

OpenNebula/Reservoir Training, January 27-28

Brussels, Belgium

# Session 1

## Introduction, Installation and Configuration


Daniel Molina & Javier Fontán  
dmolina/jfontan@opennebula.org

OpenNebula.org



Copyright 2002-2011 © OpenNebula Project Leads (OpenNebula.org). All Rights Reserved.  
Creative Commons Attribution Share Alike (CC-BY-SA)


# Cloud Computing in a Nutshell



# The IaaS Clouds a Four Point Check List

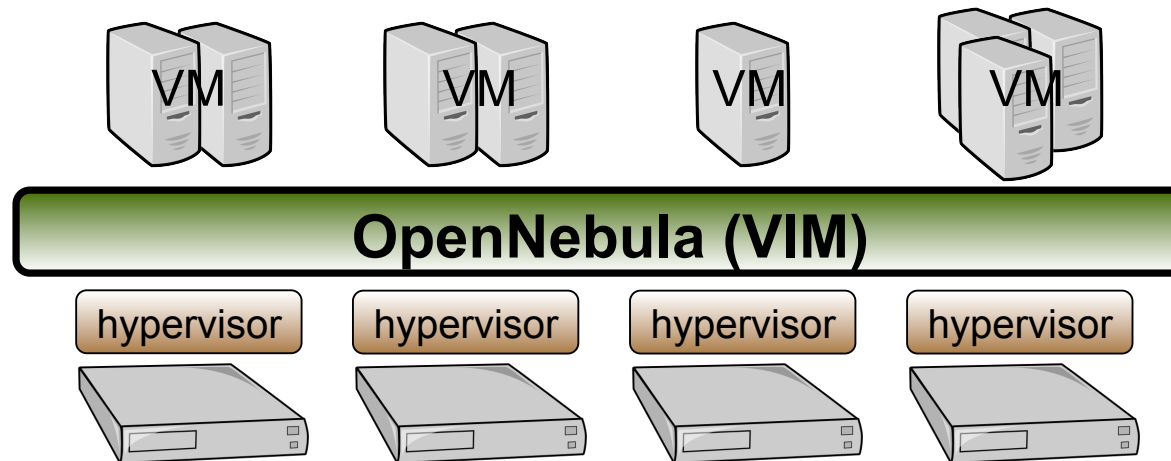
- Simple Interface
- Raw *Infrastructure* Resources
  - Total control of the resources
  - Capacity leased in the form of VMs
  - Complete Service-HW decoupling
- Pay-as-you-go
  - A single user can not get all the resources
- Elastic & “*infinite*” Capacity

# The Anatomy of an IaaS Cloud

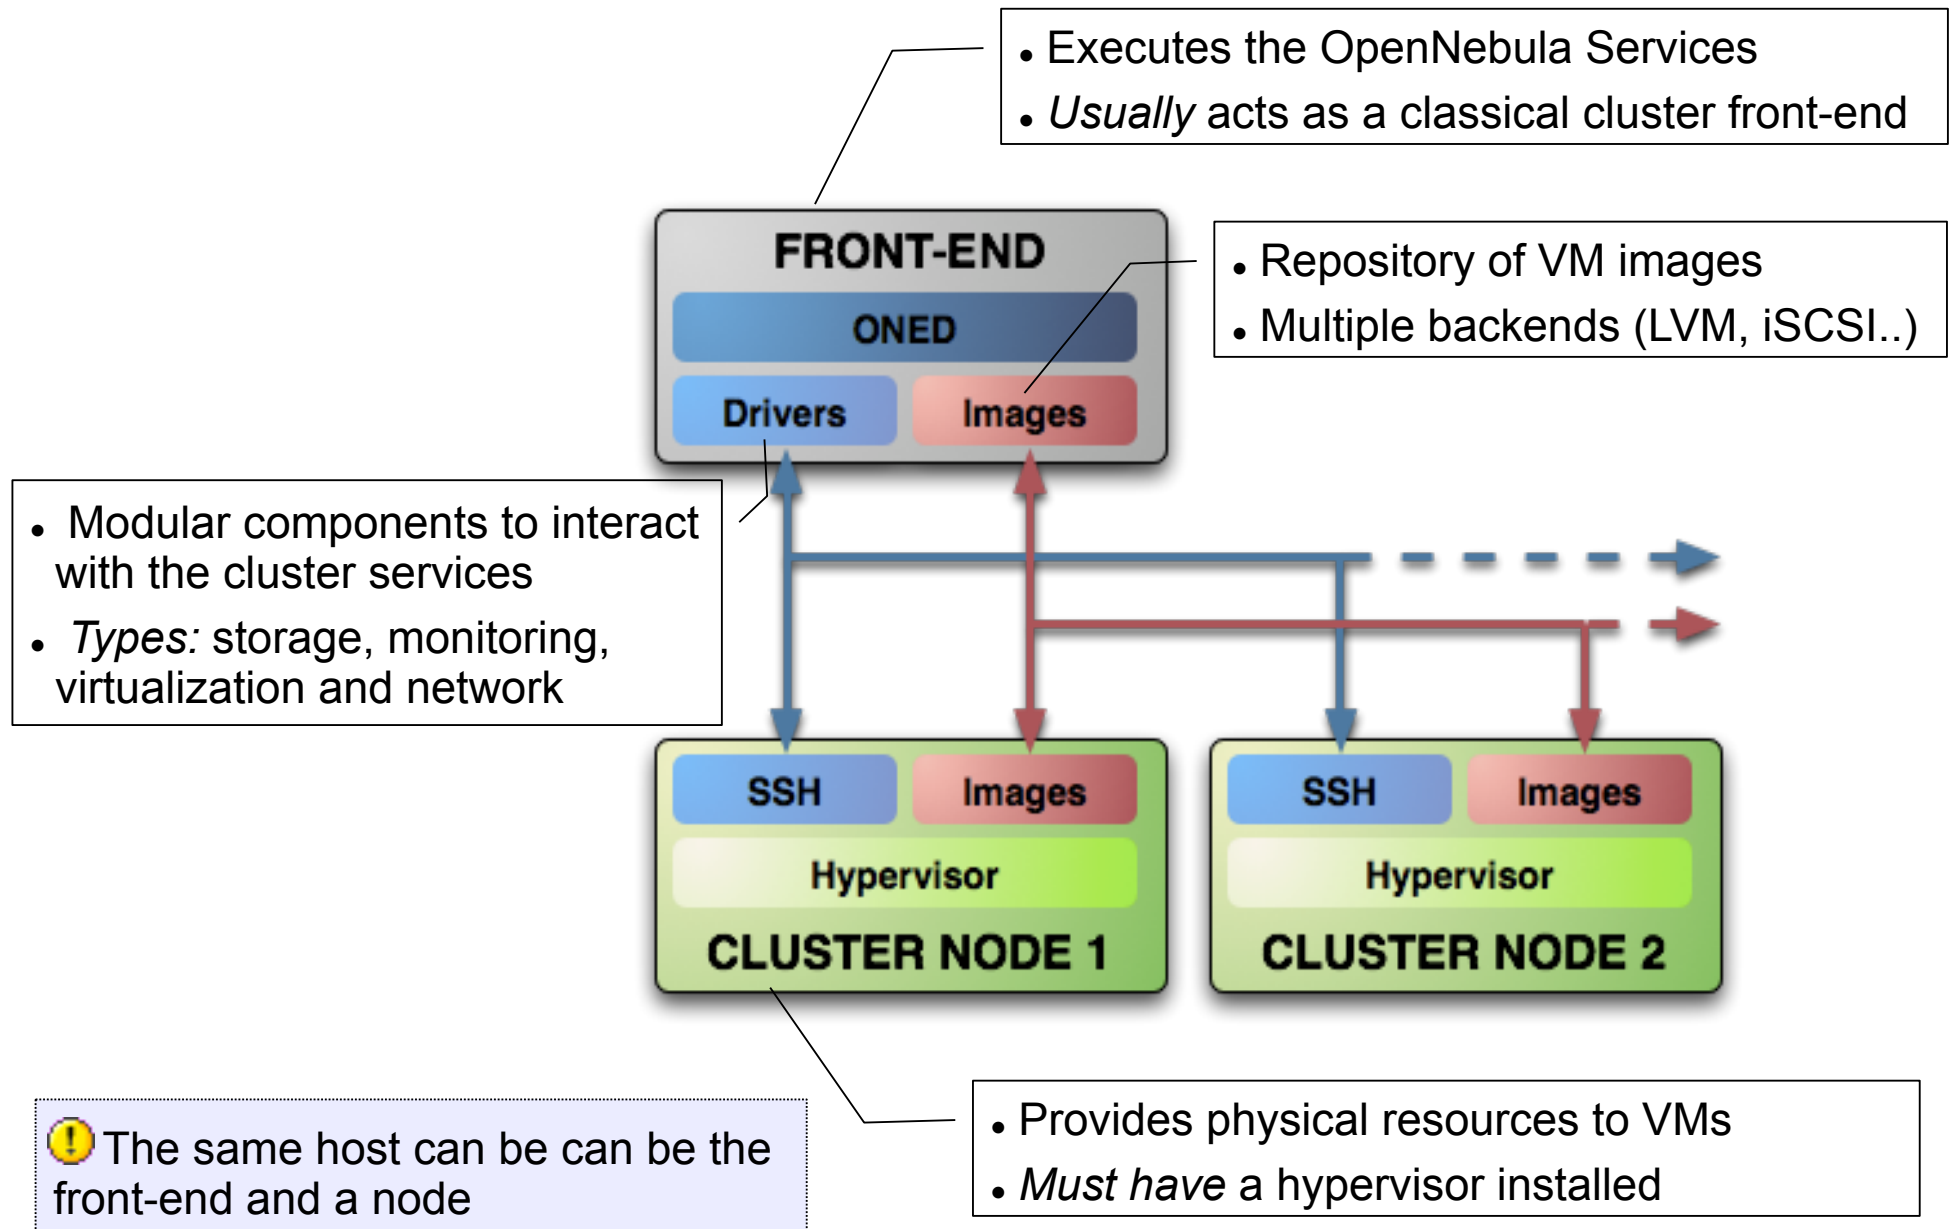


# Why a Virtual Infrastructure Manager?

- VMs are great!!...but something more is needed
  - Where did/do I put my VM? (**scheduling & monitoring**)
  - How do I provision a new cluster node? (**clone & context**)
  - What MAC addresses are available? (**networking**)
- Provides a **uniform view** of the resource pool
- **Life-cycle management** and monitoring of VM
- The VIM **integrates** Image, Network and Virtualization



# Planning the Installation: System Overview



## Planning the Installation: Working in the Front-End ...

- Choose your installation mode
  - system wide (/usr, /etc...)
  - *self-contained* (under \$ONE\_LOCATION)
- Install software dependencies.
  - Check the documentation for platform specific notes installation nodes


<http://opennebula.org/documentation:rel2.0:notes>

- Dependencies already installed in the Front-End and the Nodes

## Planning the Installation: Working in the Front-End ...

- The Users of the private cloud:
  - oneadmin: Account to run the daemons, manage the system and do all the low-level operations (e.g. start VMs, move images...).
  - Regular users: create and manage their own VMs and networks.  
*Need to be defined in OpenNebula*
- Installation layout
  - We will use the /srv/cloud directory to place the OpenNebula software
  - /srv/cloud/one will hold the OpenNebula installation

```
# tree /srv
/srv
|-- cloud
    |-- one
        |-- SRC
```

 The oneadmin account must be created system wide (i.e. front-end and all the nodes). You can use NIS, or a local account with the same ID's in all the hosts. Regular users do not need a UNIX account in the nodes, nor in the front-end.



## Planning the Installation: Working in the Front-End ...

- Hands on...

```
Fe$ su -
```

```
fe# groupadd -g 9000 oneadmin
```

```
fe# mkdir /srv/cloud
```

```
fe# useradd -d /srv/cloud/one -g oneadmin -u 9000 -s /bin/bash -m  
oneadmin
```

*Create the file-system hierarchy with the oneadmin account*

```
fe# su - oneadmin
```

```
fe$ id
```


```
uid=9000(oneadmin) gid=9000(oneadmin) groups=9000(oneadmin)
```

```
fe$ mkdir SRC
```

*We will place the OpenNebula source code in SRC*


## Planning the Installation: Working in the Front-End ...

- Preparing the storage for the private cloud...
  - **Image Repository:** Any storage medium for the VM images (usually a high performing SAN)
    - OpenNebula supports multiple back-ends (e.g. LVM for fast cloning)
    - The front-end must have access to the repository
  - **VM Directory:** The home of the VM in the cluster node
    - Stores checkpoints, description files and VM disks
    - Actual operations over the VM directory depend on the storage medium
    - Should be shared for live-migrations
    - You can go on without a shared FS and use the SSH back-end
    - Defaults to `$ONE_LOCATION/var/$VM_ID`

 **Dimensioning the Storage...** Example: A 64 core cluster will typically run around 80VMs, each VM will require an average of 10GB of disk space. So you will need ~800GB for `/srv/cloud/one`, you will also want to store 10-15 master images so ~200GB for image repository. A 1TB `/srv/cloud` will be enough for this example setup.


## Planning the Installation: Working in the Front-End ...

- In this course we will use NFS to share the VM directories
- The Image Repository is `/srv/cloud/one/var/images`



# Planning the Installation: Working in the Front-End ...

- Networking for the private cloud
  - OpenNebula management operations use ssh connections, it does not require a performing NIC
  - **Image traffic**, may require the movement of heavy files (VM images, checkpoints). Dedicated storage links may be a good idea
  - **VM demands**, consider the typical requirements of your VMs. Several NICs to support the VM traffic may be a good idea
  - OpenNebula relies on bridge networking for the VMs



## Planning the Installation: Working in the Front-End ...

- Prepare NFS

```
Export /srv/cloud to your nodes  
- only need /srv/cloud/one/var  
- we also export $HOME of oneadmin for easy SSH key configuration
```

```
fe# vi /etc/exports  
/srv/cloud cetic-nodeXX(rw,async,no_subtree_check,no_root_squash)  
  
fe# /etc/init.d/nfs reload
```

## Planning the Installation: Working in the Nodes ...

- Install software dependencies
  - We need SSH daemon running in the cluster nodes (check it!)
  - Runtime dependencies:
    - Ruby 1.8.x
- Users
  - Create the oneadmin account (**use same UID and GID**)

```
no# mkdir -p /srv/cloud
no# groupadd -g 9000 oneadmin
no# useradd -d /srv/cloud/one -g oneadmin -u 9000 -s /bin/bash oneadmin
```

- Add oneadmin to sudoers

```
no# tail -1 /etc/sudoers
oneadmin ALL=(ALL) ALL, NOPASSWD: /usr/sbin/xm, /usr/sbin/xmtop
```

# Planning the Installation: Working in the Nodes ...

- Storage
  - Recreate the installation layout and configure NFS to mount VM dirs

```
no# chown oneadmin:oneadmin /srv/cloud

no# vi /etc/fstab
frontend:/srv/cloud /srv/cloud nfs soft,intr,rsize=32768,wsiz=32768,rw
0 0

no# mount /srv/cloud
```

# Planning the Installation: SSH Configuration

- Enable password-less SSH access to cluster nodes for the oneadmin account:

```
DO NOT PROTECT PRIVATE KEY WITH A PASSWORD
```

```
fe$ ssh-keygen
```

```
Generating public/private rsa key pair.
```

```
...
```

```
Enter passphrase (empty for no passphrase):
```

```
Enter same passphrase again:
```

```
fe$ cp ~/.ssh/id_rsa.pub ~/.ssh/authorized_keys
```

```
Tell ssh client not to ask to add hosts to known_hosts (optional)
```

```
fe$ cat /srv/cloud/one/.ssh/config
```

```
Host *
```

```
    StrictHostKeyChecking no
```

```
TEST!
```


```
fe$ ssh localhost
```

```
fe$ ssh host01
```



# Planning the Installation: The Hypervisor ...

- Installing the Hypervisor
  - OpenNebula supports KVM, Xen and Vmware (even *simultaneously*). This course applies to KVM and Xen
  - Refer to the hypervisor documentation for additional (and better information) on setting up them.
  - In this course, we will use XEN.



## Planning the Installation: The Hypervisor ...

- The software bridge is essential for having different VMs in the same host with connectivity
- Let's check the bridge in the hosts

```
no$ /usr/sbin/brctl show
Bridge name      bridge id          STP enabled      interfaces
virbr0           8000.000000000000  yes
xenbr0           8000.fefffffffffff  no               peth0
                                                         vif0.0
```

## Planning the Installation: The Hypervisor ...

- Test the installation for the oneadmin account

```
no$ sudo /usr/sbin/xm list
Name      ID Mem(MiB) VCPUs State   Time(s)
Domain-0  0      256      1 r----- 8.2
```

- This ensures that oneadmin is capable of running VMs

# Planning the Installation: Checklist

Software Requirements	
ACTION	DONE/COMMENTS
Installation type: self-contained, system-wide	self-contained
Installation directory	/srv/cloud/one
OpenNebula software downloaded to /srv/cloud/one/SRC	
sqlite, g++, scons, ruby and software requirements installed	
User Accounts	
ACTION	DONE/COMMENTS
oneadmin account and cloud group ready in the nodes and front-end	
Storage Checklist	
ACTION	DONE/COMMENTS
/srv/cloud structure created in the front-end	
/srv/cloud exported and accessible from the cluster nodes	
mount point of /srv/cloud in the nodes if different	VMDIR=<mount_point>/var/
Cluster nodes Checklist	
ACTION	DONE/COMMENTS
hostnames of cluster nodes	
ruby, sshd installed in the nodes	
oneadmin can ssh the nodes passwordless	