



Technical Activities

StratusLab Kick-off Meeting, 14 – 15 June 2010



The StratusLab project is partially funded by the European Commission through Grant Agreement RI-261552



What?

• **Technical**

- Track progress of development and infrastructure activities
- Prioritize features, keep sprints and product backlogs
- Coherent technical view across WPs
- Make technical decisions

• **Administrative**

- Homogeneous technical vision across deliverables
- Track progress of technical deliverables
 - D4.1+4 (architecture), D4.3+6 (Integration),
 - D5.1, D5.2 (Infrastructure Specification & Policies)
 - D6.1 (design of Cloud-like interfaces)
 - D3s (interaction with other projects)

Who?

- Activity Leader(Cal,David,Meb,Vangelis,Juan) + Person nominated by PC (Ruben)

When?

- Regular meetings, every 3-4 weeks
- May **alternate** administrative/technical meetings
- On-demand if urgent issues arises

Development Sprint


- WP4 Integration & Distribution
- WP6 New features
- WP5 Image Appliances
- **Input**
 - Sprint Backlog
- **Output:**
 - Working increment of the StratusLab Distribution

Integration Sprint

- WP5 Infrastructure Operation
- **Input**
 - StratusLab Distro
- **Output:**
 - Devel testbed
 - Operational infrastructure*
 - Image repository
 - Feedback & issues

TSCG Meetings

- **Input**
 - PMB strategic vision
 - WP2 deliverables
 - WP5 feedback
- **Output:**
 - Updated Product Backlog
 - Sprint Backlog



Roadmap for the Group (short-term)

- **Defining process**
 - Implementation of the development cycle (All)
 - Packaging & Deployment process (WP5,WP4)
 - Development process (WP4,WP6)
- First version of **StratusLab Architecture** (All)
- **Product Backlog** (All)
- **Data model & policies** for StratusLab appliances (WP5)



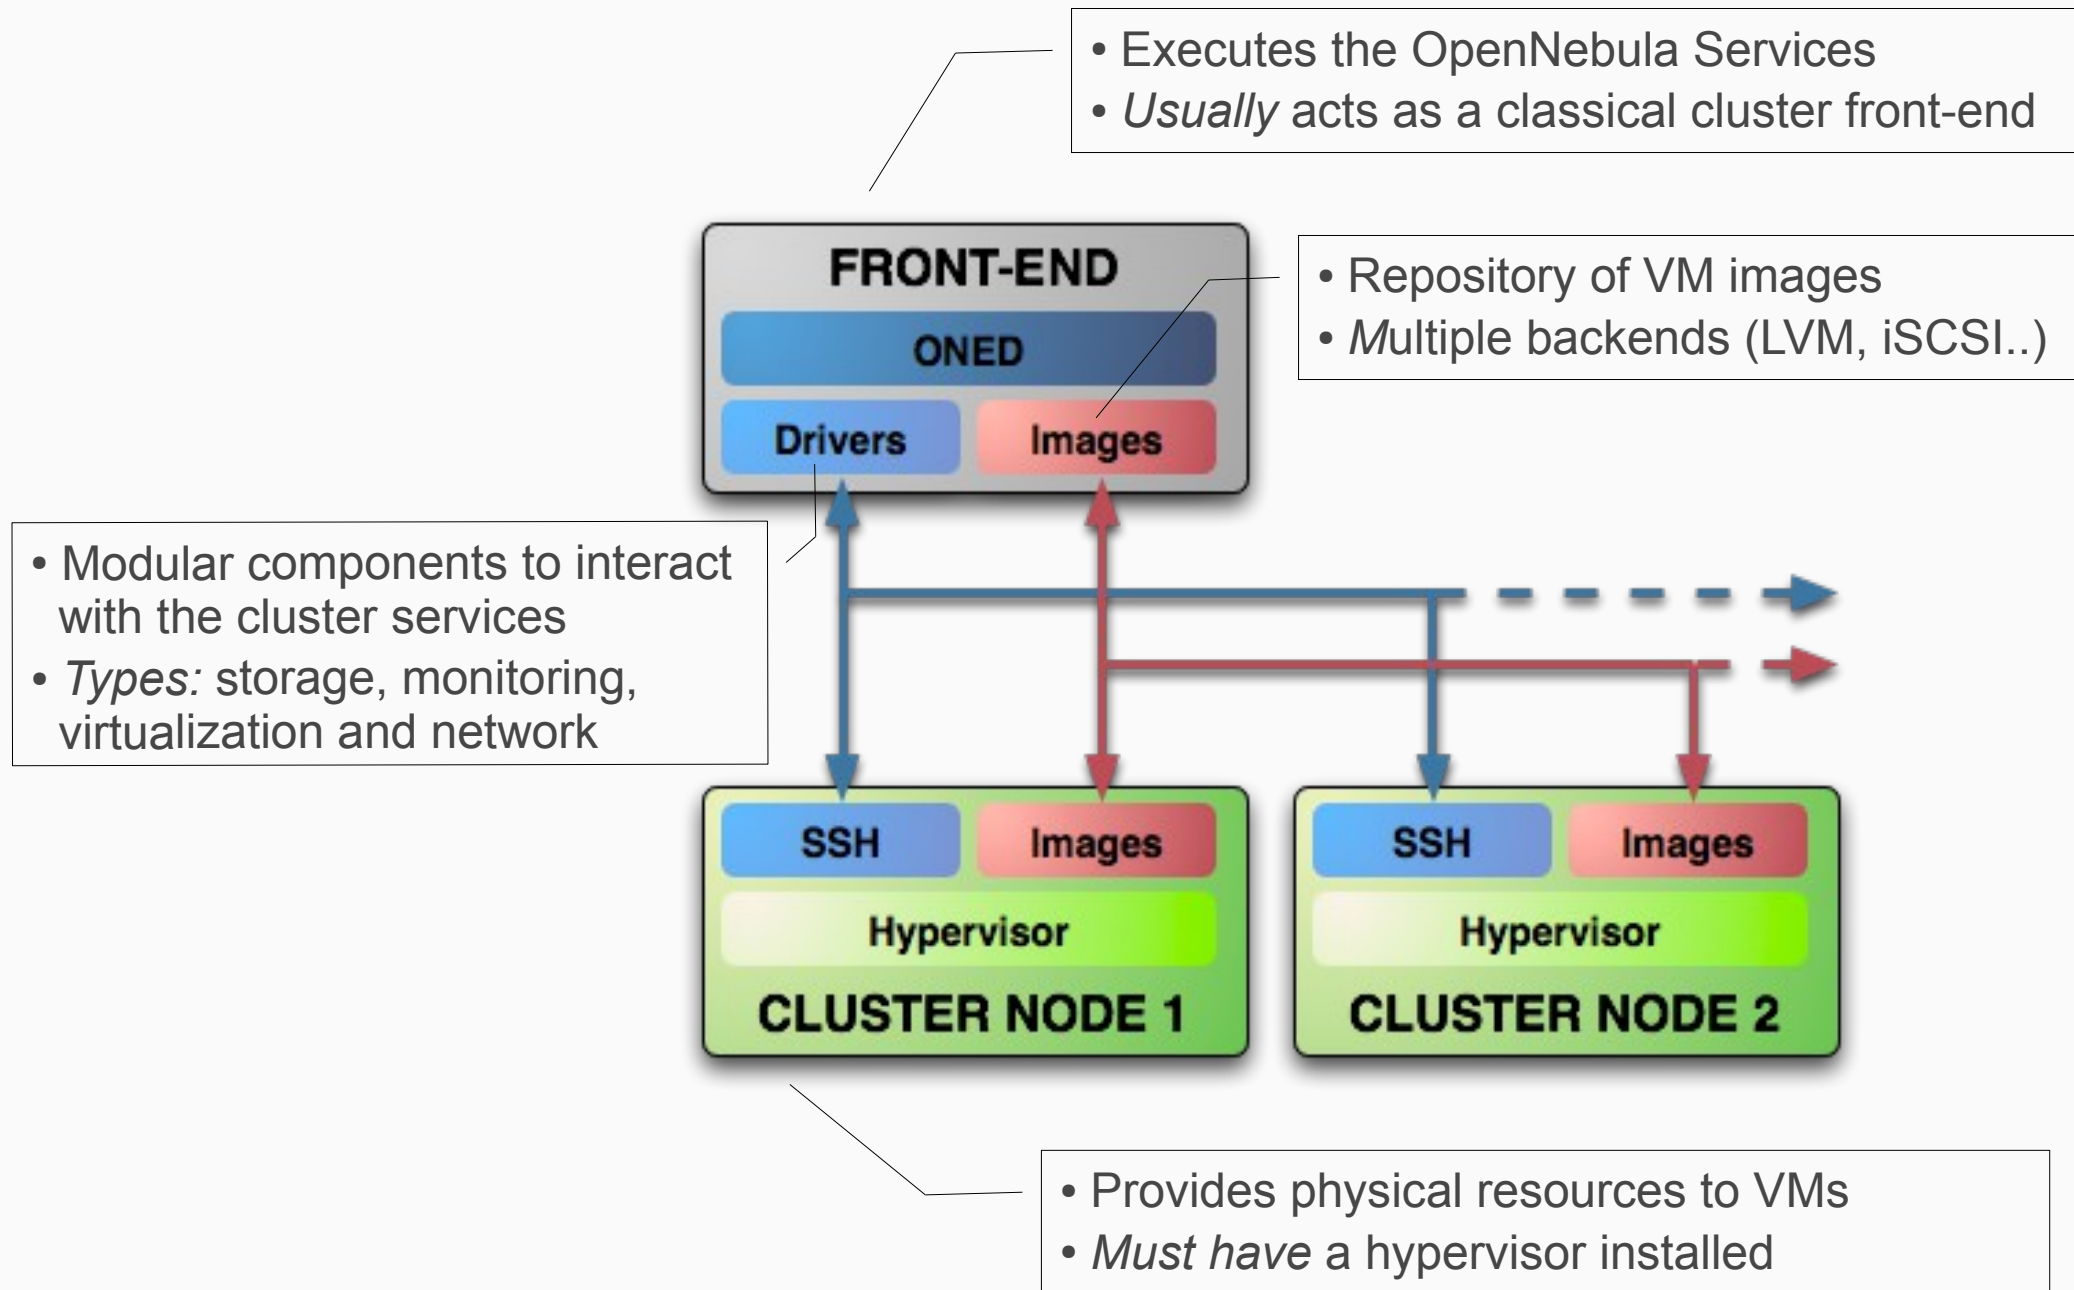
Brief Overview of OpenNebula

StratusLab Kick-off Meeting, 14 – 15 June 2010

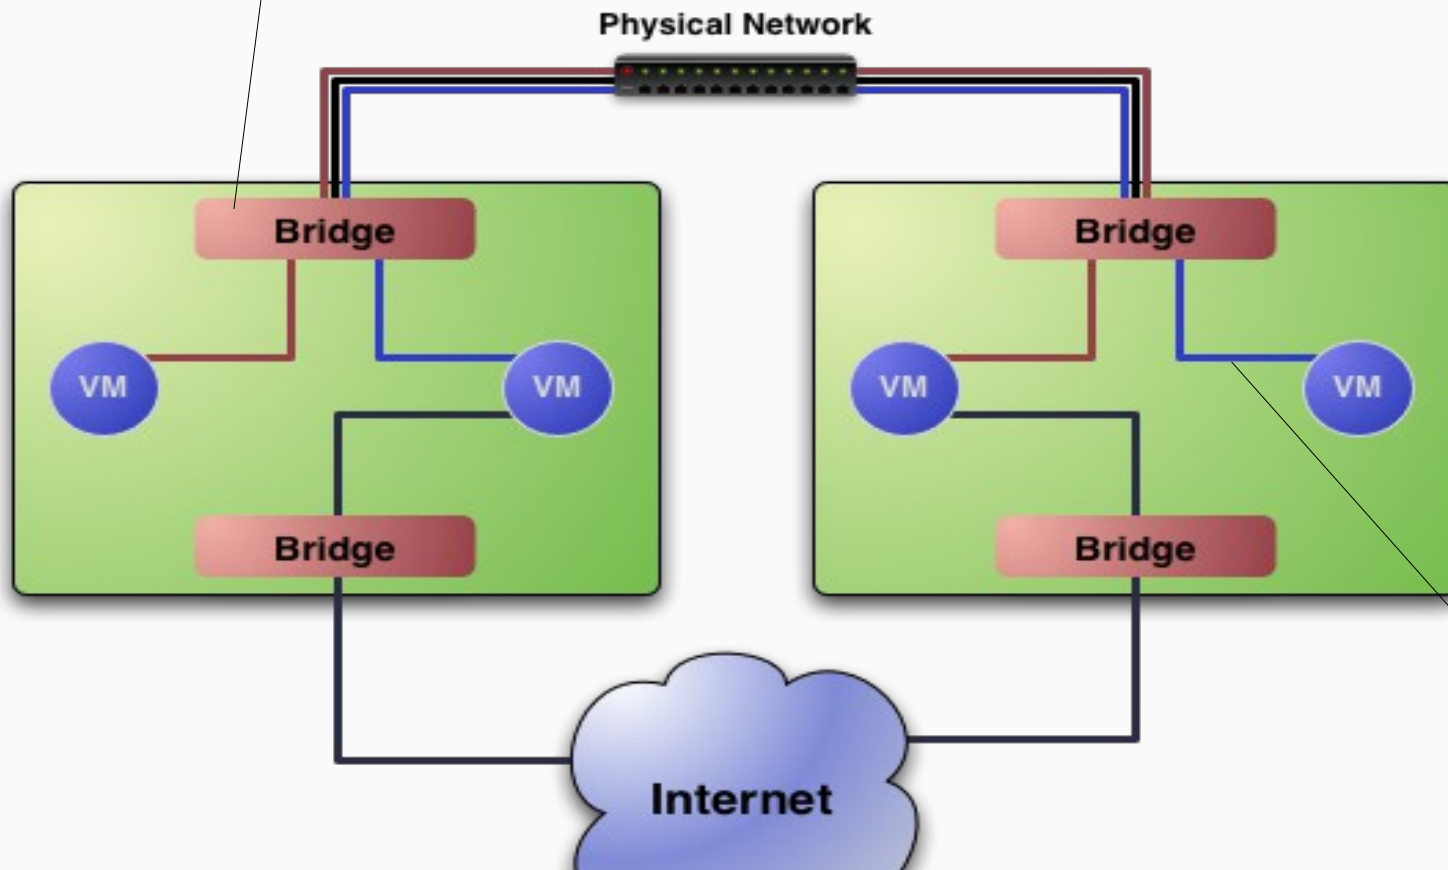


The StratusLab project is partially funded by the European Commission through Grant Agreement RI-261552





- OpenNebula uses bridge networking
- NAT, firewalling and other services are configured with hooks



- Virtual Network: MAC/IP address space + bridge binding
- Networks are isolated at layer 2 (IEEE 802.1Q, ebtables)
- You can put any TCP/IP service (e.g. DHCP, nagios...)

Using Virtual Networks with VMs

- Define VM NICs attached to a given virtual network. The VM will get a NIC with a free MAC in the network

```
#A VM with two interfaces each one in a different vlan
```

```
NIC=[NETWORK="Blue LAN"]
```

```
NIC=[NETWORK="Red LAN"]
```

```
#Ask for a specific IP/MAC of the Red vlan
```

```
NIC=[NETWORK="Red LAN", IP=192.168.0.3]
```

- Prepare the VM to use the IP. Sample scripts to set the IP based on the MAC are provided for several Linux distributions.

IP-MAC address correspondence

IP:

10. 0. 1. 2
↓ ↓ ↓ ↓

MAC: 02: 01: 0A: 00: 01: 02

oned.conf


IP Address

Defining a Virtual Machine...


- A **capacity** in terms memory and CPU
- A set of **NICs** attached to one or more virtual networks
- A set of **disk images**, to be “*transferred*” to/from the execution host.
- A **state file** (optional) or recovery file, with the memory image of a running VM plus some hypervisor specific information.
- Virtual Machines are defined in a VM template
- Each VM has an unique ID in OpenNebula the VM_ID
- All the files (logs, images, state files...) are stored in `$ONE_LOCATION/var/<VM_ID>`

The Virtual Machine Context

- Custom data to be passed to the VM at boot time
- Boot Process
 - Mount an iso image with context data
 - Source context.sh to look for variables
 - Access any file to setup VM services (e.g. ssh keys...)



- VMs can be local or remote
- VM connectivity has to be configured, usually VPNs



- External Clouds are like any other host
- Placement constraints

Using the EC2 Cloud

- Several accounts or zones can be configured
- The capacity allocated in EC2 can be limited
- VMs must be prepared to be instantiated locally or in the EC2
 - The template must provide a description for both instantiation methods.
 - The EC2 counterpart of your VM (`AMI_ID`) must be available for the driver account

```
EC2 = [  
  AMI           = "ami_id for this VM",  
  KEYPAIR      = "the keypair to use the instance",  
  AUTHORIZED_PORTS = "ports to access the instance",  
  INSTANCETYPE = "m1.small...",  
  ELASTICIP    = "the elastic ip for this instance",  
  CLOUD        = "EC2 cloud to use"  
]
```



Architecture Discussion

StratusLab Kick-off Meeting, 14 – 15 June 2010

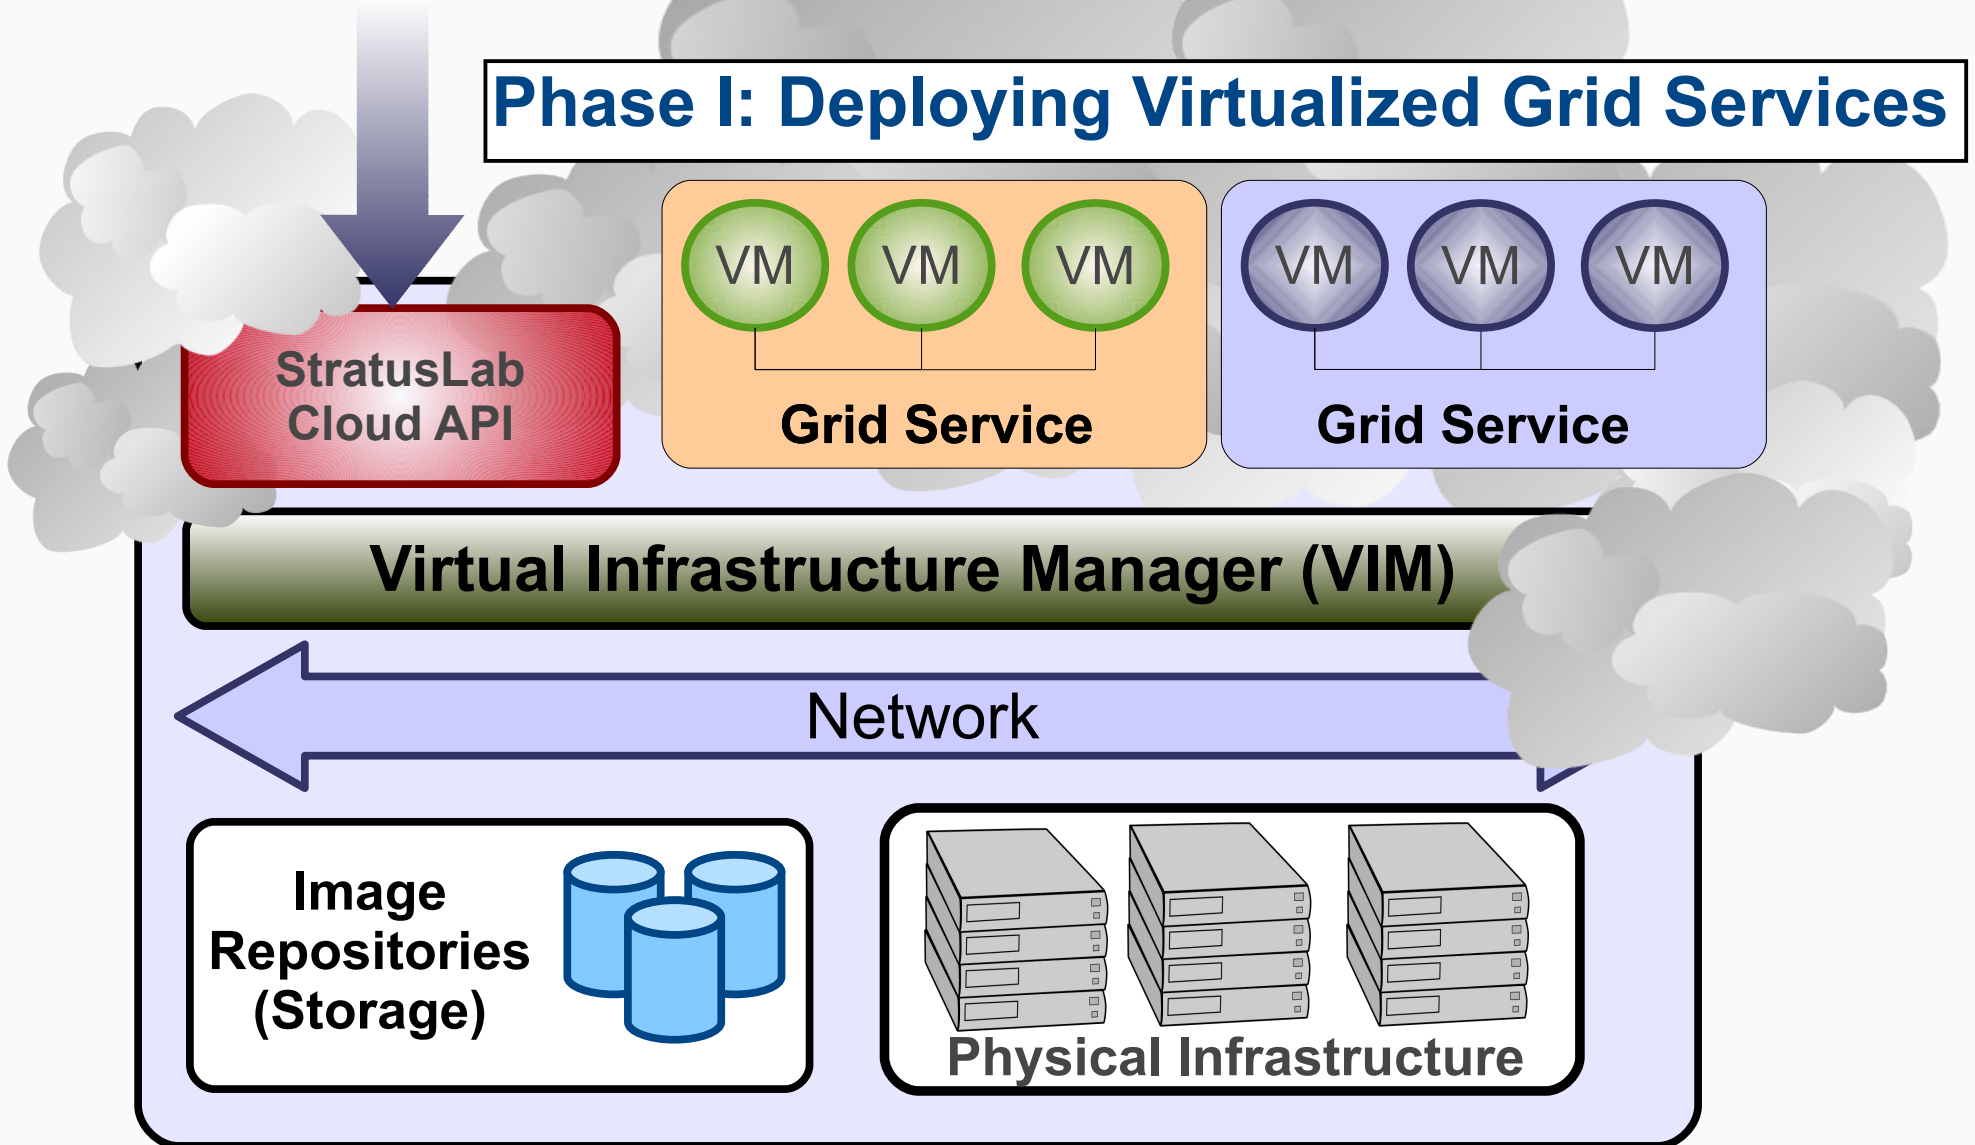


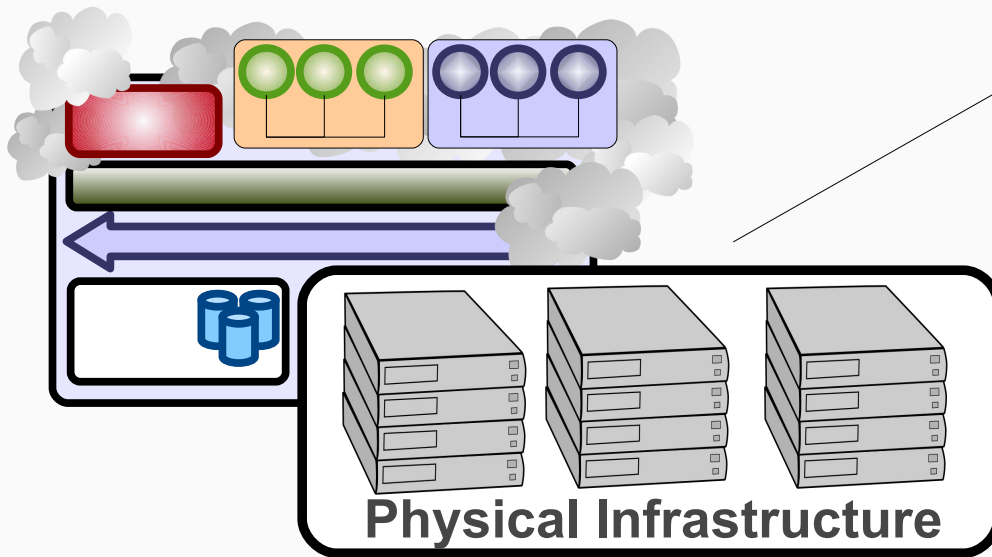
The StratusLab project is partially funded by the European Commission through Grant Agreement RI-261552



Phase II: Cloud-like Interfaces

Phase I: Deploying Virtualized Grid Services





Operating System

- Linux distribution
 - CentOS
 - Ubuntu
 - Scientific Linux
- Requirements
 - sudo
 - sshd
 - oneadmin account
 - FS related (e.g. NFS)

Hypervisor

- KVM + libvirt
- Xen
- Requirements
 - HW virtualization
 - special drivers (virtio...)

Connectivity

- No public IP needed
- Internet attached
- Bridged networking

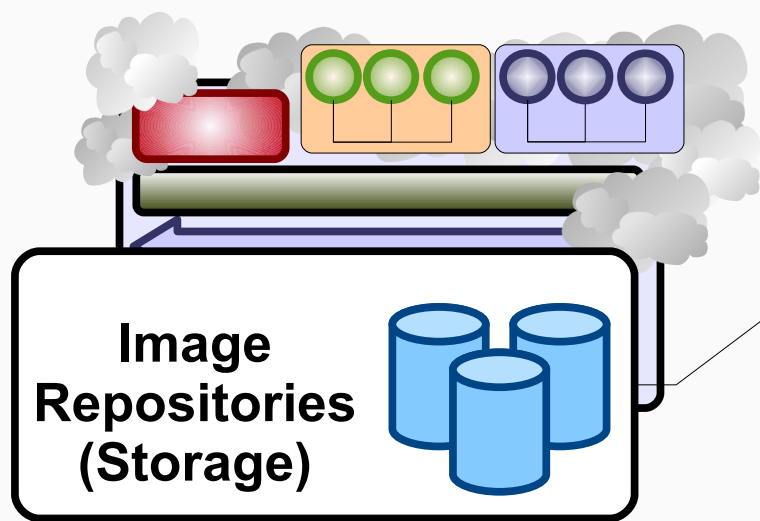


Image Distribution


- Dedicated VO Repositories (off-site)
- On-site Repository (caches VO images)
- Transfer of images to workernodes
 - Shared/Distributed FS
 - LVM-based
 - scp

Image Creation

- Format (raw, qow...)
- FS conventions (e.g. partitions, swap)
- Persistent DataBlocks (ala EBS)

Image Contextualization

- Data model:
 - pre-defined & user variables (e.g. IPs)
 - specific files (e.g. ssh keys)
 - software packages
- Access conventions (e.g. context dev)

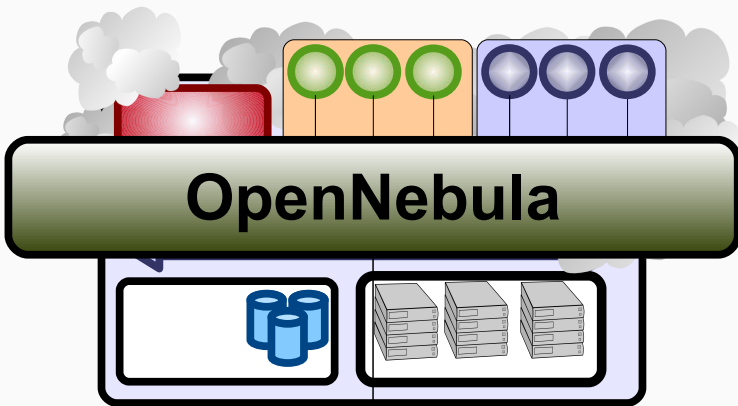


Firewalling & Security

- Site needs to open VM services
- VM access controlled “from the outside”
- Layer 2 isolation (Phase II)

IP Addresses & Networks

- Private network for Grid Services VMs
- Public network with a pool of public IPs
- Getting the IP & hostname
 - MAC2IP (preferred)
 - DHCP using dedicated network VMs
- IPv4



OpenNebula Configuration


- Deployment (dedicated machine, VM)
- SW requirements (mysql, sqlite)
- Tuning (monitor intervals, threads...)
- Simplified configuration process
- Management Interface

Scheduling

- VM placement policies:
 - Packing
 - Stripping
 - Requirements

Commission Requirements

- Keep an eye on other cloud solutions



Grid Service Virtualization

- Service Type:
 - Computing Element
 - Workernode
 - Storage Resource Manager
 - ...
- Certificate Management

Cluster Virtualization

- LRMS management (WN registration)
- Queue management
- Elasticity Management