

WNoDeS: the Open Cloud Computing
Interface implementation at INFN for a
integrated Grid/Cloud environment.

Peter Solagna

INFN

WNoDeS

WNoDeS quick facts:

- ✓ Running in production at the INFN Tier-1 for the past 12 months.
- ✓ Being installed at the INFN LNL National Laboratory.
- ✓ Currently running 2,000 on-demand VMs at the INFN Tier-1.
- ✓ Fully integrated into the 7,000 cores Tier-1 farm.
- ✓ Integrated local, grid, cloud access plus instantiation of Virtual Interactive Pools, all out of a common resource set.
- ✓ Supporting several key features like VLAN partitioning, integration with shared storage, multi-core VMs, network throttling.

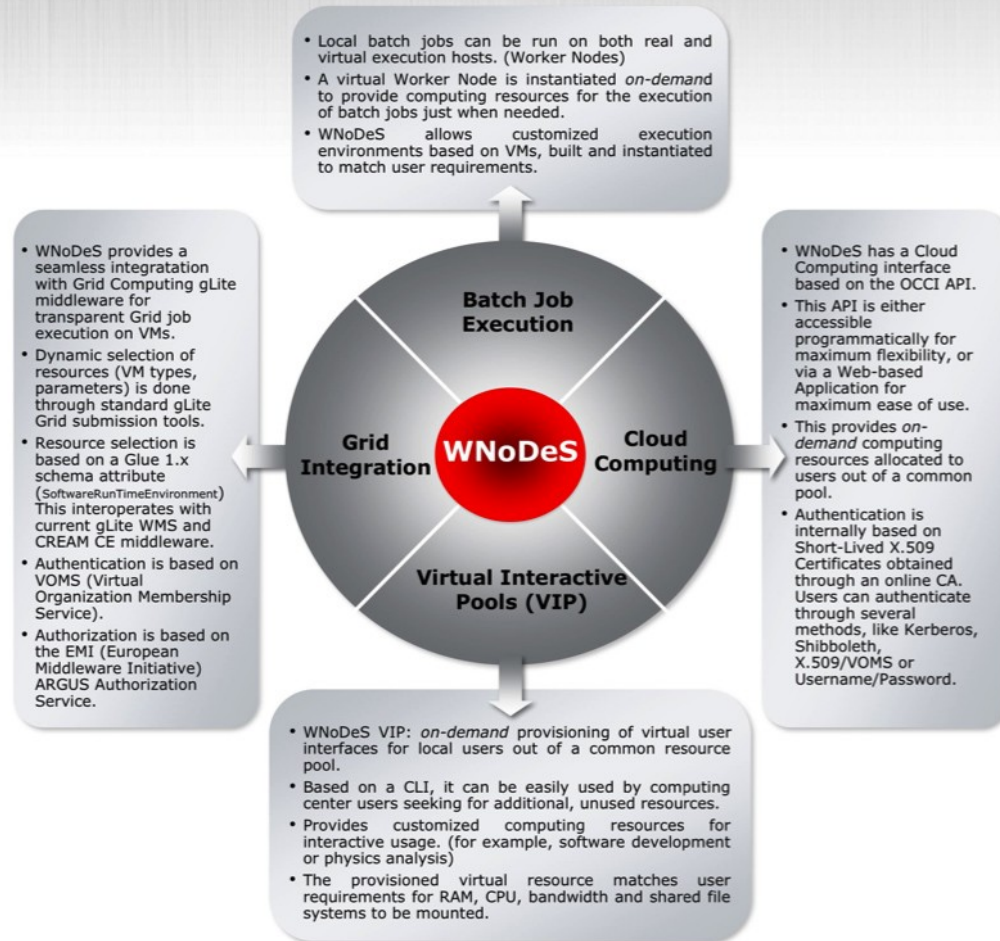
✓ **Web-site:** <http://web.infn.it/wnodes>

✓ **E-mail:** wnodes@lists.infn.it

WNoDeS Worker Nodes on Demand Service



Dynamic Provisioning of Integrated Grid/Cloud Virtual Environments



WNoDeS Fact Sheet

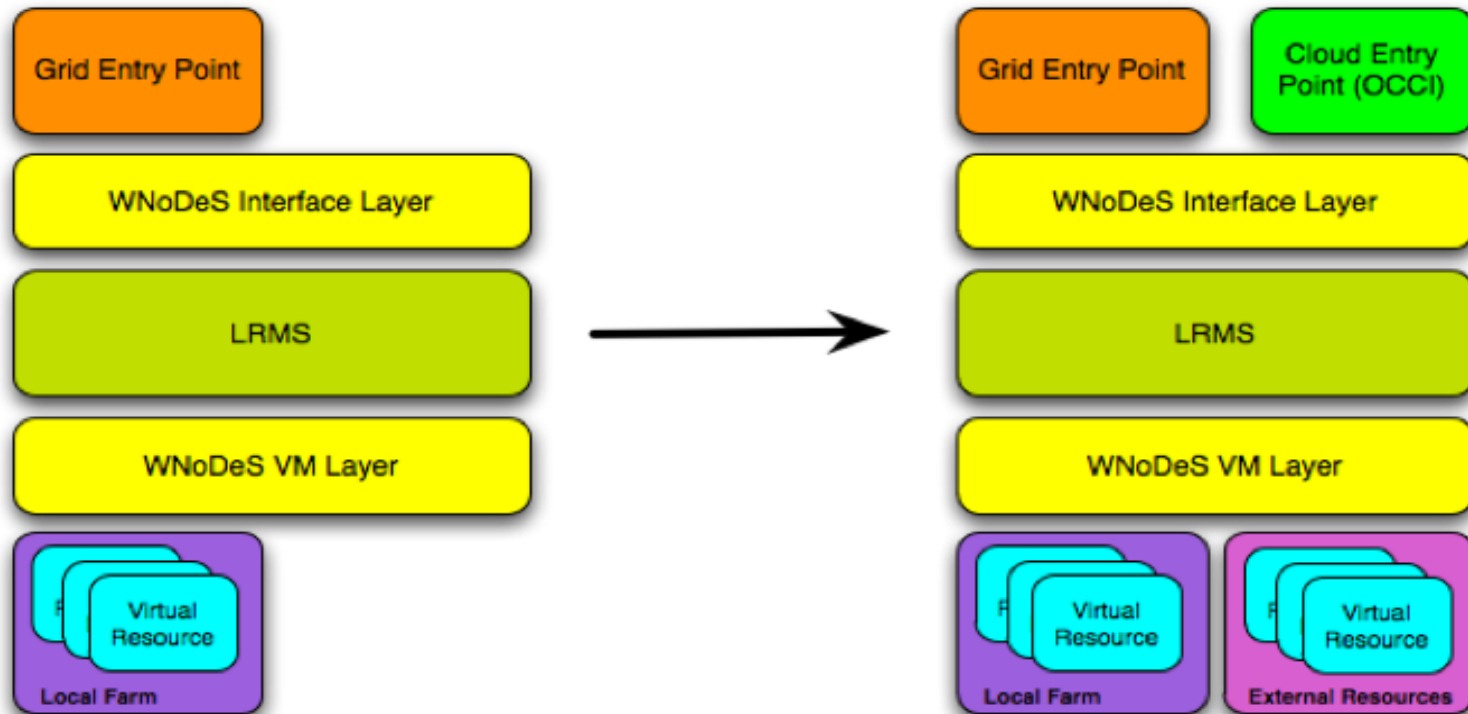
Deployment

WNoDeS is in production at the INFN Tier-1 Computing Center (Bologna - Italy), fully integrated with its 7000-cores farm. 2000 VMs (9/2010) can currently be instantiated *on-demand* out of this common farm.

Main Features

- ✓ Grid/Cloud Full Integration
- ✓ Authentication Gateway
- ✓ VM image selection
- ✓ VirtIO support
- ✓ Network throttling
- ✓ V-LAN support
- ✓ libguestFS support
- ✓ Multi-core VMs
- ✓ Cloud: OCCI API
- ✓ Cloud: Web Console
- ✓ Integration with shared distributed file systems (INFN Tier-1: IBM GPFS)

The WNoDeS architecture



- Production deployment 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

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 <1>

Compute attributes	Compute actions
architecture	start
cores	stop
hostname	restart
speed	suspend
memory	
status	

Network attributes	Network actions
vlan	down
gateway	up
address	
allocation	

Storage attributes	Storage actions
size	backup
status	online
	offline
	resize
	snapshot

OK/Supported

Develop/Deploy

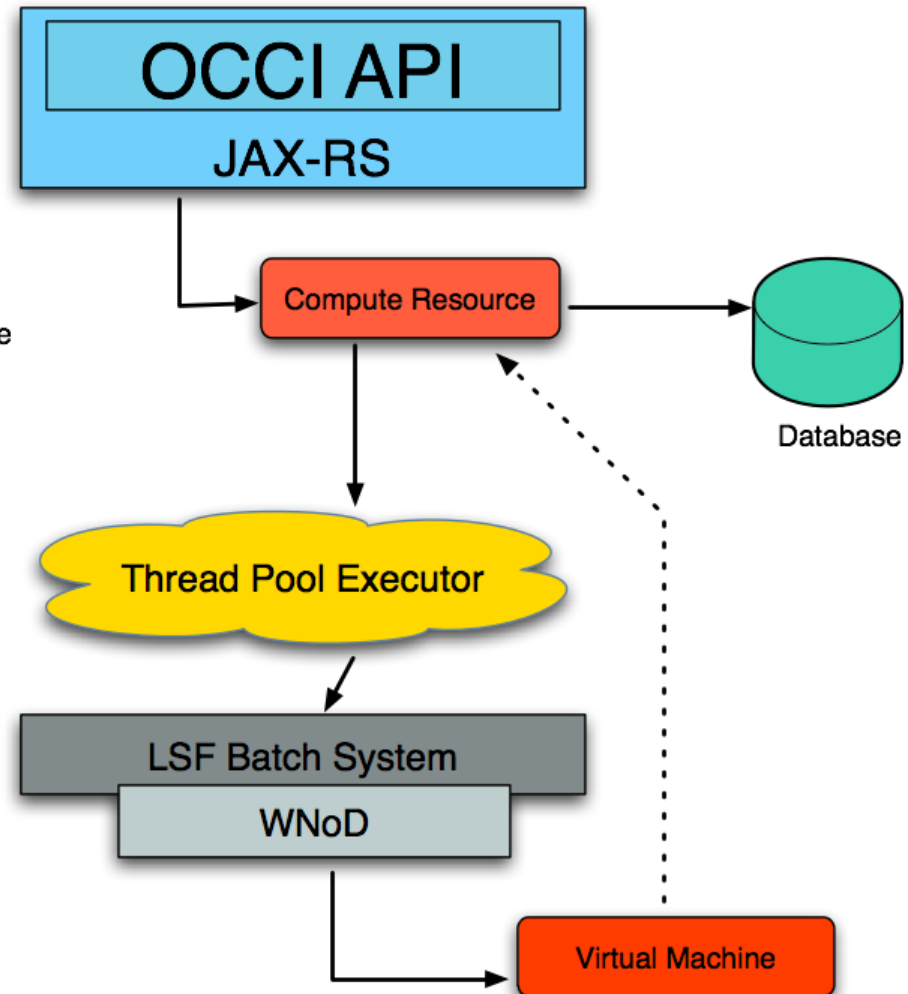
Design/Implement

The WNoDeS OCCI Implementation <2>



```
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 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:

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

- 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.

- 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.

- 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 authentication and authorization are both OK, resource is granted.

- The INFN WNoDeS project aims to:
 - Exploit existing infrastructures and previous investments
 - Support established frameworks (Grid) and new paradigms (Cloud)
 - Interoperate with diverse computing infrastructures through the adoption of open standards.

Cloud users accessing grid resources:

- Easy access of grid resources by cloud/federated users.
- Exploitation of previous investments in distributed Grid infrastructures.

Grid users accessing cloud resources:

- Easy access of cloud resources by existing Grid users.
- Grid VOs can create ad-hoc rules for users allowed to instantiate and consume Cloud resources.

Web:<http://web.infn.it/wnodes> email:wnodes@lists.infn.it