, and then we fall into the previous case.

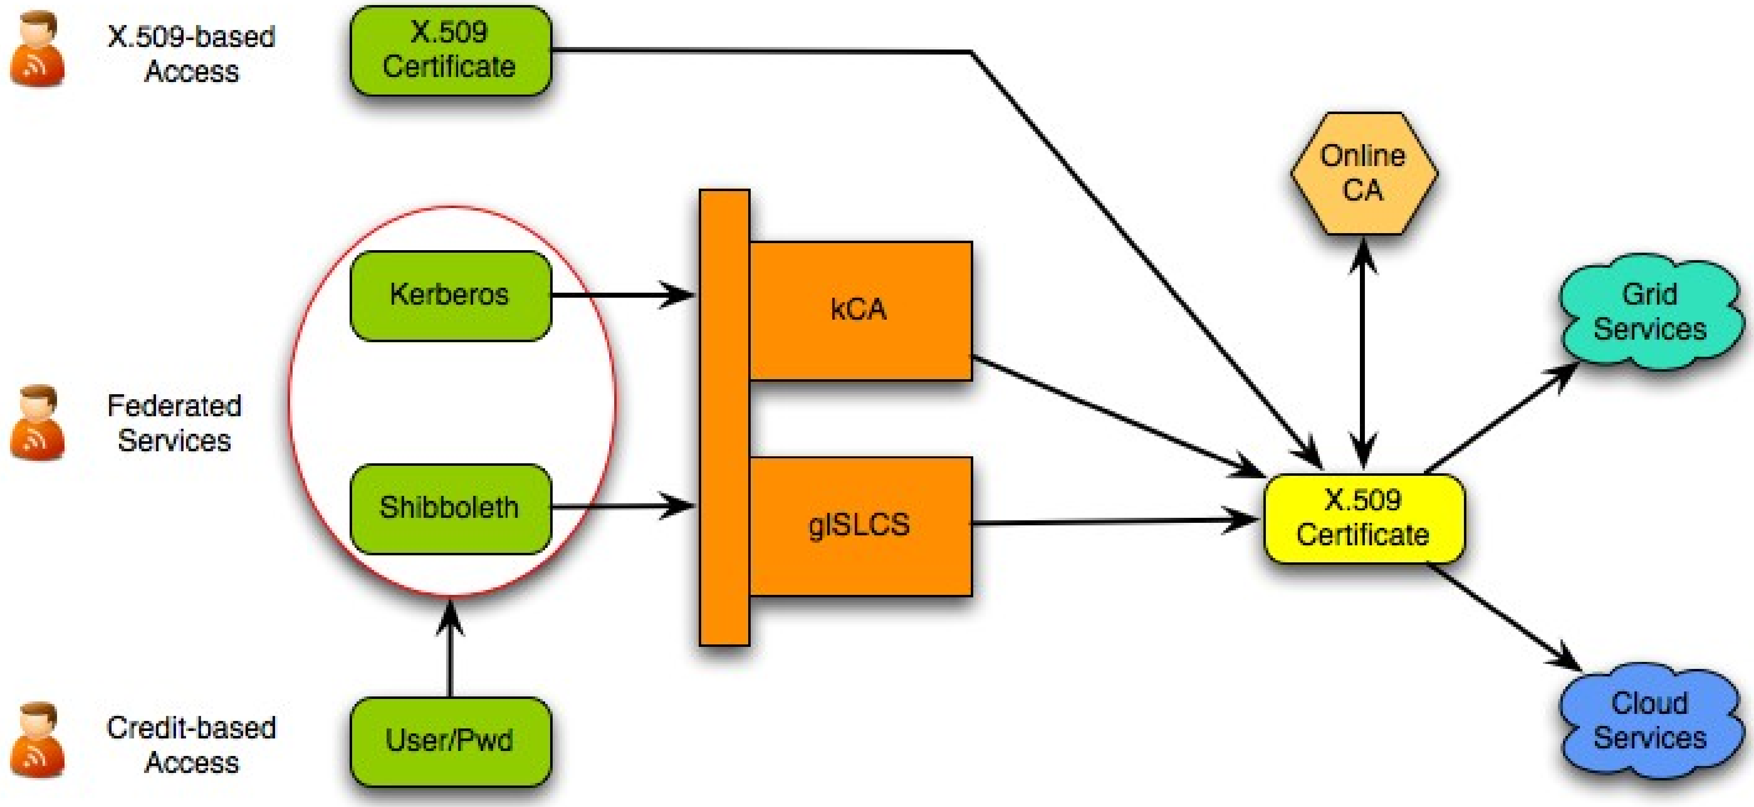

Integration: Cloud users accessing Grid resources <1>

- The Authentication Gateway provides the user an X.509 short-lived certificate – but he still needs to be a member of a VO.
  - The gateway then registers the user into a dedicated VO.
    - DN persistency is guaranteed across credential re-creation.
- Sites then need to accept the additional VO.
  - It is also possible to only accept subgroups of the VO.
  - One could have a catch-all VO, or set-up multiple VOs if the need arises (operational / business considerations apply)
- Users have gained access to Grid resources (i.e. VOMS proxies) with minimal changes to the sites.
  - Apply this to job submission portals, and/or to Cloud web portals.

Integration: Cloud users accessing Grid resources <2>

- The Authentication Gateway provides the user an X.509 short-lived certificate – but he still needs to be a member of a VO.
  - The gateway then registers the user into a dedicated VO.
    - DN persistency is guaranteed across credential re-creation.
- Sites then need to accept the additional VO.
  - It is also possible to only accept subgroups of the VO.
  - One could have a catch-all VO, or set-up multiple VOs if the need arises (operational / business considerations apply)
- Users have gained access to Grid resources (i.e. VOMS proxies) with minimal changes to the sites.
  - Apply this to job submission portals, and/or to Cloud web portals.

Integration: Grid users accessing Cloud resources <1>

- Really an application of the X.509 pass-through case
- The same services used in Grid computing for authentication and authorization are also used by the WNoDeS Cloud layer.
  - VOMS for Virtual Organization membership, gLite ARGUS for authorization policies
  - This allows us to automatically support existing Grid certificates and Virtual Organizations
  - Existing grid users are able to access Cloud resources, using just their Grid credentials.

Integration: Grid users accessing Cloud resources <2>

- User contacts the WNoDeS Cloud Web Interface (W-CWI), being authorized through a browser-installed X.509 certificate.
- A request is made by the user to create Cloud resources assigned (billed) to VO XYZ.
- W-CWI contacts the VOMS server for VO XYZ and validates user's credentials
- If successful, W-CWI contacts ARGUS to validate access policies
  - Might be e.g. per-VO, per-role, whitelist-based.
- If autentication and authorization are both OK, resource is granted.

Cloud and Networking

About Networking, there are some network topics to be considered.

- Customer wants to know what he is paying for, e.g the availability and reliability of the resources they are buying. For distributed clouds, the cloud provider should get at least part of this information from the network infrastructure provider.

- A cloud provider should be able to give his customers information about network quality and provide end- to-end paths with clear network QoS.

- WNoDeS-based cloud providers would ideally like to see the network transport layer as transparently as possible

Some network-related topics we would like to work and possibly contribute on:

- definition of network-related quality metrics and associated network quality of services

- test and definition of on-demand circuits

- test of high-volume data transfer across distributed clouds

Conclusions

The INFN WNoDeS project aims to:

- Exploit existing infrastructures and previous investments.
- Support established frameworks (e.g. Grid Computing in scientific communities) and new paradigms (e.g. Cloud computing)
- Interoperate with multiple and diverse distributed computing infrastructure projects through the adoption of open standards

Work is ongoing on the definition/implementation of key areas:

- Performance metrics associated to virtualization technologies.
- Standard authentication and authorization mechanisms (currently not implemented in the OCCI standard).

Network services must be as trasparent as possible for both resource providers and users.

Project coordinates:
http://web.infn.it/wnodes (email: wnodes@lists.infn.it)