
OpenNebula.org

OpenNebula 4.14 User Guide

Release 4.13.80


OpenNebula Project

July 29, 2015

1	Virtual Resource Management	1
1.1	Introduction to Private Cloud Computing	1
1.2	Managing Virtual Networks	4
1.3	Managing Images	13
1.4	Creating Virtual Machines	23
1.5	Managing Virtual Machines	30
2	Virtual Machine Setup	47
2.1	Contextualization Overview	47
2.2	Adding Content to Your Cloud	47
2.3	Basic Contextualization	48
2.4	Advanced Contextualization	52
2.5	Windows Contextualization	57
2.6	Cloud-init	59
3	OpenNebula Marketplace	61
3.1	Interacting with the OpenNebula Marketplace	61
3.2	Howto Create Apps for the Marketplace	65
4	References	69
4.1	Virtual Machine Definition File	69
4.2	Image Definition Template	87
4.3	Virtual Network Definition File	89
4.4	Command Line Interface	92


VIRTUAL RESOURCE MANAGEMENT

1.1 Introduction to Private Cloud Computing



The aim of a Private Cloud is not to expose to the world a cloud interface to sell capacity over the Internet, but to **provide local cloud users and administrators with a flexible and agile private infrastructure to run virtualized service workloads within the administrative domain**. OpenNebula virtual infrastructure interfaces expose **user and administrator functionality for virtualization, networking, image and physical resource configuration, management, monitoring and accounting**. This guide briefly describes how OpenNebula operates to build a Cloud infrastructure. After reading this guide you may be interested in reading the guide describing how a hybrid cloud operates and the guide describing how a public cloud operates.

1.1.1 The User View



An OpenNebula Private Cloud provides infrastructure users with an **elastic platform for fast delivery and scalability of services to meet dynamic demands of service end-users**. Services are hosted in VMs, and then submitted, monitored and controlled in the Cloud by using Sunstone or any of the OpenNebula interfaces:

- *Command Line Interface (CLI)*
- XML-RPC API
- OpenNebula Ruby and Java Cloud APIs

Lets do a **sample session to illustrate the functionality provided by the OpenNebula CLI for Private Cloud Computing**. First thing to do, **check the hosts in the physical cluster**:

```
$ onehost list
ID NAME          RVM  TCPU  FCPU  ACPU  TMEM  FMEM  AMEM  STAT
0  host01         0    800   800   800   16G   16G   16G   on
1  host02         0    800   800   800   16G   16G   16G   on
```

We can then **register an image** in OpenNebula, by using `oneimage`. We are going to build an *image template* to register the image file we had previously placed in the `/home/cloud/images` directory.

```
NAME          = Ubuntu
PATH          = /home/cloud/images/ubuntu-desktop/disk.0
PUBLIC        = YES
DESCRIPTION   = "Ubuntu 10.04 desktop for students."
```

```
$ oneimage create ubuntu.oneimg
ID: 0

$ oneimage list
```

ID	USER	GROUP	NAME	SIZE	TYPE	REGTIME	PUB	PER	STAT	RVMS
1	oneadmin	oneadmin	Ubuntu	10G	OS	09/29 07:24:35	Yes	No	rdy	0

This image is now ready to be used in a virtual machine. We need to define a *virtual machine template* to be submitted using the `onemplate` command.

```
NAME = my_vm
CPU = 1
MEMORY = 2056

DISK = [ IMAGE_ID = 0 ]

DISK = [ type = swap,
         size = 1024 ]

NIC = [ NETWORK_ID = 0 ]
```

Once we have tailored the requirements to our needs (specially, CPU and MEMORY fields), ensuring that the VM *fits* into at least one of both hosts, let's submit the VM (assuming you are currently in your home folder):

```
$ onemplate create vm
ID: 0

$ onemplate list
  ID USER   GROUP   NAME                               REGTIME PUB
  0 oneadmin oneadmin my_vm          09/29 07:28:41 No
```

The listed template is just a VM definition. To execute an instance, we can use the `onemplate` command again:

```
$ onemplate instantiate 1
VM ID: 0
```

This should come back with an ID, that we can use to identify the VM for **monitoring and controlling**, this time through the use of the `onevm` command:

```
$ onevm list
  ID USER   GROUP   NAME          STAT CPU   MEM      HOSTNAME      TIME
  0 oneadmin oneadmin one-0         runn  0     OK       host01 00 00:00:06
```

The **STAT** field tells the state of the virtual machine. If there is an **runn** state, the virtual machine is up and running. Depending on how we set up the image, we may be aware of its IP address. If that is the case we can try now and log into the VM.

To **perform a migration**, we use yet again the `onevm` command. Let's move the VM (with VID=0) to `host02` (HID=1):

```
$ onevm migrate --live 0 1
```

This will move the VM from `host01` to `host02`. The `onevm list` shows something like the following:

```
$ onevm list
  ID USER   GROUP   NAME          STAT CPU   MEM      HOSTNAME      TIME
  0 oneadmin oneadmin one-0         runn  0     OK       host02 00 00:00:48
```

You can also reproduce this sample session using the graphical interface provided by Sunstone, that will simplify the typical management operations.

OpenNebula Sunstone Virtual Machines

oneadmin OpenNebula

Dashboard System Virtual Resources Virtual Machines Templates Images Files & Kernels Infrastructure Marketplace OneFlow Support

ID	Owner	Group	Name	Status	Host	IPs	VNC
14	John	BlueVDC	Frontend_0_(service_5)	RUNNING	dev2	192.168.1.13	
13	John	BlueVDC	DB_0_(service_5)	RUNNING	dev1	192.168.1.12	
12	John	BlueVDC	Worker_0_(service_4)	RUNNING	dev2	192.168.1.11	
11	John	BlueVDC	Master_0_(service_4)	RUNNING	dev1	192.168.1.10	
10	John	BlueVDC	Slave_2_(service_3)	RUNNING	dev2	192.168.1.9	
9	John	BlueVDC	Slave_1_(service_3)	RUNNING	dev1	192.168.1.8	
8	John	BlueVDC	Slave_0_(service_3)	RUNNING	dev2	192.168.1.7	
6	Doe	BlueVDC	Sinatra - Mongo	RUNNING	dev2	192.168.1.5	
5	BlueVDC-admin	BlueVDC	Testing Env	RUNNING	dev1	192.168.1.4	
4	BlueVDC-admin	BlueVDC	Prod Machine	RUNNING	dev2	192.168.1.3	

Showing 1 to 10 of 14 entries

15 TOTAL 13 ACTIVE 2 OFF 0 PENDING 0 FAILED

OpenNebula 4.8.0 by C12G Labs.

1.1.2 Next Steps

You can now read the different guides describing how to define and manage virtual resources on your OpenNebula cloud:

- [Virtual Networks](#)
- [Virtual Machine Images](#)
- [Virtual Machine Templates](#)
- [Virtual Machine Instances](#)

You can also install OneFlow to allow users and administrators to define, execute and manage multi-tiered applications composed of interconnected Virtual Machines with auto-scaling.

1.2 Managing Virtual Networks

A host is connected to one or more networks that are available to the virtual machines through the corresponding bridges. OpenNebula allows the creation of Virtual Networks by mapping them on top of the physical ones.

1.2.1 Virtual Network Model

A virtual network (VNET) definition consists of three different parts:

- The **underlying networking infrastructure** that will support the VNET. These sections will typically include the following values:

At-tribute	Description
BRIDGE	Device to bind the virtual and physical network, depending on the network driver it may refer to different technologies or require host setups.
VLAN	Set to YES to activate VLAN isolation when supported by the network drivers
VLAN_ID	Identifier for the VLAN

Warning: The attributes above depend on the networking technology (drivers) used by the hosts when they are created. Please refer to the specific networking guide for a complete description of each value.

- The **logical address space** available. Addresses associated to a VNET can be IPv4, IPv6, dual stack IPv4-IPv6 or Ethernet; and can define a non-continuous addresses range structured in one or more Address Ranges (AR), see below
- **Context attributes.** Apart from the addresses and configuration attributes used to setup the guest network additional information can be injected into the VM at boot time. These contextualization attributes may include for example network masks, DNS servers or gateways. The following values are recognize by the contextualization packages:
 - NETWORK_ADDRESS
 - NETWORK_MASK
 - GATEWAY
 - GATEWAY6
 - DNS

These attributes can be later used in the *Virtual Machine Contextualization*. For example:

```
CONTEXT = [
  DNS = "$NETWORK[DNS, NETWORK=Private]"
]
```

Note: You can add any arbitrary data to the VNET to later use it within the VMs or just to tag the VNET with any attribute.

The Address Range (AR)

The addresses available in a VNET are defined by one or more Address Ranges (AR). Each address range defines a continuous address range and optionally, configuration attributes (context or configuration) that will override those provided by the VNET.

IPv4 Address Range

Defines a continuous IPv4 range:

At-tribute	Description
TYPE	IP 4
IP	First IP in the range in dot notation.
MAC	Optional. First MAC, if not provided it will be generated using the IP and the MAC_PREFIX in <code>oned.conf</code> .
SIZE	Number of addresses in this range.

```
#Example of an IPv4 AR, template form
AR=[
  TYPE = "IP4",
  IP   = "10.0.0.150",
  MAC  = "02:00:0a:00:00:96",
  SIZE = "51",
]
```

IPv6 Address Range

IPv6 networks assumes the auto-configuration process defined in the IPv6 standard, so the host id part is derived by the MAC address (or randomly). The global and ULA prefix needs to be provided.

Attribute	Description
TYPE	IP6
MAC	Optional. First MAC, if not provided it will be generated randomly.
GLOBAL_PREFIX	Optional. A /64 globally routable prefix
ULA_PREFIX	Optional. A /64 unique local address (ULA) prefix corresponding to the fd00::/8 block
SIZE	Number of addresses in this range.

```
#Example of an IPv6 AR, template form
AR=[
  TYPE= "IP6",
  MAC = "02:00:40:55:83:ff",
  GLOBAL_PREFIX = "2001:a::",
  ULA_PREFIX    = "fd01:a:b::",
  SIZE="1000"
]
```

Note: You can define either GLOBAL_PREFIX or ULA_PREFIX, or both as needed to generate global, or ULA, or both IP6 addresses. Also note that the prefix are 64 bit long, including any subnet ID.

Dual IPv4-IPv6 Address Range

The dual stack IP4-IP6 is just a combination of the two previous ARs, thus generating both a IPv4 and IPv6 addresses.

Attribute	Description
TYPE	IP4_6
IP	First IP in the range in dot notation.
MAC	Optional. First MAC, if not provided it will be generated using the IP and the MAC_PREFIX in oned.conf.
GLOBAL_PREFIX	Optional. A /64 globally routable prefix
ULA_PREFIX	Optional. A /64 unique local address (ULA) prefix corresponding to the fd00::/8 block
SIZE	Number of addresses in this range.

```
# Example of a dual IP6 IP4 range, template form
AR=[
  TYPE = "IP4_6",
  IP   = "192.1.0.1",
  GLOBAL_PREFIX = "2001:a::",
  SIZE = "60"
]
```

Ethernet Address Range

This is the simplest AR, just MAC addresses are generated for the VM guests. You should use this AR when an external service is providing the IP addresses, such a DHCP server.

Attribute	Description
TYPE	ETHER
MAC	Optional. First MAC, if not provided it will be generated randomly.
SIZE	Number of addresses in this range.

```
# Example of Ethernet range, template form
AR=[
  TYPE = "ETHER",
  SIZE = "25"
]
```

This guide uses the CLI command `onevnet`, but you can also manage your virtual networks using Sunstone. Select the Network tab, and there you will be able to create and manage your virtual networks in a user friendly way.

1.2.2 Adding and Deleting Virtual Networks

A VNET is created through a template definition file containing the previous set of attributes: configuration, context and address ranges. The following example shows how to define a pure IPv4.

Create a file with the network configuration: `priv.net`

```
# Configuration attributes (dummy driver)
NAME = "Private Network"
DESCRIPTION = "A private network for VM inter-communication"

BRIDGE = "bond-br0"

# Context attributes
NETWORK_ADDRESS = "10.0.0.0"
```

```
NETWORK_MASK = "255.255.255.0"
DNS           = "10.0.0.1"
GATEWAY      = "10.0.0.1"

#Address Ranges, only these addresses will be assigned to the VMs
AR=[
  TYPE = "IP4",
  IP   = "10.0.0.10",
  SIZE = "100",
]

AR=[
  TYPE = "IP4",
  IP   = "10.0.0.200",
  SIZE = "10",
]
```

Once the file has been created, we can create the VNET executing:

```
$ onevnet create priv.net
ID: 0
```

You can remove a VNET when no longer needed using its ID or NAME:

```
$ onevnet delete 0
$ onevnet delete "Private Network"
```

Also, onevnet can be used to query OpenNebula about available VNETs:

```
$ onevnet list
```

ID	USER	GROUP	NAME	CLUSTER	BRIDGE	LEASES
0	admin	oneadmin	Private	-	vbr1	0
1	admin	oneadmin	Public	-	vbr0	0

In the output above, USER is the owner of the network and LEASES the number of addresses assigned to a VM or reserved from each VNET.

You can also check the IPs leased in a network with the onevnet show command

```
$ onevnet show 1
VIRTUAL NETWORK 1 INFORMATION
ID           : 1
NAME        : Public
USER        : admin
GROUP       : oneadmin
CLUSTER     : -
BRIDGE      : vbr0
VLAN        : No
USED LEASES : 1

PERMISSIONS
OWNER       : um-
GROUP      : u--
OTHER      : u--

VIRTUAL NETWORK TEMPLATE
BRIDGE="vbr0"
DESCRIPTION="Network with Internet connection through NAT"
NETWORK_ADDRESS="10.0.0.0"
NETWORK_MASK="255.255.255.0"
```

```

PHYDEV=""
VLAN="NO"
VLAN_ID=""

ADDRESS RANGE POOL
AR  TYPE      SIZE LEASES          MAC              IP              GLOBAL_PREFIX
0  IP4         51      1 02:00:0a:00:00:96  10.0.0.150

LEASES
AR  OWNER      MAC              IP              IP6_GLOBAL
0  VM : 43      02:00:0a:00:00:96  10.0.0.150

```

Check the `onevnet` command help or the *reference guide* for more options to list the virtual networks.

VNET Definition Tips

- You may have some used IPs in a VNET so you do not want them to be assigned. You can add as many ARs as you need to implement these address gaps. Alternatively you can put address on hold to prevent them to be assigned.
- ARs can be of SIZE = 1 to define single addresses lease scheme. This set up is equivalent to the previous FIXED VNET type.
- ARs does not need to be of the same type or belong to the same IP network. To accommodate this use case you can overwrite context attributes in the AR, this is you can include attributes like NETWORK_MASK or DNS in the AR definition.
- *Super-netting*, you can even combine ARs overwriting the VNET BRIDGE or with a different VLAN_ID. This way a VNET can be a logical network, e.g. DMZ, that can be implemented through multiple VLANs or host interfaces.
- There are no need to plan all your IP assignment plan beforehand, ARs can be added and modified after the VNET is created, see below.

1.2.3 Updating a Virtual Network

The following attributes can be changed after creating the network, using `onevnet update` command: - Any attribute corresponding to the context or VNET description. - Network configuration attributes, in particular: PHYDEV, VLAN_ID, VLAN and BRIDGE

Also the name of the VNET can be changed with `onevnet rename` command.

1.2.4 Managing Address Ranges

Addresses of a VNET are structured in Address Ranges (AR), VNETs are quite flexible in terms of addition and removal of addresses. In this way, you can easily add new addresses to an existing VNET if the current addresses are exhausted.

Adding and Removing Address Ranges

A new AR can be added to the VNET using exactly the same definition parameters as described above. For example the following command will add a new AR of 20 IP addresses to the VNET:

```
onevnet addar Private --ip 10.0.0.200 --size 20
```

In the same way you can remove an AR:

```
onevnet rmar Private 2
```

Using Sunstone you can manage ARs (add, remove or update) in the Addresses tab of the VNET information.

The screenshot shows the Sunstone interface for a Virtual Network. The 'Addresses' tab is active, displaying a table with the following data:

Address Range	Type	Start	IPv6 Prefix	Leases
0	IP4	192.168.4.1	--	1 / 254

Below the table, there is a large information icon and the text: "Select an Address Range to see more information".

Updating Address Ranges

You can update the following attributes of an AR:

- SIZE, assigned addresses cannot fall outside of the range.
- IPv6 prefix: GLOBAL_PREFIX and ULA_PREFIX
- Any custom attribute that may override the VNET defaults.

The following command shows how to update an AR using the CLI

```
#Update the AR 0 of VNET "Private"  
onevnet updatear Private 0
```

Hold and Release Leases

Addresses can be temporarily be marked on hold state. They are part of the network, but they will not be assigned to any VM.

To do so, use the 'onevnet hold' and 'onevnet release' commands. By default, the address will be put on hold in all ARs containing it; if you need to hold the IP of a specific AR you can specified it with the '-a <AR_ID>' option.

```
#Hold IP 10.0.0.120 in all ARs  
$ onevnet hold "Private Network" 10.0.0.120  
  
#Hold IP 10.0.0.123 in AR 0  
$ onevnet hold 0 10.0.0.123 -a 0
```

You see the list of leases on hold with the ‘onevnet show’ command, they’ll show up as used by VM -1, ‘VM: -1’

1.2.5 Using a VNET

Getting an address for a VM

An address lease from a VNET can be obtained by simply specifying the virtual network name in the NIC attribute.

For example, to define VM with a network interfaces connected to the `Private Network` just include in the template:

```
# Reference by ID
NIC = [ NETWORK_ID = 0 ]
# Reference by NAME
NIC = [ NETWORK      = "Private Network" ]
```

Networks can be referred in a NIC in two different ways, see the *Simplified Virtual Machine Definition File documentation* for more information:

- `NETWORK_ID`, using its ID as returned by the create operation
- `NETWORK`, using its name. In this case the name refers to one of the virtual networks owned by the user (names cannot be repeated for the same user). If you want to refer to a VNET of other user you can specify that with `NETWORK_UID` (by the uid of the user) or `NETWORK_UNAME` (by the name of the user).

You can also request a specific address just by adding the `IP` or `MAC` attributes to `NIC`. If no address is requested the first free address (in any AR) will be used.

```
NIC = [ NETWORK_ID = 1, IP = 192.168.0.3 ]
```

When the VM is submitted, OpenNebula will look for available IPs, leases on hold, reserved or in use by other VMs will be skipped. If successful, the `onevm show` command should return information about the machine, including network information.

```
$ onevm show 0
VIRTUAL MACHINE 0 INFORMATION
VMID=0
...
VM NICS
  ID NETWORK          VLAN BRIDGE          IP          MAC
  0 Public            no vbr0             10.0.0.150  02:00:0a:00:00:96
```

Warning: Note that if OpenNebula is not able to obtain a lease from a network the submission will fail.

Using the address within the VM

Hypervisors can set the MAC address of a virtual NIC, but VMs need to obtain an IP address for it. In order to configure the IP inside the guest, you need to use one of the two available methods:

- Instantiate a Virtual Router inside each Virtual Network. The Virtual Router appliance contains a DHCP server that knows the IP assigned to each VM.
- Contextualize the VM. Please visit the *contextualization guide* to learn how to configure your Virtual Machines to automatically obtain an IP derived from the MAC.

- Use an custom external service (e.g. your own DHCP server)

Apply Firewall Rules to VMs

You can apply firewall rules on your VMs, to filter TCP and UDP ports, and to define a policy for ICMP connections.

Read more about this feature here.

1.2.6 VNET Self-Provisioning: Reservations

VNETs implement a simple self-provisioning scheme, that allows users to create their own networks consisting of portions of an existing VNET. Each portion is called a Reservation. To implement this you need to:

- **Define a VNET**, with the desired ARs and configuration attributes. These attributes will be inherited by any Reservation made on the VNET. Final users does not need to deal with low-level networking details.
- **Setting up access**. In order to make a Reservation, users needs USE rights on the VNET, just as if they would use it to directly to provision IPs from it.
- **Make Reservations**. Users can easily request specific addresses or just a number of addresses from a VNET. Reservations are placed in their own VNET for the user.
- **Use Reservations**. Reservations offer the same interface as a regular VNET so you can just point your VM templates to the new VNET. The number of addresses and usage stats are shown also in the same way.

Setting up access to VNETs

Once a VNET is setup by a Cloud admin, she needs to make it available to other users in the cloud. See the Managing Permissions documentation for more information.

Let's see a quick example. The following command allows users in the same group **USE** and **MANAGE** the virtual network, and the rest of the users **USE** it:

```
$ onevnet chmod 0 664
$ onevnet show 0
...
PERMISSIONS
OWNER      : um-
GROUP      : um-
OTHER      : u--
```

Note: Users can only attach VMs or make reservations from VNETs with **USE** rights on the VNET

Make and delete Reservations

In its simplest form you can make a reservations just by defining the source VNET, the number of addresses and the name of the reservation. For example to reserve 10 addresses from VNET Private and place it on MyVNET just:

```
$ onevnet reserve Private -n MyVNET -s 10
```

As a result a new VNET has been created:


```
$ onevnet list
ID USER      GROUP      NAME      CLUSTER  BRIDGE  LEASES
0 admin      oneadmin   Private   -        vbr1    10
1 helen      users      MyVNET    -        vbr1    0
```

Note that VNET Private shows 10 address leases in use, and leased to VNET 1. Also note that both VNETs share the same configuration, e.g. BRIDGE vbr1. You can verify this details with `onevnet show` command.

Reservations can include advanced options such as:

- The AR where you want to make the reservation from in the source VNET
- The starting IP or MAC to make the reservation from

A reservation can be remove just as a regular VNET:

```
$ onevnet delete MyVNET
```

Using Reservations

To use a reservation you can use it as any other VNET; as they expose the same interface, i.e. you can refer to VNET variables in context, add NICs...

```
#Use a reservation in a VM
NIC = [ NETWORK = "MyVNET" ]
```

A Reservation can be also extended with new addresses. This is, you can add a new reservation to an existing one. This way a user can refer to its own network with a controlled and deterministic address space.

Note: Reservation increase leases counters on the user and group, and they can be limited through a quota.

Note: The reservation interface is exposed by Sunstone in a very convenient way.

1.3 Managing Images

The Storage system allows OpenNebula administrators and users to set up images, which can be operative systems or data, to be used in Virtual Machines easily. These images can be used by several Virtual Machines simultaneously, and also shared with other users.

If you want to customize the Storage in your system, visit the Storage subsystem guide.

1.3.1 Image Types

There are six different types of images. Using the command `oneimage chtype`, you can change the type of an existing Image.

- **OS:** An OS image contains a working operative system. Every *VM template* must define one DISK referring to an image of this type.
- **CDROM:** These images are readonly data. Only one image of this type can be used in each *VM template*. These type of images are not cloned when using shared storage.
- **DATABLOCK:** A datablock image is a storage for data, which can be accessed and modified from different Virtual Machines. These images can be created from previous existing data, or as an empty drive.

- **KERNEL:** A plain file to be used as kernel (VM attribute OS/KERNEL_DS). Note that KERNEL file images can be registered only in File Datastores.
- **RAMDISK:** A plain file to be used as ramdisk (VM attribute OS/INITRD_DS). Note that RAMDISK file images can be registered only in File Datastores.
- **CONTEXT:** A plain file to be included in the context CD-ROM (VM attribute CONTEXT/FILES_DS). Note that CONTEXT file images can be registered only in File Datastores.


The Virtual Machines can use as many datablocks as needed. Refer to the *VM template* documentation for further information.

Warning: Note that some of the operations described below do not apply to KERNEL, RAMDISK and CONTEXT images, in particular: clone and persistent.


1.3.2 Image Life-cycle

Short state	State	Meaning
lock	LOCKED	The image file is being copied or created in the Datastore.
rdy	READY	Image ready to be used.
used	USED	Non-persistent Image used by at least one VM. It can still be used by other VMs.
used	USED_PERS	Persistent Image is use by a VM. It cannot be used by new VMs.
disa	DISABLED	Image disabled by the owner, it cannot be used by new VMs.
err	ERROR	Error state, a FS operation failed. See the Image information with <code>oneimage show</code> for an error message.
dele	DELETE	The image is being deleted from the Datastore.

This is the state diagram for **persistent** images:



And the following one is the state diagram for **non-persistent** images:



1.3.3 Managing Images

Users can manage their images using the command line interface command `oneimage`. The complete reference is [here](#).

You can also manage your images using Sunstone. Select the Images tab, and there you will be able to create, enable, disable, delete your images and even manage their persistence and publicity in a user friendly way. From Sunstone 3.4, you can also upload images directly from the web UI.

Create Image

Wizard | Advanced mode

Name: ⓘ Type: ⓘ

Description: ⓘ Datastore: ⓘ

Persistent: ⓘ

Image location:

Provide a path Upload Empty datablock

No file chosen

▼ Advanced options

Create Images

Note: For VMWare images, please read **also** the VMware Drivers guide.

The three types of images can be created from an existing file, but for **datablock** images you can specify a size and filesystem type and let OpenNebula create an empty image in the datastore.

If you want to create an **OS image**, you need to prepare a contextualized virtual machine, and extract its disk.

Please read first the documentation about the MAC to IP mechanism in the [virtual network management documentation](#), and how to use contextualization [here](#).

Once you have a disk you want to upload, you need to create a new [image template](#), and submit it using the `oneimage create` command.

The complete reference for the image template is [here](#). This is how a sample template looks like:

```
$ cat ubuntu_img.one
NAME           = "Ubuntu"
PATH           = /home/cloud/images/ubuntu-desktop/disk.0
```

```
TYPE          = OS
DESCRIPTION   = "Ubuntu 10.04 desktop for students."
```

You need to choose the Datastore where to register the new Image. To know the available datastores, use the `onedatastore list` command. In this case, only the ‘default’ one is listed:

```
$ onedatastore list
ID NAME          CLUSTER IMAGES TYPE  TM
1 default        -        1      fs   shared
```

To submit the template, you just have to issue the command

```
$ oneimage create ubuntu_img.one --datastore default
ID: 0
```

You can also create images using just parameters in the `oneimage create` call. The parameters to generate the image are as follows:

Parameter	Description
<code>--name name</code>	Name of the new image
<code>--description description</code>	Description for the new Image
<code>--type type</code>	Type of the new Image: OS, CDROM or DATABLOCK, FILE
<code>--persistent</code>	Tells if the image will be persistent
<code>--prefix prefix</code>	Device prefix for the disk (eg. hd, sd, xvd or vd)
<code>--target target</code>	Device the disk will be attached to
<code>--path path</code>	Path of the image file
<code>--driver driver</code>	Driver to use image (raw, qcow2, tap:aio:...)
<code>--disk_type disk_type</code>	Type of the image (BLOCK, CDROM or FILE)
<code>--source source</code>	Source to be used. Useful for not file-based images
<code>--size size</code>	Size in MB. Used for DATABLOCK type
<code>--fstype fstype</code>	Type of file system to be built: ext2, ext3, ext4, ntfs, reiserfs, jfs, swap, qcow2

To create the previous example image you can do it like this:

```
$ oneimage create --datastore default --name Ubuntu --path /home/cloud/images/ubuntu-desktop/disk.0 \
--description "Ubuntu 10.04 desktop for students."
```

Note: You can use **gz** compressed image files (i.e. as specified in path) when registering them in OpenNebula.

Uploading Images from Sunstone

Image file upload to the server via the client browser is possible with the help of a vendor library. The process is as follow:

- Step 1: The client uploads the whole image to the server in a temporal file in the `tmpdir` folder specified in the configuration.
- Step 2: OpenNebula registers an image setting the `PATH` to that temporal file.
- Step 3: OpenNebula copies the images to the datastore.
- Step 4: The temporal file is deleted and the request returns successfully to the user (a message pops up indicating that image was uploaded correctly).

Note that when file sizes become big (normally over 1GB), and depending on your hardware, it may take long to complete the copying in step 3. Since the upload request needs to stay pending until copying is successful (so it can delete the temp file safely), there might be Ajax timeouts and/or lack of response from the server. This may cause errors, or trigger re-uploads (which reinitiate the loading progress bar).

As of Firefox 11 and previous versions, uploads seem to be limited to 2GB. Chrome seems to work well with images > 4 GB.

Clone Images

Existing images can be cloned to a new one. This is useful to make a backup of an Image before you modify it, or to get a private persistent copy of an image shared by other user. Note that persistent images with snapshots cannot be cloned. In order to do so, the user would need to flatten it first, see the *snapshots* section for more information.

To clone an image, execute

```
$ oneimage clone Ubuntu new_image
```

You can optionally clone the Image to a different Datastore. The new Datastore must be compatible with the current one, i.e. have the same DS_MAD drivers.

```
$ oneimage clone Ubuntu new_image --datastore new_img_ds
```

Listing Available Images

You can use the `oneimage list` command to check the available images in the repository.

```
$ oneimage list
  ID USER      GROUP   NAME      DATASTORE   SIZE TYPE PER STAT  RVMS
   0 oneuser1  users   Ubuntu    default      8M  OS  No  rdy    0
```

To get complete information about an image, use `oneimage show`, or list images continuously with `oneimage top`.

Publishing Images

The users can share their images with other users in their group, or with all the users in OpenNebula. See the *Managing Permissions* documentation for more information.

Let's see a quick example. To share the image 0 with users in the group, the **USE** right bit for **GROUP** must be set with the `chmod` command:

```
$ oneimage show 0
...
PERMISSIONS
OWNER       : um-
GROUP      : ---
OTHER      : ---

$ oneimage chmod 0 640

$ oneimage show 0
...
PERMISSIONS
OWNER       : um-
```

```
GROUP      : u--
OTHER     : ---
```

The following command allows users in the same group **USE** and **MANAGE** the image, and the rest of the users **USE** it:

```
$ oneimage chmod 0 664

$ oneimage show 0
...
PERMISSIONS
OWNER       : um-
GROUP      : um-
OTHER      : u--
```

The commands `oneimage publish` and `oneimage unpublish` are still present for compatibility with previous versions. These commands set/unset the **GROUP USE** bit.

Making Images Persistent

Use the `oneimage persistent` and `oneimage nonpersistent` commands to make your images persistent or not.

A persistent image saves back to the datastore the changes made inside the VM after it is shut down. More specifically, the changes are correctly preserved only if the VM is ended with the `onevm shutdown` or `onevm shutdown --hard` commands. Note that depending on the Datastore type a persistent image can be a link to the original image, so any modification is directly made on the image.

```
$ oneimage list
ID USER      GROUP      NAME          DATASTORE    SIZE TYPE PER  STAT  RVMS
  0 oneadmin oneadmin Ubuntu      default      10G  OS  No  rdy   0
$ oneimage persistent Ubuntu
$ oneimage list
ID USER      GROUP      NAME          DATASTORE    SIZE TYPE PER  STAT  RVMS
  0 oneadmin oneadmin Ubuntu      default      10G  OS  Yes rdy   0
$ oneimage nonpersistent 0
$ oneimage list
ID USER      GROUP      NAME          DATASTORE    SIZE TYPE PER  STAT  RVMS
  0 oneadmin oneadmin Ubuntu      default      10G  OS  No  rdy   0
```

Warning: When images are public (GROUP or OTHER USE bit set) they are always cloned, and persistent images are never cloned. Therefore, an image cannot be public and persistent at the same time. To manage a public image that won't be cloned, unpublish it first and make it persistent.

Note that persistent images with snapshots cannot be made nonpersistent. In order to do so, the user would need to flatten it first, see the *snapshots* section for more information.

Managing Snapshots in Persistent Images

Persistent images can have associated snapshots if the user *created them* during the life-cycle of VM that used the persistent image. The following are operations that allow the user to manage these snapshots directly:

- `oneimage snapshot-revert <image_id> <snapshot_id>`: The active state of the image is overwritten by the specified snapshot. Note that this operation discards any unsaved data of the disk state.

- `oneimage snapshot-delete <image_id> <snapshot_id>`: Deletes a snapshot. This operation is only allowed if the snapshot is not the active snapshot and if it has no children.
- `oneimage snapshot-flatten <image_id> <snapshot_id>`: This operation effectively converts the image to an image without snapshots. The saved disk state of the image is the state of the specified snapshot. It's an operation similar to running `snapshot-revert` and then deleting all the snapshots.

Images with snapshots **cannot** be cloned or made non-persistent. To run either of these operations the user would need to flatten the image first.

1.3.4 How to Use Images in Virtual Machines

This is a simple example on how to specify images as virtual machine disks. Please visit the [virtual machine user guide](#) and the [virtual machine template](#) documentation for a more thorough explanation.

Assuming you have an OS image called *Ubuntu desktop* with ID 1, you can use it in your *virtual machine template* as a DISK. When this machine is deployed, the first disk will be taken from the image repository.

Images can be referred to in a DISK in two different ways:

- `IMAGE_ID`, using its ID as returned by the create operation
- `IMAGE`, using its name. In this case the name refers to one of the images owned by the user (names can not be repeated for the same user). If you want to refer to an `IMAGE` of another user you can specify that with `IMAGE_UID` (by the uid of the user) or `IMAGE_UNAME` (by the name of the user).

```
CPU      = 1
MEMORY  = 3.08

DISK = [ IMAGE_ID    = 1 ]

DISK = [ type       = swap,
        size       = 1024 ]

NIC     = [ NETWORK_ID = 1 ]
NIC     = [ NETWORK_ID = 0 ]

# FEATURES=[ acpi="no" ]

GRAPHICS = [
  type     = "vnc",
  listen   = "1.2.3.4",
  port     = "5902" ]

CONTEXT = [
  files    = "/home/cloud/images/ubuntu-desktop/init.sh" ]
```

Save Changes

Once the VM is deployed you can snapshot a disk, i.e. save the changes made to the disk as a new image. There are two types of disk snapshots in OpenNebula:

- **Deferred snapshots** (disk-snapshot), changes to a disk will be saved as a new Image in the associated datastore when the VM is shutdown.

- **Hot snapshots** (hot disk-snapshot), just as the deferred snapshots, but the disk is copied to the datastore the moment the operation is triggered. Therefore, you must guarantee that the disk is in a consistent state during the `save_as` operation (e.g. by unmounting the disk from the VM).

To save a disk, use the `onevm disk-snapshot` command. This command takes three arguments: The VM name (or ID), the disk ID to save and the name of the new image to register. And optionally the `-live` argument to not defer the disk-snapshot operation. A deferred snapshot can be canceled with the `onevm disk-snapshot-cancel` command.

To know the ID of the disk you want to save, just take a look at the `onevm show` output for your VM, you are interested in the ID column in the VM DISK section.

```
$ onevm show 11
VIRTUAL MACHINE 11 INFORMATION
ID                : 11
NAME              : ttylinux-11
USER              : ruben
GROUP             : oneadmin
STATE             : PENDING
LCM_STATE         : LCM_INIT
RESCHED           : No
START TIME        : 03/08 22:24:57
END TIME          : -
DEPLOY ID         : -

[..]

VM DISKS
ID TARGET IMAGE          TYPE SAVE SAVE_AS
 0   hda ttylinux        file  NO     -
 1   hdb raw - 100M      fs    NO     -

VM NICS
ID NETWORK      VLAN BRIDGE   IP           MAC
 0 net_172      no  vbr0    172.16.0.201 02:00:ac:10:00:c9
                                     fe80::400:acff:fe10:c9
```

The IDs are assigned in the same order the disks were defined in the *VM template*.

The next command will register a new image called “SO upgraded”, that will be ready as soon as the VM is shut down. Till then the image will be locked, and so you cannot use it.

```
$ onevm disk-snapshot ttylinux-11 0 "SO upgraded"
```

This command copies disk 1 to the datastore with name *Backup of DB volume*, the image will be available once the image copy end:

```
$ onevm disk-snapshot --live ttylinux-11 1 "Backup of DB volume"
```

1.3.5 How to Use File Images in Virtual Machines

KERNEL and RAMDISK

KERNEL and RAMDISK type Images can be used in the `OS/KERNEL_DS` and `OS/INITRD_DS` attributes of the VM template. See the *complete reference* for more information.

Example:

```
OS = [ KERNEL_DS   = "$FILE[IMAGE=kernel3.6]",
        INITRD_DS  = "$FILE[IMAGE_ID=23]",
        ROOT       = "sda1",
        KERNEL_CMD = "ro xencons=tty console=tty1" ]
```

CONTEXT

The *contextualization cdrom* can include CONTEXT type Images. Visit the *complete reference* for more information.

```
CONTEXT = [
  FILES_DS = "$FILE[IMAGE_ID=34] $FILE[IMAGE=kernel]",
]
```

1.4 Creating Virtual Machines

In OpenNebula the Virtual Machines are defined with Template files. This guide explains **how to describe the wanted-to-be-ran Virtual Machine, and how users typically interact with the system.**

The Template Repository system allows OpenNebula administrators and users to register Virtual Machine definitions in the system, to be instantiated later as Virtual Machine instances. These Templates can be instantiated several times, and also shared with other users.

1.4.1 Virtual Machine Model

A Virtual Machine within the OpenNebula system consists of:

- A capacity in terms memory and CPU
- A set of NICs attached to one or more virtual networks
- A set of disk images
- A state file (optional) or recovery file, that contains the memory image of a running VM plus some hypervisor specific information.

The above items, plus some additional VM attributes like the OS kernel and context information to be used inside the VM, are specified in a template file.

1.4.2 Defining a VM in 3 Steps

Virtual Machines are defined in an OpenNebula Template. Templates are stored in a repository to easily browse and instantiate VMs from them. To create a new Template you have to define 3 things

- **Capacity & Name**, how big will the VM be?

Attribute	Description	Mandatory	Default
NAME	Name that the VM will get for description purposes.	Yes	one-<vmid>
MEMORY	Amount of RAM required for the VM, in Megabytes.	Yes	
CPU	CPU ratio (e..g half a physical CPU is 0.5).	Yes	
VCPU	Number of virtual cpus.	No	1

- **Disks**. Each disk is defined with a DISK attribute. A VM can use three types of disk:
 - **Use a persistent Image** changes to the disk image will persist after the VM is shutdown.

- Use a **non-persistent Image** images are cloned, changes to the image will be lost.
- **Volatile** disks are created on the fly on the target host. After the VM is shutdown the disk is disposed.

• **Persistent and Clone Disks**

Attribute	Description	Mandatory	Default
IMAGE_ID and IMAGE	The ID or Name of the image in the datastore	Yes	
IMAGE_UID	Select the IMAGE of a given user by her ID	No	self
IMAGE_UNAME	Select the IMAGE of a given user by her NAME	No	self

• **Volatile**

Attribute	Description	Mandatory	Default
TYPE	Type of the disk: <code>swap</code> , <code>fs</code> . <code>swap</code> type will set the label to <code>swap</code> so it is easier to mount and the context packages will automatically mount it.	Yes	
SIZE	size in MB	Yes	
FORMAT	filesystem for fs images: <code>ext2</code> , <code>ext3</code> , etc. <code>raw</code> will not format the image. For VMs to run on <code>vmfs</code> or <code>vmware shared</code> configurations, the valid values are: <code>vmdk_thin</code> , <code>vmdk_zeroedthick</code> , <code>vmdk_eagerzeroedthick</code>	Yes	

• **Network Interfaces.** Each network interface of a VM is defined with the `NIC` attribute.

Attribute	Description	Mandatory	Default
NETWORK_ID and NETWORK	The ID or Name of the network	Yes	
NETWORK_UID	Select the NETWORK of a given user by her ID	No	self
NETWORK_UNAME	Select the NETWORK of a given user by her NAME	No	self

The following example shows a VM Template file with a couple of disks and a network interface, also a VNC section was added.

```

NAME = test-vm
MEMORY = 128
CPU = 1

DISK = [ IMAGE = "Arch Linux" ]
DISK = [ TYPE = swap,
        SIZE = 1024 ]

NIC = [ NETWORK = "Public", NETWORK_UNAME="oneadmin" ]

GRAPHICS = [
  TYPE = "vnc",
  LISTEN = "0.0.0.0" ]
    
```

Simple templates can be also created using the command line instead of creating a template file. The parameters to do this for `oneteplate` are:

Parameter	Description
<code>--name name</code>	Name for the VM
<code>--cpu cpu</code>	CPU percentage reserved for the VM (1=100% one CPU)
<code>--vcpu vcpu</code>	Number of virtualized CPUs
<code>--arch arch</code>	Architecture of the VM, e.g.: i386 or x86_64
<code>--memory memory</code>	Memory amount given to the VM
<code>--disk disk0,disk1</code>	Disks to attach. To use a disk owned by other user use <code>user[disk]</code>
<code>--nic vnet0,vnet1</code>	Networks to attach. To use a network owned by other user use <code>user[network]</code>
<code>--raw string</code>	Raw string to add to the template. Not to be confused with the RAW attribute. If you want to provide more than one element, just include an enter inside quotes, instead of using more than one <code>-raw</code> option
<code>--vnc</code>	Add VNC server to the VM
<code>--ssh [file]</code>	Add an ssh public key to the context. If the file is omitted then the user variable <code>SSH_PUBLIC_KEY</code> will be used.
<code>--net_context</code>	Add network contextualization parameters
<code>--context line1,line2*</code>	Lines to add to the context section
<code>--boot device</code>	Select boot device (hd, fd, cdrom or network)

A similar template as the previous example can be created with the following command:

```
$ onetemplate create --name test-vm --memory 128 --cpu 1 --disk "Arch Linux" --nic Public
```

Warning: You may want to add VNC access, input hw or change the default targets of the disks. Check the *VM definition file for a complete reference*

Warning: OpenNebula Templates are designed to be hypervisor-agnostic, but there are additional attributes that are supported for each hypervisor. Check the Xen, KVM and VMware configuration guides for more details

Warning: Volatile disks can not be saved as. Pre-register a DataBlock image if you need to attach arbitrary volumes to the VM

1.4.3 Managing Templates

Users can manage the Template Repository using the command `onetemplate`, or the graphical interface Sunstone. For each user, the actual list of templates available are determined by the ownership and permissions of the templates.

Listing Available Templates

You can use the `onetemplate list` command to check the available Templates in the system.

```
$ onetemplate list a
ID USER      GROUP      NAME                REGTIME
  0 oneadmin  oneadmin  template-0          09/27 09:37:00
  1 oneuser   users     template-1          09/27 09:37:19
  2 oneadmin  oneadmin  Ubuntu_server       09/27 09:37:42
```

To get complete information about a Template, use `onemplate show`.

Here is a view of templates tab in Sunstone:

ID	Owner	Group	Name	Registration time
6	John	BlueVDC	My Saved Template-4	15:39:05 05/08/2014
5	BlueVDC-admin	BlueVDC	Dev Environment-3	11:41:49 05/08/2014
4	BlueVDC-admin	BlueVDC	CentOS with Apache-2	11:41:30 05/08/2014
3	BlueVDC-admin	BlueVDC	My saved template-1	11:40:53 05/08/2014
2	oneadmin	BlueVDC	Fedora 20	16:57:52 04/08/2014
1	oneadmin	BlueVDC	Ubuntu 14.04	16:55:29 04/08/2014
0	oneadmin	BlueVDC	CentOS 6.6	16:54:32 04/08/2014

Adding and Deleting Templates

Using `onemplate create`, users can create new Templates for private or shared use. The `onemplate delete` command allows the Template owner -or the OpenNebula administrator- to delete it from the repository.

For instance, if the previous example template is written in the `vm-example.txt` file:

```
$ onemplate create vm-example.txt
ID: 6
```

You can also clone an existing Template, with the `onemplate clone` command:


```
$ onemplate clone 6 new_template
ID: 7
```


Via Sunstone, you can easily add templates using the provided wizards (or copy/pasting a template file) and delete them clicking on the delete button:


✕


Create VM Template


Wizard
Advanced mode



General



Storage



Network


OS Booting


Input/Output


Context


Scheduling



Other

Name ?

Logo ?

CentOS
▼

Description ?



Memory ?

▢

MB
▼

CPU ?

▢

Advanced options

Reset

Create

Updating a Template

It is possible to update a template by using the `oneteplate update`. This will launch the editor defined in the variable `EDITOR` and let you edit the template.

```
$ oneteplate update 3
```

Publishing Templates

The users can share their Templates with other users in their group, or with all the users in OpenNebula. See the [Managing Permissions](#) documentation for more information.

Let's see a quick example. To share the Template 0 with users in the group, the **USE** right bit for **GROUP** must be set with the **chmod** command:

```
$ oneteplate show 0
...
PERMISSIONS
OWNER       : um-
GROUP      : ---
OTHER      : ---

$ oneteplate chmod 0 640
```

```
$ onetemplate show 0
...
PERMISSIONS
OWNER      : um-
GROUP      : u--
OTHER      : ---
```

The following command allows users in the same group **USE** and **MANAGE** the Template, and the rest of the users **USE** it:

```
$ onetemplate chmod 0 664

$ onetemplate show 0
...
PERMISSIONS
OWNER      : um-
GROUP      : um-
OTHER      : u--
```

The commands `onetemplate publish` and `onetemplate unpublish` are still present for compatibility with previous versions. These commands `set/unset` the `GROUP USE` bit.

1.4.4 Instantiating Templates

The `onetemplate instantiate` command accepts a Template ID or name, and creates a VM instance (you can define the number of instances using the `-multiple num_of_instances` option) from the given template.

```
$ onetemplate instantiate 6
VM ID: 0

$ onevm list
  ID USER      GROUP      NAME          STAT CPU    MEM      HOSTNAME      TIME
  0 oneuser1  users      one-0         pend  0      0K           00 00:00:16
```

You can also merge another template to the one being instantiated. The new attributes will be added, or will replace the ones from the source template. This can be more convenient than cloning an existing template and updating it.

```
$ cat /tmp/file
MEMORY = 512
COMMENT = "This is a bigger instance"

$ onetemplate instantiate 6 /tmp/file
VM ID: 1
```

The same options to create new templates can be used to be merged with an existing one. See the above table, or execute `'onetemplate instantiate -help'` for a complete reference.

```
$ onetemplate instantiate 6 --cpu 2 --memory 1024
VM ID: 2
```

Ask for User Inputs

The User Inputs functionality provides the template creator with the possibility to dynamically ask the user instantiating the template for dynamic values that must be defined.

Update VM Template

Wizard

Advanced mode

General	Storage	Network	OS Booting	Input/Output	Context	Scheduling	Other
---------	---------	---------	------------	--------------	---------	------------	-------

Network & SSH

Files

User Inputs

Custom vars

User Inputs

Name	Type	Description
BLOG_TITLE	text	Blog title
MYSQL_PASSWORD	password	MySQL Password

+ Add another attribute

These inputs will be presented to the user when the Template is instantiated. The VM guest needs to be *contextualized* to make use of the values provided by the user.

Custom Attributes

* Blog title

* MySQL Password

Create

If a VM Template with user inputs is used by a Service Template Role, the user will be also asked for these inputs when the Service is created.

Merge Use Case

The template merge functionality, combined with the restricted attributes, can be used to allow users some degree of customization for predefined templates.

Let's say the administrator wants to provide base templates that the users can customize, but with some restrictions. Having the following restricted attributes in `oned.conf`:

```
VM_RESTRICTED_ATTR = "CPU"
VM_RESTRICTED_ATTR = "VPU"
VM_RESTRICTED_ATTR = "NIC"
```

And the following template:

```
CPU      = "1"
VCPU    = "1"
MEMORY  = "512"
DISK=[
  IMAGE_ID = "0" ]
NIC=[
  NETWORK_ID = "0" ]
```

Users can instantiate it customizing anything except the CPU, VCPU and NIC. To create a VM with different memory and disks:

```
$ onetemplate instantiate 0 --memory 1G --disk "Ubuntu 12.10"
```

Warning: The merged attributes replace the existing ones. To add a new disk, the current one needs to be added also.

```
$ onetemplate instantiate 0 --disk 0,"Ubuntu 12.10"
```

1.4.5 Deployment

The OpenNebula Scheduler will deploy automatically the VMs in one of the available Hosts, if they meet the requirements. The deployment can be forced by an administrator using the `onevm deploy` command.

Use `onevm shutdown` to shutdown a running VM.

Continue to the *Managing Virtual Machine Instances Guide* to learn more about the VM Life Cycle, and the available operations that can be performed.

1.5 Managing Virtual Machines

This guide follows the *Creating Virtual Machines guide*. Once a Template is instantiated to a Virtual Machine, there are a number of operations that can be performed using the `onevm` command.

1.5.1 Virtual Machine Life-cycle

The life-cycle of a Virtual Machine within OpenNebula includes the following stages:

Warning: Note that this is a simplified version. If you are a developer you may want to take a look at the complete diagram referenced in the `xml-rpc api` page):

Short state	State	Meaning
pend	Pending	By default a VM starts in the pending state, waiting for a resource to run on. It will stay in this state until the scheduler decides to deploy it, or the user deploys it using the <code>onevm deploy</code> command.
hold	Hold	The owner has held the VM and it will not be scheduled until it is released. It can be, however, deployed manually.
prol	Prolog	The system is transferring the VM files (disk images and the recovery file) to the host in which the virtual machine will be running.
boot	Boot	OpenNebula is waiting for the hypervisor to create the VM.
runn	Running	The VM is running (note that this stage includes the internal virtualized machine booting and shutting down phases). In this state, the virtualization driver will periodically monitor it.
migr	Migrate	The VM is migrating from one resource to another. This can be a live migration or cold migration (the VM is saved and VM files are transferred to the new resource).
hotp	Hotplug	A disk attach/detach, nic attach/detach operation is in process.
snap	Snapshot	A system snapshot is being taken.
save	Save	The system is saving the VM files after a migration, stop or suspend operation.
epil	Epilog	In this phase the system cleans up the Host used to virtualize the VM, and additionally disk images to be saved are copied back to the system datastore.
shut	Shutdown	OpenNebula has sent the VM the shutdown ACPI signal, and is waiting for it to complete the shutdown process. If after a timeout period the VM does not disappear, OpenNebula will assume that the guest OS ignored the ACPI signal and the VM state will be changed to running , instead of done .
stop	Stopped	The VM is stopped. VM state has been saved and it has been transferred back along with the disk images to the system datastore.
susp	Suspended	Same as stopped, but the files are left in the host to later resume the VM there (i.e. there is no need to re-schedule the VM).
poff	PowerOff	Same as suspended, but no checkpoint file is generated. Note that the files are left in the host to later boot the VM there. When the VM guest is shutdown, OpenNebula will put the VM in this state.
unde	Undeploy	The VM is shut down. The VM disks are transferred to the system datastore. The VM can be resumed later.
fail	Failed	The VM failed.
unkn	Unknown	The VM couldn't be reached, it is in an unknown state.
done	Done	The VM is done. VMs in this state won't be shown with <code>onevm list</code> but are kept in the database for accounting purposes. You can still get their information with the <code>onevm show</code> command.

Warning: Read the complete reference for *Virtual Machine templates*.

Assuming we have a VM Template registered called *vm-example* with ID 6, then we can instantiate the VM issuing a:

```
$ onetemplate list
ID USER      GROUP      NAME                REGTIME
 6 oneadmin  oneadmin  vm_example          09/28 06:44:07

$ onetemplate instantiate vm-example --name my_vm
VM ID: 0
```

If the template has *USER INPUTS* defined the CLI will prompt the user for these values:

```
$ onetemplate instantiate vm-example --name my_vm
There are some parameters that require user input.
* (BLOG_TITLE) Blog Title: <my_title>
* (DB_PASSWORD) Database Password:
VM ID: 0
```

Afterwards, the VM can be listed with the `onevm list` command. You can also use the `onevm top` command to list VMs continuously.

```
$ onevm list
ID USER      GROUP      NAME                STAT CPU    MEM      HOSTNAME           TIME
 0 oneadmin  oneadmin  my_vm               pend  0        0K              00 00:00:03
```

After a Scheduling cycle, the VM will be automatically deployed. But the deployment can also be forced by `oneadmin` using `onevm deploy`:

```
$ onehost list
ID NAME                RVM  TCPU  FCPU  ACPU  TMEM  FMEM  AMEM  STAT
 2 testbed              0    800   800   800   16G   16G   16G   on

$ onevm deploy 0 2

$ onevm list
ID USER      GROUP      NAME                STAT CPU    MEM      HOSTNAME           TIME
 0 oneadmin  oneadmin  my_vm               runn  0        0K              testbed 00 00:02:40
```

and details about it can be obtained with `show`:

```
$ onevm show 0
VIRTUAL MACHINE 0 INFORMATION
ID                : 0
NAME              : my_vm
USER              : oneadmin
GROUP             : oneadmin
STATE             : ACTIVE
LCM_STATE         : RUNNING
START TIME        : 04/14 09:00:24
END TIME          : -
DEPLOY ID:        : one-0

PERMISSIONS
OWNER             : um-
GROUP             : ---
OTHER             : ---

VIRTUAL MACHINE MONITORING
```

```
NET_TX           : 13.05
NET_RX           : 0
USED MEMORY      : 512
USED CPU         : 0

VIRTUAL MACHINE TEMPLATE
...

VIRTUAL MACHINE HISTORY
SEQ      HOSTNAME REASON      START      TIME      PTIME
0        testbed  none    09/28 06:48:18 00 00:07:23 00 00:00:00
```

Terminating VM Instances...

You can terminate a running instance with the following operations (either as `onevm` commands or through Sunstone):

- `shutdown`: Gracefully shuts down a running VM, sending the ACPI signal. Once the VM is shutdown the host is cleaned, and persistent and deferred-snapshot disk will be moved to the associated datastore. If after a given time the VM is still running (e.g. guest ignoring ACPI signals), OpenNebula will returned the VM to the `RUNNING` state.
- `shutdown --hard`: Same as above but the VM is immediately destroyed. Use this action instead of `shutdown` when the VM doesn't have ACPI support.

If you need to terminate an instance in any state use:

- `delete`: The VM is immediately destroyed no matter its state. Hosts are cleaned as needed but no images are moved to the repository, leaving then in error. Think of `delete` as `kill -9` for a process, and so it should be only used when the VM is not responding to other actions.

All the above operations free the resources used by the VM

Pausing VM Instances...

There are two different ways to temporarily stop the execution of a VM: short and long term pauses. A **short term** pause keeps all the VM resources allocated to the hosts so its resume its operation in the same hosts quickly. Use the following `onevm` commands or Sunstone actions:

- `suspend`: the VM state is saved in the running Host. When a suspended VM is resumed, it is immediately deployed in the same Host by restoring its saved state.
- `poweroff`: Gracefully powers off a running VM by sending the ACPI signal. It is similar to `suspend` but without saving the VM state. When the VM is resumed it will boot immediately in the same Host.
- `poweroff --hard`: Same as above but the VM is immediately powered off. Use this action when the VM doesn't have ACPI support.

Note: When the guest is shutdown from within the VM, OpenNebula will put the VM in the `poweroff` state.

You can also plan a **long term pause**. The Host resources used by the VM are freed and the Host is cleaned. Any needed disk is saved in the system datastore. The following actions are useful if you want to preserve network and storage allocations (e.g. IPs, persistent disk images):

- `undeploy`: Gracefully shuts down a running VM, sending the ACPI signal. The Virtual Machine disks are transferred back to the system datastore. When an undeployed VM is resumed, it is be moved to the pending state, and the scheduler will choose where to re-deploy it.
- `undeploy --hard`: Same as above but the running VM is immediately destroyed.

- `stop`: Same as `undeploy` but also the VM state is saved to later resume it.

When the VM is successfully paused you can resume its execution with:

- `resume`: Resumes the execution of VMs in the stopped, suspended, undeployed and poweroff states.

Resetting VM Instances...

There are two ways of resetting a VM: in-host and full reset. The first one does not free any resources and reset a RUNNING VM instance at the hypervisor level:

- `reboot`: Gracefully reboots a running VM, sending the ACPI signal.
- `reboot --hard`: Performs a 'hard' reboot.

A VM instance can be reset in any state with:

- `delete --recreate`: Deletes the VM as described above, but instead of disposing it the VM is moving again to PENDING state. As the delete operation this action should be used when the VM is not responding to other actions. Try `undeploy` or `undeploy --hard` first.

Delaying VM Instances...

The deployment of a PENDING VM (e.g. after creating or resuming it) can be delayed with:

- `hold`: Sets the VM to hold state. The scheduler will not deploy VMs in the hold state. Please note that VMs can be created directly on hold, using `'onemplate instantiate --hold'` or `'onevm create --hold'`.

Then you can resume it with:

- `release`: Releases a VM from hold state, setting it to pending. Note that you can automatically release a VM by scheduling the operation as explained below

Disk Snapshots

There are two kinds of operations related to disk snapshots:

- `disk-snapshot-create`, `disk-snapshot-revert`, `disk-snapshot-delete`: Allows the user to take snapshots of the disk states and return to them during the VM life-cycle. It is also possible to delete snapshots.
- `disk-saveas`: To export a VM disk to an image. This is a live action.

Managing disk snapshots

Todo

Check VM states

A user can take snapshots of the disk states at any moment in time (if the VM is in RUNNING, POWEROFF or SUSPENDED states). These snapshots are organized in a tree-like structure, meaning that every snapshot has a parent, except for the first snapshot whose parent is `-1`. At any given time a user can revert the disk state to a previously taken snapshot. The active snapshot, the one the user has last reverted to, or taken, will act as the parent of the next snapshot. In addition, it's possible to delete snapshots that are not active and that have no children.

- `disk-snapshot-create <vmid> <diskid> <tag>`: Creates a new snapshot of the specified disk.

- `disk-snapshot-revert <vmid> <diskid> <snapshot_id>`: Reverts to the specified snapshot. The snapshots are immutable, therefore the user can revert to the same snapshot one and again, the disk will return always to the state of the snapshot at the time it was taken.
- `disk-snapshot-delete <vmid> <diskid> <snapshot_id>`: Deletes a snapshot if it has no children and is not active.

These actions are available for both persistent and non-persistent images. In the case of persistent images the snapshots **will** be preserved upon VM termination and will be able to be used by other VMs using that image. See the *snapshots* section in the Images guide for more information.

Warning: These actions are not in sync with the hypervisor. If the VM is in `RUNNING` state make sure the disk is unmounted (preferred), synced or quiesced in some way or another before taking the snapshot.

Todo

Review the following list. QCow2, etc. Extend info on Filesystem DS snapshots

The snapshots are implemented differently depending on the storage backend:

- **Ceph:** They are actual protected snapshots. When a revert is performed a clone of the snapshot is issued. It requires RBD format 2.
- **Filesystem:** Being ssh, raw, shared and Qcow2.
- **Others:** Not implemented.

Exporting disk images with `disk-saveas`

Any VM disk can be exported to a new image (if the VM is in `RUNNING`, `POWEROFF` or `SUSPENDED` states). This is a live operation that happens immediately. This operation accepts `--snapshot <snapshot_id>` as an optional argument, which specifies a disk snapshot to use as the source of the clone, instead of the current disk state (value by default).

Note: This action is called `onevm disk-snapshot --live` in OpenNebula \leq 4.14 but has been renamed to `onevm disk-saveas`

Warning: This action is not in sync with the hypervisor. If the VM is in `RUNNING` state make sure the disk is unmounted (preferred), synced or quiesced in some way or another before taking the snapshot.

Disk Hotplugging

New disks can be hot-plugged to running VMs with the `onevm disk-attach` and `disk-detach` commands. For example, to attach to a running VM the Image named **storage**:

```
$ onevm disk-attach one-5 --image storage
```

To detach a disk from a running VM, find the disk ID of the Image you want to detach using the `onevm show` command, and then simply execute `onevm detach vm_id disk_id`:

```
$ onevm show one-5
...
DISK=[
  DISK_ID="1",
```



```

...
]
...
$ onevm disk-detach one-5 1


```

x

Attach new disk

Virtual Machine ID:
4

Image Volatile Disk



ID	Name	Datastore	Type	Status	#VMS
4	My Saved Template	default	OS	READY	0
3	Dev Environment	default	OS	READY	0
2	CentOS with Apache	default	OS	READY	0
1	My saved template	default	OS	READY	0

You selected the following image: Dev Environment

« 1 2 »

▼ Advanced options

Attach

NIC Hotplugging

You can hotplug network interfaces to VMs in the `RUNNING`, `POWEROFF` or `SUSPENDED` states. Simply specify the network where the new interface should be attach to, for example:

```

$ onevm show 2

VIRTUAL MACHINE 2 INFORMATION
ID                : 2
NAME              : centos-server
USER              : ruben
GROUP             : oneadmin
STATE             : ACTIVE
LCM_STATE         : RUNNING
RESCHED          : No
HOST              : cloud01

...

VM NICS
ID NETWORK        VLAN BRIDGE   IP           MAC
0 net_172         no vbr0     172.16.0.201 02:00:ac:10:0

```

```

VIRTUAL MACHINE HISTORY
SEQ HOST          REASON          START          TIME          PROLOG_TIME
 0 cloud01        none           03/07 11:37:40 0d 00h02m14s 0d 00h00m00s
...
$ onevm nic-attach 2 --network net_172
    
```

After the operation you should see two NICs, 0 and 1:

```

$ onevm show 2
VIRTUAL MACHINE 2 INFORMATION
ID           : 2
NAME        : centos-server
USER        : ruben
GROUP       : oneadmin
...

VM NICS
ID NETWORK    VLAN BRIDGE    IP           MAC
0 net_172     no vbr0      172.16.0.201 02:00:ac:10:00:c9
              fe80::400:acff:fe10:c9
1 net_172     no vbr0      172.16.0.202 02:00:ac:10:00:ca
              fe80::400:acff:fe10:ca
...
    
```

Also, you can detach a NIC by its ID. If you want to detach interface 1 (MAC=02:00:ac:10:00:ca), just execute:

```
$ onevm nic-detach 2 1
```

✕

Attach new nic

Virtual Machine ID:

ID	Owner	Group	Name	Reservation	Cluster	Leases	VLAN ID
1	oneadmin	BlueVDC	Private Network	No	HPC	13 / 400	-
0	oneadmin	BlueVDC	Public Network	No	HPC	1 / 100	-

« 1 »

Please select a network from the list Public Network

▼ Advanced options

Attach

Snapshotting

You can create, delete and restore snapshots for running VMs. A snapshot will contain the current disks and memory state.

Warning: The snapshots will only be available during the `RUNNING` state. If the state changes (stop, migrate, etc...) the snapshots **will** be lost.


```
$ onevm snapshot-create 4 "just in case"

$ onevm show 4
...
SNAPSHOTS
  ID      TIME NAME                                HYPERVISOR_ID
  0  02/21 16:05 just in case                onesnap-0

$ onevm snapshot-revert 4 0 --verbose
VM 4: snapshot reverted
```

Please take into consideration the following limitations:

- **The snapshots are lost if any life-cycle operation is performed, e.g. a suspend, migrate, delete request.**
- KVM: Snapshots are only available if all the VM disks use the *qcow2 driver*.
- VMware: the snapshots will persist in the hypervisor after any life-cycle operation is performed, but they will not be available to be used with OpenNebula.
- Xen: does not support snapshotting



Virtual Machine 4 oneadmin OpenNebula

Refresh Back Play Pause Stop Reset Menu
VNC User Delete

Info Capacity Storage Network Snapshots Placement Actions Template Log

ID	Target	Image / Format-Size	Persistent	Save as	Actions	Attach disk
0	hda	ttylinux - kvm_file0	NO	-	Snapshot Detach	

OpenNebula 4.8.0 by C12G Labs.

Resizing a VM Capacity

You may re-size the capacity assigned to a Virtual Machine in terms of the virtual CPUs, memory and CPU allocated. VM re-sizing can be done when the VM is not `ACTIVE`, and so in any of the following states: `PENDING`, `HOLD`, `FAILED` and specially in `POWEROFF`.

If you have created a Virtual Machine and you need more resources, the following procedure is recommended:

- Perform any operation needed to prepare your Virtual Machine for shutting down, e.g. you may want to manually stop some services...
- Poweroff the Virtual Machine

- Re-size the VM
- Resume the Virtual Machine using the new capacity

Note that using this procedure the VM will preserve any resource assigned by OpenNebula (e.g. IP leases)

The following is an example of the previous procedure from the command line (the Sunstone equivalent is straight forward):

```
> onevm poweroff web_vm  
> onevm resize web_vm --memory 2G --vcpu 2  
> onevm resume web_vm
```

From Sunstone:

The screenshot shows the 'Instantiate VM Template' window in Sunstone. At the top, there are input fields for 'VM Name' (containing 'ttylinux virtio'), 'Number of instances' (containing '1'), and a 'Hold' checkbox. Below this, the 'Disks' section shows 'ttylinux-vd' selected with a size of 10 GBs. The 'Custom Attributes' section has a text input field with the placeholder 'This is a password'. A green 'Instantiate' button is located at the bottom right of the form.

Resizing a VM Disks

If the disks assigned to a Virtual Machine need more size, this can be achieved at instantiation time of the VM. The `SIZE` parameter of the disk can be adjusted and, if it is bigger than the original size of the image, OpenNebula will:

- Increase the size of the disk container prior to launching the VM
- Using the *contextualization packages*, at boot time the VM will grow the filesystem to adjust to the new size.

This can be

```
DISK=[IMAGE_ID=4,
      SIZE=2000] # If Image 4 is 1 GB, OpenNebula will resize it to 2 GB
```

This can also be achieved from Sunstone, both in Cloud and Admin View, at the time of instantiating a VM Template:

The screenshot shows a web interface for instantiating a VM template. At the top, there are input fields for 'VM Name' and 'Number of instances' (set to 1), and a 'Hold' checkbox. Below this, the template name 'ttylinux virtio' is displayed. A 'Disks' section shows a slider for 'ttylinux-vd' set to 10 GBs, with a cost of 120.00 COST / HOUR. A 'Custom Attributes' section has a text input field with the placeholder 'This is a password'. An 'Instantiate' button is at the bottom right.

Cloning a VM

A VM instance can be saved back to a new VM Template. To do that, `poweroff` the VM and then use the `onevm save` command:

```
$ onevm save web_vm copy_of_web_vm
Template ID: 26
```

The clone takes into account the customization available to end users through Sunstone. This action clones the VM source Template, replacing the disks with snapshots of the current disks (see the `disk-snapshot` action). If the VM instance was resized, the current capacity is also used. NIC interfaces are also overwritten with the ones from the VM instance, to preserve any attach/detach action.

Please bear in mind the following limitations:

- The VM's source Template will be used. If this Template was updated since the VM was instantiated, the new contents will be used.

- Volatile disks cannot be saved, and the current contents will be lost. The cloned VM Template will contain the definition for an empty volatile disk.
- Disks and NICs will only contain the target Image/Network ID. If your Template requires extra configuration (such as DISK/DEV_PREFIX), you will need to update the new Template.

Todo

command in Sunstone

Scheduling Actions

Most of the `onevm` commands accept the `--schedule` option, allowing users to delay the actions until the given date and time.

Here is an usage example:

```
$ onevm suspend 0 --schedule "09/20"
VM 0: suspend scheduled at 2013-09-20 00:00:00 +0200

$ onevm resume 0 --schedule "09/23 14:15"
VM 0: resume scheduled at 2013-09-23 14:15:00 +0200

$ onevm show 0
VIRTUAL MACHINE 0 INFORMATION
ID                : 0
NAME              : one-0

[...]

SCHEDULED ACTIONS
ID ACTION          SCHEDULED          DONE MESSAGE
0 suspend          09/20 00:00         -
1 resume           09/23 14:15         -
```

These actions can be deleted or edited using the `onevm update` command. The time attributes use Unix time internally.

```
$ onevm update 0

SCHED_ACTION=[
  ACTION="suspend",
  ID="0",
  TIME="1379628000" ]
SCHED_ACTION=[
  ACTION="resume",
  ID="1",
  TIME="1379938500" ]
```

These are the commands that can be scheduled:

- `shutdown`
- `shutdown --hard`
- `undeploy`
- `undeploy --hard`
- `hold`

- release
- stop
- suspend
- resume
- delete
- delete-recreate
- reboot
- reboot --hard
- poweroff
- poweroff --hard
- snapshot-create

User Defined Data

Custom tags can be associated to a VM to store metadata related to this specific VM instance. To add custom attributes simply use the `onevm update` command.

```
$ onevm show 0
...
VIRTUAL MACHINE TEMPLATE
...
VMID="0"

$ onevm update 0
ROOT_GENERATED_PASSWORD="1234"
~
~

$onevm show 0
...
VIRTUAL MACHINE TEMPLATE
...
VMID="0"

USER TEMPLATE
ROOT_GENERATED_PASSWORD="1234"
```

Manage VM Permissions

OpenNebula comes with an advanced ACL rules permission mechanism intended for administrators, but each VM object has also implicit permissions that can be managed by the VM owner. To share a VM instance with other users, to allow them to list and show its information, use the `onevm chmod` command:

```
$ onevm show 0
...
PERMISSIONS
OWNER      : um-
GROUP      : ---
```

```
OTHER          : ---
$ onevm chmod 0 640
$ onevm show 0
...
PERMISSIONS
OWNER          : um-
GROUP         : u--
OTHER         : ---
```

Administrators can also change the VM's group and owner with the `chgrp` and `chown` commands.

Life-Cycle Operations for Administrators

There are some `onevm` commands operations meant for the cloud administrators:

Scheduling:

- `resched`: Sets the reschedule flag for the VM. The Scheduler will migrate (or migrate `-live`, depending on the Scheduler configuration) the VM in the next monitorization cycle to a Host that better matches the requirements and rank restrictions. Read more in the Scheduler documentation.
- `unresched`: Clears the reschedule flag for the VM, canceling the rescheduling operation.

Deployment:

- `deploy`: Starts an existing VM in a specific Host.
- `migrate --live`: The Virtual Machine is transferred between Hosts with no noticeable downtime. This action requires a shared file system storage.
- `migrate`: The VM gets stopped and resumed in the target host. In an infrastructure with multiple system datastores, the VM storage can be also migrated (the datastore id can be specified).

Note: By default, the above operations do not check the target host capacity. You can use the `-e` (`-enforce`) option to be sure that the host capacity is not overcommitted.

Troubleshooting:

- `recover`: If the VM is stuck in any other state (or the boot operation does not work), you can recover the VM by simulating the failure or success of the missing action, or you can launch it with the `--retry` flag (and optionally the `--interactive` if its a Transfer Manager problem) to replay the driver actions. Read the Virtual Machine Failures guide for more information.
- `migrate` or `resched`: A VM in the UNKNOWN state can be booted in a different host manually (`migrate`) or automatically by the scheduler (`resched`). This action must be performed only if the storage is shared, or manually transferred by the administrator. OpenNebula will not perform any action on the storage for this migration.

1.5.3 Sunstone

You can manage your virtual machines using the `onevm` command or Sunstone.

In Sunstone, you can easily instantiate currently defined *templates* by clicking `New` on the Virtual Machines tab and manage the life cycle of the new instances

OpenNebula Sunstone Virtual Machines

oneadmin OpenNebula

Dashboard System Virtual Resources Virtual Machines Templates Images Files & Kernels Infrastructure Marketplace OneFlow Support

ID	Owner	Group	Name	Status	Host	IPs	VNC
14	John	BlueVDC	Frontend_0_(service_5)	RUNNING	dev2	192.168.1.13	
13	John	BlueVDC	DB_0_(service_5)	RUNNING	dev1	192.168.1.12	
12	John	BlueVDC	Worker_0_(service_4)	RUNNING	dev2	192.168.1.11	
11	John	BlueVDC	Master_0_(service_4)	RUNNING	dev1	192.168.1.10	
10	John	BlueVDC	Slave_2_(service_3)	RUNNING	dev2	192.168.1.9	
9	John	BlueVDC	Slave_1_(service_3)	RUNNING	dev1	192.168.1.8	
8	John	BlueVDC	Slave_0_(service_3)	RUNNING	dev2	192.168.1.7	
6	Doe	BlueVDC	Sinatra - Mongo	RUNNING	dev2	192.168.1.5	
5	BlueVDC-admin	BlueVDC	Testing Env	RUNNING	dev1	192.168.1.4	
4	BlueVDC-admin	BlueVDC	Prod Machine	RUNNING	dev2	192.168.1.3	

Showing 1 to 10 of 14 entries

15 TOTAL 13 ACTIVE 2 OFF 0 PENDING 0 FAILED

OpenNebula 4.8.0 by C12G Labs.

Using the noVNC Console

In order to use this feature, make sure that:

- The VM template has a GRAPHICS section defined, that the TYPE attribute in it is set to VNC.
- The specified VNC port on the host on which the VM is deployed is accessible from the Sunstone server host.
- The VM is in running state.

If the VM supports VNC and is running, then the VNC icon on the Virtual Machines view should be visible and clickable:


ID	Owner	Group	Name	Status	Host	IPs	VNC
0	oneadmin	oneadmin	ttlinux-0	RUNNING	localhost	192.168.122.2	

Showing 1 to 1 of 1 entries

« 1 » 10

When clicking the VNC icon, the process of starting a session begins:

- A request is made and if a VNC session is possible, Sunstone server will add the VM Host to the list of allowed vnc session targets and create a random token associated to it.
- The server responds with the session token, then a noVNC dialog pops up.
- The VNC console embedded in this dialog will try to connect to the proxy either using websockets (default) or emulating them using Flash. Only connections providing the right token will be successful. Websockets are supported from Firefox 4.0 (manual activation required in this version) and Chrome. The token expires and cannot be reused.



In order to close the VNC session just close the console dialog.

Note: From Sunstone 3.8, a single instance of the VNC proxy is launched when Sunstone server starts. This instance will listen on a single port and proxy all connections from there.

1.5.4 Information for Developers and Integrators

- Although the default way to create a VM instance is to register a Template and then instantiate it, VMs can be created directly from a template file using the `onevm create` command.
- When a VM reaches the `done` state, it disappears from the `onevm list` output, but the VM is still in the database and can be retrieved with the `onevm show` command.
- OpenNebula comes with an accounting tool that reports resource usage data.
- The monitoring information, shown with nice graphs in Sunstone, can be retrieved using the XML-RPC methods `one.vm.monitoring` and `one.vmpool.monitoring`.

VIRTUAL MACHINE SETUP

2.1 Contextualization Overview

OpenNebula provides different methods to pass information to a newly created Virtual Machine. This information can be the network configuration of the VM, user credentials, init scripts and free form data.

- *Basic Contextualization*: If you only want to configure networking and root ssh keys read this guide.
- *Advanced Contextualization*: For additional topics in contextualization like adding custom init scripts and variables also read this guide.
- *Cloud-init*: To know how to use the `cloud-init` functionality with OpenNebula check this guide.
- *Windows Contextualization*: Contextualization guide specific for Windows guests. From provisioning to contextualization.

2.2 Adding Content to Your Cloud

Once you have setup your OpenNebula cloud you'll have ready the infrastructure (clusters, hosts, virtual networks and datastores) but you need to add contents to it for your users. This basically means two different things:

- Add base disk images with OS installations of your choice. Including any software package of interest.
- Define virtual servers in the form of VM Templates. We recommend that VM definitions are made by the admins as it may require fine or advanced tuning. For example you may want to define a LAMP server with the capacity to be instantiated in a remote AWS cloud.

When you have basic virtual server definitions the users of your cloud can use them to easily provision VMs, adjusting basic parameters, like capacity or network connectivity.

There are three basic methods to bootstrap the contents of your cloud, namely:

- **External Images**. If you already have disk images in any supported format (raw, qcow2, vmdk...) you can just add them to a datastore. Alternatively you can use any virtualization tool (e.g. virt-manager) to install an image and then add it to a OpenNebula datastore.
- **Install within OpenNebula**. You can also use OpenNebula to prepare the images for your cloud. The process will be as follows:
 - Add the installation medium to a OpenNebula datastore. Usually it will be a OS installation CD-ROM/DVD.
 - Create a DATABLOCK image of the desired capacity to install the OS. Once created change its type to OS and make it persistent.

- Create a new template using the previous two images. Make sure to set the OS/BOOT parameter to cdrom and enable the VNC console.
- Instantiate the template and install the OS and any additional software
- Once you are done, shutdown the VM
- **Use the OpenNebula Marketplace.** Go to the marketplace tab in Sunstone, and simply pick a disk image with the OS and Hypervisor of your choice.

Once the images are ready, just create VM templates with the relevant configuration attributes, including default capacity, networking or any other preset needed by your infrastructure.

You are done, make sure that your cloud users can access the images and templates you have just created.

2.3 Basic Contextualization

This guide shows how to automatically configure networking in the initialization process of the VM. Following are the instructions to contextualize your images to configure the network. For more in depth information and information on how to use this information for other duties head to the *Advanced Contextualization* guide.

If you are using `vcenter` drivers continue to *vcenter Contextualization*.

2.3.1 Preparing the Virtual Machine Image

To enable the Virtual Machine images to use the contextualization information written by OpenNebula we need to add to it a series of scripts that will trigger the contextualization.

You can use the images available in the Marketplace, that are already prepared, or prepare your own images. To make your life easier you can use a couple of Linux packages that do the work for you.

The contextualization package will also mount any partition labeled `swap` as swap. OpenNebula sets this label for volatile swap disks.

- Start a image (or finish its installation)
- Download and install the package for your distribution:
 - **DEB:** Compatible with Ubuntu 11.10 to 14.04 and Debian 6/7
 - **RPM:** Compatible with CentOS and RHEL 6/7
- Shutdown the VM

2.3.2 Automatic Image Resizing

OpenNebula supports image resizing on clone. This only makes the disks bigger not the partitions and filesystems it contains. The context packages can automate the resizing of the root filesystem of your linux images but some extra considerations must be taken into account.

The root filesystem must be the last partition as the grow script can not shuffle the other partitions around.

Here are specific needs for some Linux distributions:

CentOS 6

With kernels < 3.8 partitions can not be changed online and they must be modified in `initrd`. To do this there is a `dracut` module that resizes the root partition. The package is in `epel` repository so it must be configured in the machine. After installing it `dracut` must be run so `initrd` is rebuilt.

```
# yum install -y epel-release
# yum install -i dracut-modules-growroot
# dracut -f
```

CentOS 7

This distribution comes with a recent kernel that supports online partition table modification so no `initrd` modification is needed. The package needed is in `epel` repository and must be configured before attempting to install `growpart` package:

```
# yum install -y epel-release
# yum install -y cloud-utils-growpart
```

You also need an updated version of `util-linux` packages as old versions of `partx` can not update partitions when they are in use.

Debian and Ubuntu

Like CentOS 7 it needs a recent version of `util-linux` (version 2.23) and `cloud-utils` package:

```
# apt-get install -y cloud-utils
```

2.3.3 Manual Disk Resizing

If your images don't have context packages that support image resizing (version 4.14.1) or you want to resize a partition that is not the root filesystem you can follow the instructions from the previous section to install needed tools and use these commands:

```
# growpart <disk> <partition number>
```

For example to resize `sdb3` the command is:

```
# growpart /dev/sdb 3
```

This command changes the partition parameters. Make sure that the partition is the last one or the command won't succeed. The next step is growing the filesystem with the FS tools. For example, for `ext4` you'll use this command:

```
# resize2fs /dev/sdb3
```

2.3.4 Preparing the Template

Network Configuration

These packages install a generic network configuration script that will set network parameters extracted from the virtual networks, the VM Template NIC section and the Context section. OpenNebula will automatically generate network configuration information for the configuration script if `NETWORK="yes"` is added to the `CONTEXT` attribute, or alternately the Network Context checkbox is set in the Context tab of the VM Template creation dialog.

The precedence is the following, from weaker to stronger priority:

- Virtual Network Template
- Address Range Template
- Virtual Machine NIC section
- Virtual Machine Context section

Following that, setting a Gateway in the NIC section will override the Virtual Network default Gateway, and writing a DNS in the Context section for a particular NIC will override the Address Range the NIC is pulling its IP from.

Let's see an example. We define a Virtual Network like the following:

```
NAME=public
NETWORK_ADDRESS=80.0.0.0
NETWORK_MASK=255.255.255.0
GATEWAY=80.0.0.1
DNS="8.8.8.8 8.8.4.4"
```

And then in the VM template contextualization we set:

```
CONTEXT=[
  NETWORK=YES ]
```

When the template is instantiated, those parameters for eth0 are automatically set in the VM as:

```
CONTEXT=[
  DISK_ID="0",
  ETH0_DNS="8.8.8.8 8.8.4.4",
  ETH0_GATEWAY="80.0.0.1",
  ETH0_IP="80.0.0.2",
  ETH0_MASK="255.255.255.0",
  ETH0_NETWORK="80.0.0.0",
  NETWORK="YES",
  TARGET="hda" ]
```

We can override some of the parameters, for instance let's set a different Gateway and DNS for eth0 in the NIC section, and a different DNS in the Context section. So, in the VM Template

```
NIC=[ GATEWAY="80.0.0.27", DNS="80.0.0.26" ]

CONTEXT=[
  ETH0_DNS="80.0.0.80"
  NETWORK=YES ]
```

When the template is instantiated, the values that will be setting the final network configuration are:

```
CONTEXT=[
  DISK_ID="0",
  ETH0_DNS="80.0.0.80",
  ETH0_GATEWAY="80.0.0.27",
  ETH0_IP="80.0.0.2",
  ETH0_MASK="255.255.255.0",
  ETH0_NETWORK="80.0.0.0",
  NETWORK="YES",
  TARGET="hda" ]
```

If you add more than one interface to a Virtual Machine you will end with same parameters changing ETH0 to ETH1, ETH2, etc.

A complete list of parameters that can be used for network contextualization are:

Attribute	Description
<DEV>_MAC	MAC address of the interface
<DEV>_IP	IP assigned to the interface
<DEV>_NETWORK	Interface network
<DEV>_MASK	Interface net mask
<DEV>_GATEWAY	Interface gateway
<DEV>_DNS	DNS servers for the network
<DEV>_SEARCH_DOMAIN	DNS domain search path
<DEV>_IPV6	Global IPv6 assigned to the interface
<DEV>_GATEWAY6	IPv6 gateway for this interface
<DEV>_CONTEXT_FORCE_IPV4	Configure IPv4 even if IPv6 values are present
DNS	main DNS server for the machine

SSH Configuration

You can add `SSH_PUBLIC_KEY` parameter to the context to add a SSH public key to the `authorized_keys` file of the root user.

```
CONTEXT=[
  SSH_PUBLIC_KEY = "ssh-rsa AAAAB3NzaC1yc2EAAAADAQABAAQOC+vPFFwem49zcepQxsyO51YMSpuywwt6GazgpJe9vQ
]
```

If the `SSH_PUBLIC_KEY` exists as a User Template attribute, and the template is instantiated in Sunstone, this value will be used to populate `SSH_PUBLIC_KEY` value of the `CONTEXT` section. This way templates can be made generic.

If you want to know more in deep the contextualization options head to the [Advanced Contextualization guide](#).

2.3.5 vcenter Contextualization

Contextualization with `vcenter` does not have all the features available for `kvm`, `xen` or `vmware` drivers. Here is a table with the parameters supported:

Parameter	Description
<code>SET_HOST</code>	Change the hostname of the VM. In Windows the machine needs to be restarted.
<code>SSH_PUBLIC_KEY</code>	SSH public keys to add to <code>authorized_keys</code> file. This parameter only works with Linux guests.
<code>USERNAME</code>	Create a new administrator user with the given user name. Only for Windows guests.
<code>PASSWORD</code>	Password for the new administrator user. Used with <code>USERNAME</code> and only for Windows guests.
<code>DNS</code>	Add DNS entries to <code>resolv.conf</code> file. Only for Linux guests.

In Linux guests, the information can be consumed using the following command (and acted accordingly):

```
$ vmttoolsd --cmd 'info-get guestinfo.opennebula.context' | base64 -d
MYSQLPASSWORD = 'MyPassword'
ENABLEWORDPRESS = 'YES'
```

Linux Packages

The linux packages can be downloaded from its [project page](#) and installed in the guest OS. There is one rpm file for Debian and Ubuntu and an rpm for RHEL and CentOS. After installing the package shutdown the machine and create a new template.

Alternative Linux packages:

- **Arch Linux:** AUR package `one-context`

Windows Package

The official `addon-opennebula-context` provides all the necessary files to run the contextualization in Windows 2008 R2.

The contextualization procedure is as follows:

1. Download `startup.vbs` and `context.ps1` to the Windows VM and save them in `C:\`.
2. Open the Local Group Policy Dialog by running `gpedit.msc`. Under: Computer Configuration -> Windows Settings -> Scripts -> startup (right click); browse to the `startup.vbs` file and enable it as a startup script.

After that power off the VM and create a new template from it.

2.4 Advanced Contextualization

There are two contextualization mechanisms available in OpenNebula: the automatic IP assignment, and a more generic way to give any file and configuration parameters. You can use any of them individually, or both.

You can use already made packages that install context scripts and prepare udev configuration in your appliances. This is described in *Contextualization Packages for VM Images* section.

Please note that you will need to edit your template file, which you can do via `onetemplate update TEMPLATE_ID`.

2.4.1 Automatic IP Assignment

With OpenNebula you can derive the IP address assigned to the VM from the MAC address using the `MAC_PREFIX:IP` rule. In order to achieve this we provide context scripts for Debian, Ubuntu, CentOS and openSUSE based systems. These scripts can be easily adapted for other distributions, check dev.opennebula.org.

To configure the Virtual Machine follow these steps:

Warning: These actions are to configure the VM, the commands refer to the VMs root file system

- Copy the script `$ONE_SRC_CODE_PATH/share/scripts/vmcontext.sh` into the `/etc/init.d` directory in the VM root file system.
- Execute the script at boot time before starting any network service, usually runlevel 2 should work.


```
$ ln /etc/init.d/vmcontext.sh /etc/rc2.d/S01vmcontext.sh
```

Having done so, whenever the VM boots it will execute this script, which in turn would scan the available network interfaces, extract their MAC addresses, make the MAC to IP conversion and construct a `/etc/network/interfaces` that will ensure the correct IP assignment to the corresponding interface.

2.4.2 Generic Contextualization

The method we provide to give configuration parameters to a newly started virtual machine is using an ISO image (OVF recommendation). This method is network agnostic so it can be used also to configure network interfaces. In the VM description file you can specify the contents of the iso file (files and directories), tell the device the ISO image

will be accessible and specify the configuration parameters that will be written to a file for later use inside the virtual machine.



In this example we see a Virtual Machine with two associated disks. The Disk Image holds the filesystem where the Operating System will run from. The ISO image has the contextualization for that VM:

- `context.sh`: file that contains configuration variables, filled by OpenNebula with the parameters specified in the VM description file
- `init.sh`: script called by VM at start that will configure specific services for this VM instance
- `certificates`: directory that contains certificates for some service
- `service.conf`: service configuration

Warning: This is just an example of what a contextualization image may look like. Only `context.sh` is included by default. You have to specify the values that will be written inside `context.sh` and the files that will be included in the image.

Warning: To prevent regular users to copy system/secure files, the `FILES` attribute within `CONTEXT` is only allowed to OpenNebula users within the `oneadmin` group. `FILES_DS` can be used to include arbitrary files from Files Datastores.

Defining Context

In VM description file you can tell OpenNebula to create a contextualization image and to fill it with values using `CONTEXT` parameter. For example:

```
CONTEXT = [
  hostname = "MAINHOST",
  ip_private = "$NIC[IP, NETWORK=\"public net\"]",
  dns = "$NETWORK[DNS, NETWORK_ID=0]",
  root_pass = "$IMAGE[ROOT_PASS, IMAGE_ID=3]",
  ip_gen = "10.0.0.$VMID",
  files_ds = "$FILE[IMAGE=\"certificate\"] $FILE[IMAGE=\"server_license\"]"
]
```

Variables inside CONTEXT section will be added to `context.sh` file inside the contextualization image. These variables can be specified in three different ways:

Hardcoded variables

```
hostname = "MAINHOST"
```

Using template variables

`<template_variable>`: any single value variable of the VM template, like for example:

```
ip_gen = "10.0.0.$VMID"
```

`<template_variable>[<attribute>]`: Any single value contained in a multiple value variable in the VM template, like for example:

```
ip_private = $NIC[IP]
```

`<template_variable>[<attribute>, <attribute2>=<value2>]`: Any single value contained in a multiple value variable in the VM template, setting one attribute to discern between multiple variables called the same way, like for example:

```
ip_public = "$NIC[IP, NETWORK=\"Public\"]"
```

You can use any of the attributes defined in the variable, NIC in the previous example.

Using Virtual Network template variables

`$NETWORK[<vnet_attribute>, <NETWORK_ID|NETWORK|NIC_ID>=<vnet_id|vnet_name|nic_id>]`: Any single value variable in the Virtual Network template, like for example:

```
dns = "$NETWORK[DNS, NETWORK_ID=3]"
```

Using Image template variables

`$IMAGE[<image_attribute>, <IMAGE_ID|IMAGE>=<img_id|img_name>]`: Any single value variable in the Image template, like for example:

```
root = "$IMAGE[ROOT_PASS, IMAGE_ID=0]"
```

Note that the image MUST be in used by any of the DISKS defined in the template. The `image_attribute`

Using User template variables

`$USER[<user_attribute>]`: Any single value variable in the user (owner of the VM) template, like for example:

```
ssh_key = "$USER[SSH_KEY]"
```

The `user_attribute` can be `TEMPLATE` to include the whole user template in XML (base64 encoded).

Pre-defined variables, apart from those defined in the template you can use:

- `$UID`, the uid of the VM owner

- \$UNAME, the VM owner user name
- \$GID, the id of the VM group
- \$GNAME, the VM group name
- \$TEMPLATE, the whole template in XML format and encoded in base64

The file generated will be something like this:

```
# Context variables generated by OpenNebula
hostname="MAINHOST"
ip_private="192.168.0.5"
dns="192.168.4.9"
ip_gen="10.0.0.85"
files_ds="/home/cloud/var/datastores/2/3fae86a862b7539b41de350e8fa56100 /home/cloud/var/datastores/2"
target="sdb"
root="13.0"
```

Some of the variables have special meanings, but none of them are mandatory:

Attribute	Description
files_ds	Files that will be included in the contextualization image. Each file must be stored in a FILE_DS Datastore and must be of type CONTEXT
target	device where the contextualization image will be available to the VM instance. Please note that the proper device mapping may depend on the guest OS, e.g. ubuntu VMs should use hd* as the target device
file	Files and directories that will be included in the contextualization image. Specified as absolute paths, by default this can be used only by onedadmin .
init_scripts	If you want the VM to execute a script that is not called init.sh (or if you want to call more than just one script),this list contains the scripts to run and their order. Ex. <code>init.sh users.sh mysql.sh</code> will force the VM to execute init.sh , then users.sh and lastly mysql.sh at boot time'
START_SCRIPT	Text of the script executed when the machine starts up. It can contain shebang in case it is not shell script. For example <code>START_SCRIPT="yum upgrade"</code>
START_SCRIPT_BASE64	BASE64 as START_SCRIPT but encoded in Base64
TOKEN	YES to create a token.txt file for OneGate monitorization
NETWORK	YES to fill automatically the networking parameters for each NIC, used by the <i>Contextualization packages</i>
SET_HOSTNAME	This parameter value will be the hostname of the VM.
DNS_HOSTNAME	YES to set the VM hostname to the reverse dns name (from the first IP)
GATEWAY_IFACE	This variable can be set to the interface number you want to configure the gateway. It is useful when several networks have GATEWAY parameter and you want yo choose the one that configures it. For example to set the first interface to configure the gateway you use <code>GATEWAY_IFACE=0</code>

Warning: A default target attribute is *generated automatically* by OpenNebula, based on the default device prefix set at oned.conf.

Contextualization Packages for VM Images

The VM should be prepared to use the contextualization image. First of all it needs to mount the contextualization image somewhere at boot time. Also a script that executes after boot will be useful to make use of the information provided.

The file `context.sh` is compatible with `bash` syntax so you can easily source it inside a shellscript to get the variables that it contains.

Contextualization packages are available to several distributions so you can prepare them to work with OpenNebula without much effort. These are the changes they do to your VM:

- Disables `udev net` and `cd persistent` rules
- Deletes `udev net` and `cd persistent` rules
- Unconfigures the network
- Adds OpenNebula contextualization scripts to startup

Warning: These packages are destructive. The configuration for networking will be deleted. Make sure to use this script on copies of your images.

Instructions on how to install the contextualization packages are located in the *contextualization overview documentation*.

After the installation of these packages the images on start will configure the network using the mac address generated by OpenNebula. They will also try to mount the `cdrom context` image from `/dev/cdrom` and if `init.sh` is found it will be executed.

2.4.3 Generating Custom Contextualization Packages

Network configuration is a script located in `/etc/one-context.d/00-network`. Any file located in that directory will be executed on start, in alphabetical order. This way we can add any script to configure or start processes on boot. For example, we can have a script that populates `authorized_keys` file using a variable in the `context.sh`. Remember that those variables are exported to the environment and will be easily accessible by the scripts:

```
#!/bin/bash
echo "$SSH_PUBLIC_KEY" > /root/.ssh/authorized_keys
```

OpenNebula source code comes with the scripts and the files needed to generate contextualization packages. This way you can also generate custom packages tweaking the scripts that will go inside your images or adding new scripts that will perform other duties.

The files are located in `share/scripts/context-packages`:

- `base`: files that will be in all the packages. Right now it contains empty `udev` rules and the `init` script that will be executed on startup.
- `base_<type>`: files specific for linux distributions. It contains the contextualization scripts for the network and comes in `rpm` and `deb` flavors. You can add here your own contextualization scripts and they will be added to the package when you run the generation script.
- `generate.sh`: The script that generates the packages.
- `postinstall`: This script will be executed after the package installation and will clean the network and `udev` configuration. It will also add the `init` script to the started services on boot.

To generate the packages you will need:

- Ruby `>= 1.8.7`
- `gem fpm`
- `dpkg utils` for `deb` package creation
- `rpm utils` for `rpm` package creation

You can also give to the generation script some parameters using env variables to generate the packages. For example, to generate an rpm package you will execute:

```
$ PACKAGE_TYPE=rpm ./generate.sh
```

These are the default values of the parameters, but you can change any of them the same way we did for PACKAGE_TYPE:

```
VERSION=4.4.0
MAINTAINER=C12G Labs <support@c12g.com>
LICENSE=Apache
PACKAGE_NAME=one-context
VENDOR=C12G Labs
DESCRIPTION="
This package prepares a VM image for OpenNebula:
  * Disables udev net and cd persistent rules
  * Deletes udev net and cd persistent rules
  * Unconfigures the network
  * Adds OpenNebula contextualization scripts to startup

To get support use the OpenNebula mailing list:
  http://opennebula.org/community:mailinglists
"
PACKAGE_TYPE=deb
URL=http://opennebula.org
```

For more information check the README.md file from that directory.

2.5 Windows Contextualization

This guide describes the standard process of provisioning and contextualizing a Windows guest.

Note: This guide has been tested for Windows 2008 R2, however it should work with Windows systems \geq Windows 7.

2.5.1 Provisioning

Installation

Provisioning a Windows VM is performed the standard way in OpenNebula:

1. Register the Installation media (typically a DVD) into a Datastore
2. Create an empty datablock with an appropriate size, at least 10GB. Change the type to OS. If you are using a qcow2 image, don't forget to add DRIVER=qcow2 and FORMAT=qcow2.
3. Create a template that boots from CDROM, enables VNC, and references the Installation media and the Image created in step 2.
4. Follow the typical installation procedure over VNC.
5. Perform a deferred disk-snapshot of the OS disk, which will be saved upon shutdown.
6. Shutdown the VM.

The resulting image will boot under any OpenNebula cloud that uses KVM or VMware, and for any storage subsystem. However it hasn't been contextualized, therefore it will only obtain its IP via DHCP. To apply contextualization please follow the *Contextualization* section.

Sysprep

If you are adapting a pre-existing Windows VM to run in an OpenNebula environment, and you want to remove all the pre-existing sensitive data in order to be able to clone and deliver it to third party users, it's highly recommended to run [Sysprep](#) on the image. To do so execute the installation steps but instead of step 6 run `c:\Windows\System32\sysprep\sysprep.exe`. Select OOBE, Generalize and Shutdown as the action when it finishes. The image generated with this procedure will be able to configure its network and will also be a pristine Windows installation.

2.5.2 Contextualization

Enabling Contextualization

The official `addon-opennebula-context` provides all the necessary files to run the contextualization in Windows 2008 R2.

The contextualization procedure is as follows:

1. Download `startup.vbs` to the Windows VM (you can also send it via Context files) and write it to a path under `C:`.
2. Open the Local Group Policy Dialog by running `gpedit.msc`. Under: Computer Configuration -> Windows Settings -> Scripts -> startup (right click); browse to the `startup.vbs` file and enable it as a startup script.

Save the image by performing a deferred disk-snapshot of the OS disk, which will be saved upon shutdown.

To use the Windows contextualization script you need to use the previously prepared Windows image and include into the CONTEXT files the `context.ps1` script available [here](#).

Warning: The `context.ps1` name matters. If changed, the script will not run.

Features

The `context.ps1` script will:

- Add a new user (using `USERNAME` and `PASSWORD`).
- Rename the server (using `SET_HOSTNAME`).
- Enable Remote Desktop.
- Enable Ping.
- Configure the Network, using the automatically generated networking variables in the CONTEXT CD-ROM.
- Run arbitrary PowerShell scripts available in the CONTEXT CD-ROM and referenced by the `INIT_SCRIPTS` variable.

Variables

The contextualization variables supported by the Windows context script are very similar to the ones in *Linux* except for a few Windows-specific exceptions.

This is the list of supported variables:

- `<DEV>_MAC`: MAC address of the interface.
- `<DEV>_IP`: IP assigned to the interface.
- `<DEV>_NETWORK`: Interface network.
- `<DEV>_MASK`: Interface net mask.
- `<DEV>_GATEWAY`: Interface gateway.
- `<DEV>_DNS`: DNS servers for the network.
- `<DEV>_SEARCH_DOMAIN`: DNS domain search path.
- `DNS`: main DNS server for the machine.
- `SET_HOSTNAME`: Set the hostname of the machine.
- `INIT_SCRIPTS`: List of PowerShell scripts to be executed. Must be available in the CONTEXT CD-ROM.
- `USERNAME`: Create a new user.
- `PASSWORD`: Password for the new user.

Customization

The `context.ps1` script has been designed to be easily hacked and modified. Perform any changes to that script and use it locally.

2.6 Cloud-init

Since version 0.7.3 of cloud-init packages the OpenNebula context CD is supported. It is able to get and configure networking, hostname, ssh key for root and cloud-init user data. Here are the options in a table:

Option	Description
standard network options	OpenNebula network parameters in the context added by <code>NETWORK=yes</code>
HOSTNAME	VM hostname
SSH_PUBLIC_KEY	ssh public key added to root's authorized keys
USER_DATA	Specific user data for cloud-init
DSMODE	Can be set to local, net or disabled to change cloud-init datasource mode

You have more information on how to use it at the [cloud-init documentation page](#).

There are plenty of examples on what can go in the `USER_DATA` string at the [cloud-init examples page](#).

Warning: The current version of cloud-init configures the network before running cloud-init configuration. This makes the network configuration not reliable. Until a new version that fixes this is released you can add OpenNebula context packages or this user data to reboot the machine so the network is properly configured.

```
CONTEXT=[
  USER_DATA="#cloud-config
power_state:
  mode: reboot
" ]
```

2.6.1 Platform Specific Notes

CentOS

Works correctly for cloud-init \geq 0.7.4.

Ubuntu/Debian

To make it configure the network correctly it needs to be down so the network configuration part makes its work:

```
CONTEXT=[
  NETWORK="YES",
  SSH_PUBLIC_KEY="$USER[SSH_PUBLIC_KEY]",
  USER_DATA="#cloud-config
bootcmd:
  - ifdown -a
runcmd:
  - curl http://10.0.1.1:8999/I_am_alive
write_files:
-   encoding: b64
    content: RG9lcYBpdCB3b3JrPwo=
    owner: root:root
    path: /etc/test_file
    permissions: '0644'
packages:
  - ruby2.0" ]
```


OPENNEBULA MARKETPLACE

3.1 Interacting with the OpenNebula Marketplace

The OpenNebula Marketplace is a catalog of third party virtual appliances ready to run in OpenNebula environments. The OpenNebula Marketplace only contains appliances metadata. The images and files required by an appliance will not be stored in the Marketplace, but links to them.

OpenNebula Marketplace BETA


About
The OpenNebula Marketplace is a catalog of third party virtual appliances ready to run in OpenNebula environments.

Appliances
The OpenNebula Marketplace only contains appliances metadata. The images and files required by an appliance will not be stored in the Marketplace, but links to them.

Development
Our work is currently focused on providing new features that will enable users to share their own appliances. We will also provide a CLI in order to fully integrate the Marketplace with OpenNebula

Appliances [Advanced search](#)

CentOS 6.2

 **CentOS**

PUBLISHER
OpenNebula.org

CATALOG
public

DOWNLOADS
42

linux, centos [More info](#)


— This image has been created using OpenNebula. It has been tested with KVM, although since it's a RAW image it should work with Xen and even VMware (using qemu-img convert). Install Media: CentOS-6.2-x86_64-minimal.iso Disk layout: Single partition 10G ext4 Architecture: x86_64 The login information for this VM is login: root password: centos It has been prepared to work with OpenNebula, i.e. it's contextualized.

HYPERVISOR
KVM

ARCH
x86_64

FORMAT
raw

OpenNebula Sandbox KVM 3.4.1



— This image is meant to be run on a KVM hypervisor. It comes with a preconfigured OpenNebula 3.4.1, ready to manage a KVM farm. Several resources are created within OpenNebula (images, virtual networks, VM templates) to build a pilot cloud under 30 minutes. More information on the SandBox guide: <http://opennebula.org/cloud/sandbox/>. The login information for this VM is login: root

HYPERVISOR
KVM

ARCH
x86_64

FORMAT
raw

3.1.1 Using Sunstone to Interact with the OpenNebula Marketplace

Since the release 3.6, Sunstone includes a new tab that allows OpenNebula users to interact with the OpenNebula Marketplace:

The screenshot shows the OpenNebula Sunstone interface. The top navigation bar includes the OpenNebula Sunstone logo, the 'OpenNebula Marketplace' title, and user information for 'oneadmin'. The left sidebar contains navigation links: Dashboard, System, Virtual Resources (with sub-links for Virtual Machines, Templates, Images, and Files & Kernels), Infrastructure, Marketplace, OneFlow, and Support. The main content area features a table of virtual machine images. At the top right of this area, there is an 'Import' button and a search input field. The table has the following data:

Name	Publisher	Hypervisor	Arch	Format
<input type="checkbox"/> Ubuntu Server 12.04 (Precise Pangolin) - kvm	OpenNebula.org	KVM	x86_64	raw
<input type="checkbox"/> CentOS Server 6.2 - kvm	OpenNebula.org	KVM	x86_64	raw
<input type="checkbox"/> ttylinux - kvm	OpenNebula.org	KVM	x86_64	raw
<input type="checkbox"/> openSUSE Server 12.1	C12G Labs	XEN	x86_64	raw
<input type="checkbox"/> CentOS Server 5.8 - xen	C12G Labs	XEN	x86_64	raw
<input type="checkbox"/> Ubuntu Server 12.04 (Precise Pangolin) - xen	C12G Labs	XEN	x86_64	raw
<input type="checkbox"/> Debian Squeeze - xen	C12G Labs	XEN	x86_64	raw
<input type="checkbox"/> CentOS Server 6.2 - VMware	C12G Labs	VMWARE	x86_64	vmdk
<input type="checkbox"/> ttylinux - VMware	C12G Labs	VMWARE	i686	vmdk
<input type="checkbox"/> Debian Squeeze - kvm	C12G Labs	KVM	x86_64	raw

Below the table, it says 'Showing 1 to 10 of 45 entries'. At the bottom right, there is a pagination control showing page 1 of 10, with a search dropdown set to 10.

If you want to import a new appliance into your local infrastructure, you just have to select an image and click the button `import`. A new dialog box will prompt you to create a new image.



Import Appliance

The following images will be created in OpenNebula. If you want to edit parameters of the image you can do it later in the images tab

Select the datastore for the images

1: default

↓ 0 - Image Name

10GB

Ubuntu-14.04

The following template will be created in OpenNebula and the previous images will be referenced in the disks If you want to edit parameters of the template you can do it later in the templates tab

📄 Template Name

Ubuntu 14.04 - KVM

Import

After that you will be able use that image in a template in order to create a new instance.

3.1.2 Using the CLI to Interact with the OpenNebula Marketplace

You can also use the CLI to interact with the OpenNebula Marketplace:

- List appliances:

```
$ onemarket list --server http://marketplace.c12g.com
      ID                               NAME                               PUBLISHER
4fc76a938fb81d3517000001      Ubuntu Server 12.04 LTS (Precise Pangolin)  OpenNebula.org
4fc76a938fb81d3517000002                               CentOS 6.2  OpenNebula.org
4fc76a938fb81d3517000003                               ttylinux  OpenNebula.org
4fc76a938fb81d3517000004      OpenNebula Sandbox VMware 3.4.1  C12G Labs
4fcf5d0a8fb81d1bb8000001      OpenNebula Sandbox KVM 3.4.1  C12G Labs
```

- Show an appliance:

```
$ onemarket show 4fc76a938fb81d3517000004 --server http://marketplace.c12g.com
{
  "_id": {"$oid": "4fc76a938fb81d3517000004"},
  "catalog": "public",
  "description": "This image is meant to be run on a ESX hypervisor, and comes with a preconfigured C
  "downloads": 90,
```

```

"files": [
  {
    "type": "OS",
    "hypervisor": "ESX",
    "format": "VMDK",
    "size": 693729120,
    "compression": "gzip",
    "os-id": "CentOS",
    "os-release": "6.2",
    "os-arch": "x86_64",
    "checksum": {
      "md5": "2dba351902bffb4716168f3693e932e2"
    }
  }
],
"logo": "/img/logos/view_dashboard.png",
"name": "OpenNebula Sandbox VMware 3.4.1",
"opennebula_template": "",
"opennebula_version": "",
"publisher": "C12G Labs",
"tags": [
  "linux",
  "vmware",
  "sandbox",
  "esx",
  "frontend"
],
"links": {
  "download": {
    "href": "http://marketplace.c12g.com/appliance/4fc76a938fb81d3517000004/download"
  }
}
}

```

- **Create a new image:** You can use the download link as PATH in a new Image template to create an Image.

```

$ onemarket show 4fc76a938fb81d3517000004 --server http://marketplace.c12g.com
{
  ...
  "links": {
    "download": {
      "href": "http://marketplace.c12g.com/appliance/4fc76a938fb81d3517000004/download"
    }
  }
}

$ cat marketplace_image.one
NAME          = "OpenNebula Sandbox VMware 3.4.1"
PATH          = http://marketplace.c12g.com/appliance/4fc76a938fb81d3517000004/download
TYPE         = OS

$ oneimage create marketplace_image.one
ID: 1231

```

3.2 Howto Create Apps for the Marketplace

In this section some general guidelines on creating OpenNebula compatible images for the marketplace are described. Following this you will find a tutorial showing how to create an Ubuntu 12.04 image ready to distribute it in the marketplace.

3.2.1 Image Creation Guidelines

Images in the marketplace are just direct installation of OS, prepared to run with OpenNebula. There are two basic things you need to do (apart from the standard OS installation):

- Add OpenNebula contextualization script, so the image is able to receive and use context information
- Disable udev network rule writing, usually images are cloned multiple times, using different MAC addresses each time. In this case, you'll need to disable udev to prevent getting a new interface each time.

These both steps can be automated in some distributions (Debian, Ubuntu, CentOS and RHEL) using preparation packages. You can find the packages and more information about them at the *Contextualization Packages for VM Images* section.

Add OpenNebula Contextualization Script

The contextualization scripts configure the VM on startup. You can find the scripts for different distributions at the [OpenNebula repository](#). Depending on the distribution the method of installation is different so refer to the distribution documentation to do so. Make sure that these scripts are executed before the network is initialized.

You can find more information about contextualization in the *Contextualizing Virtual Machines* guide.

Disable udev Network Rule Writing

Most linux distribution upon start search for new devices and write the configuration for them. This fixes the network device for each different network mac address. This is a bad behavior in VM images as they will be used to run with very different mac addresses. You need to disable this udev configuration saving and also delete any udev network rule that could be already saved.

3.2.2 Tutorial: Preparing an Ubuntu 12.04 Xen for the Marketplace

The installation is based on the [Ubuntu documentation](#).

You will need a machine where xen is correctly configured, a bridge with internet connection and a public IP or a private IP with access to a router that can connecto the internet.

First we create an empty disk, in this case it will be 8 Gb:

```
$ dd if=/dev/zero of=ubuntu.img bs=1 count=1 seek=8G
```

Then we download netboot kernel and initrd compatible with Xen. We are using a mirror near to us but you can select one from the [Ubuntu mirrors list](#):

```
$ wget http://ftp.dat.etsit.upm.es/ubuntu/dists/precise/main/installer-amd64/current/images/netboot/
$ wget http://ftp.dat.etsit.upm.es/ubuntu/dists/precise/main/installer-amd64/current/images/netboot/
```

Now we can create a file describing the VM where the ubuntu will be installed:

```
name = "ubuntu"

memory = 256

disk = ['file:PATH/ubuntu.img,xvda,w']
vif = ['bridge=BRIDGE']

kernel = "PATH/vmlinuz"
ramdisk = "PATH/initrd.gz"
```

Change PATH to the path where the VM files are located and BRIDGE to the name of the network bridge you are going to use. After this we can start the VM:

```
$ sudo xm create ubuntu.xen
```

To connect to the VM console and proceed with the installation you can use xm console command:

```
$ sudo xm console ubuntu
```

Use the menus to configure your VM. Make sure that you configure the network correctly as this installation will use it to download packages.

After the installation is done it will reboot again into the installation. You can exit the console pressing <CTRL>+< >. Now you should shutdown the machine:

```
$ sudo xm shutdown ubuntu
```

The system is now installed in the disk image and now we must start it to configure it so it plays nice with OpenNebula. The configuratio we are going to do is:

- Disable udev network generation rules and clean any that could be saved
- Add contextualization scripts

To start the VM we will need a new xen description file:

```
name = "ubuntu1204"

memory = 512

disk = ['file:PATH/ubuntu.img,xvda,w']
vif = ['bridge=BRIDGE']

bootloader = "pygrub"
```

It is pretty similar to the other one but notice that we no longer specify kernel nor initrd and we also add the bootloader option. This will make out VM use the kernel and initrd that reside inside out VM image.

We can start it using the same command as before:

```
$ sudo xm create ubuntu-new.xen
```

And the console also works the same as before:

```
$ sudo xm console ubuntu
```

We log and become root. To disable udev network rule generation we should edit the file /lib/udev/rules.d/75-persistent-net-generator.rules and comment the line that says:

```
DRIVERS=="?* ", IMPORT{program}="write_net_rules"
```

Now to make sure that no network rules are saved we can empty the rules file:

```
# echo '' > /etc/udev/rules.d/70-persistent-net.rules
```

Copy the contextualiza located at the [OpenNebula repository](#) to `/etc/init.d` and give it write permissions. This is the script that will contextualize the VM on start.

Now we modify the file `/etc/init/networking.conf` and change the line:

```
pre-start exec mkdir -p /run/network
```

with

```
pre-start script
  mkdir -p /run/network
  /etc/init.d/vmcontext
end script
```

and also in `/etc/init/network-interface.conf` we add the line:

```
/etc/init.d/vmcontext
```

so it looks similar to:

```
pre-start script
  /etc/init.d/vmcontext
  if [ "$INTERFACE" = lo ]; then
    # bring this up even if /etc/network/interfaces is broken
    ifconfig lo 127.0.0.1 up || true
    initctl emit -n net-device-up \
      IFACE=lo LOGICAL=lo ADDR FAM=inet METHOD=loopback || true
  fi
  mkdir -p /run/network
  exec ifup --allow auto $INTERFACE
end script
```


REFERENCES

4.1 Virtual Machine Definition File

A template file consists of a set of attributes that defines a Virtual Machine. Using the command `onetemplate create`, a template can be registered in OpenNebula to be later instantiated. For compatibility with previous versions, you can also create a new Virtual Machine directly from a template file, using the `onevm create` command.

Warning: There are some template attributes that can compromise the security of the system or the security of other VMs, and can be used **only** by users in the `oneadmin` group. These attributes can be configured in `oned.conf`, the default ones are labeled with `*` in the following tables. See the complete list in the *Restricted Attributes* section.

4.1.1 Syntax

The syntax of the template file is as follows:

- Anything behind the pound or hash sign `#` is a **comment**.
- **Strings** are delimited with double quotes `"`, if a double quote is part of the string it needs to be escaped `\"`.
- **Single Attributes** are in the form:

```
NAME=VALUE
```

- **Vector Attributes** that contain several values can be defined as follows:

```
NAME=[NAME1=VALUE1, NAME2=VALUE2]
```

- **Vector Attributes** must contain at least one value.
- Attribute names are case insensitive, in fact the names are converted to uppercase internally.

4.1.2 XML Syntax

Since OpenNebula 3.4, template files can be in XML, with the following syntax:

- The root element must be `TEMPLATE`
- **Single Attributes** are in the form:

```
<NAME>VALUE</NAME>
```

- **Vector Attributes** that contain several values can be defined as follows:

```
<NAME>
  <NAME1>VALUE1</NAME1>
  <NAME2>VALUE2</NAME2>
</NAME>
```

A simple example:

```
<TEMPLATE>
  <NAME>test_vm</NAME>
  <CPU>2</CPU>
  <MEMORY>1024</MEMORY>
  <DISK>
    <IMAGE_ID>2</IMAGE_ID>
  </DISK>
  <DISK>
    <IMAGE>Data</IMAGE>
    <IMAGE_UNAME>oneadmin</IMAGE_UNAME>
  </DISK>
</TEMPLATE>
```

4.1.3 Capacity Section

The following attributes can be defined to specify the capacity of a VM.

Attribute	Description	Mandatory
NAME	Name that the VM will get for description purposes. If NAME is not supplied a name generated by one will be in the form of <code>one-<VID></code> . NOTE: When defining a Template it is the name of the VM Template. The actual name of the VM will be set when the VM Template is instantiated.	YES For Templates NO For VMs - will be set to <code>one-<vmid></code> if omitted
MEMORY	Amount of RAM required for the VM, in Megabytes.	YES
CPU	Percentage of CPU divided by 100 required for the Virtual Machine, half a processor is written 0.5. This value is used by OpenNebula and the scheduler to guide the host overcommitment.	YES
VCPU	Number of virtual cpus. This value is optional , the default hypervisor behavior is used, usually one virtual CPU.	YES - will be set to 1 if omitted, this can be changed in the driver configuration

Example:

```
NAME    = test-vm
MEMORY = 128
CPU     = 1
```

4.1.4 Showback Section

The following attributes can be defined to set the cost of a VM. Read the showback documentation for more information.

Attribute	Description	Mandatory
MEMORY_COST	Cost of each memory MB per hour.	NO
CPU_COST	Cost of each CPU per hour.	NO

4.1.5 OS and Boot Options Section

The OS system is defined with the `OS` vector attribute. The following sub-attributes are supported:

Note the hypervisor column states that the attribute is **Optional**, **Mandatory**, or **-** not supported for that hypervisor

OS Sub-Attribute	Description	XEN	KVM	VMWARE
ARCH	CPU architecture to virtualize	.	M (default i686)	M (default i686)
MACHINE	libvirt machine type. Check libvirt capabilities for the list of available machine types.	.	O	.
KERNEL	path to the OS kernel to boot the image in the host	O see (*)	O	.
KERNEL_DS	image to be used as kernel (see !!)	O see (*)	O	.
INITRD	path to the initrd image in the host	O (for kernel)	O (for kernel)	.
INITRD_DS	image to be used as ramdisk (see !!)	O (for kernel)	O (for kernel)	.
ROOT	device to be mounted as root	O (for kernel)	O (for kernel)	.
KERNEL_CMD	arguments for the booting kernel	O (for kernel)	O (for kernel)	.
BOOTLOADER	path to the bootloader executable	O see (*)	O	.
BOOT	comma separated list of boot devices types, by order of preference (first device in the list is the first device used for boot). Possible values: hd,fd,cdrom, network	O (only HVM)	M	.

(*) If no kernel/initrd or bootloader are specified a Xen HVM will be created.

(!!) Use one of `KERNEL_DS` or `KERNEL` (and `INITRD` or `INITRD_DS`).

KERNEL_DS and INITRD_DS refer to an image registered in a File Datastore and must be of type KERNEL and RAMDISK, respectively. The image should be referred using one of the following:

- \$FILE[IMAGE=<image name>], to select own files
- \$FILE[IMAGE=<image name>, <IMAGE_UNAME|IMAGE_UID>=<owner name|owner id>], to select images owned by other users, by user name or uid.
- \$FILE[IMAGE_ID=<image id>], global file selection

Example, a VM booting from sda1 with kernel /vmlinuz :

```
OS = [ KERNEL      = /vmlinuz,
      INITRD      = /initrd.img,
      ROOT        = sda1,
      KERNEL_CMD  = "ro xencons=tty console=tty1"]
```

```
OS = [ KERNEL_DS   = "$FILE[IMAGE=\"kernel 3.6\"]",
      INITRD_DS   = "$FILE[IMAGE=\"initrd 3.6\"]",
      ROOT        = sda1,
      KERNEL_CMD  = "ro xencons=tty console=tty1"]
```

4.1.6 Features Section

This section configures the features enabled for the VM.

Note the hypervisor column states that the attribute is **Optional** or **-** not supported for that hypervisor

Sub-Attribute	Description	XEN HVM	KVM
PAE	Physical address extension mode allows 32-bit guests to address more than 4 GB of memory	O	O
ACPI	Useful for power management, for example, with KVM guests it is required for graceful shutdown to work	O	O
APIC	Enables the advanced programmable IRQ management. Useful for SMP machines.	O	O
LOCALTIME	The guest clock will be synchronized to the host's configured timezone when booted. Useful for Windows VMs	.	O
HYPERV	Add hyperv extensions to the VM. The options can be configured in the driver configuration, HYPERV_OPTIONS	.	O
DEVICE_MODE	Used to change the IO emulator in Xen HVM.	O	.

```
FEATURES = [  
  PAE = "yes",  
  ACPI = "yes",  
  APIC = "no",  
  DEVICE_MODE = "qemu-dm"  
]
```

4.1.7 Disks Section

The disks of a VM are defined with the `DISK` vector attribute. You can define as many `DISK` attributes as you need. There are three types of disks:

- Persistent disks, uses an Image registered in a Datastore mark as persistent.
- Clone disks, uses an Image registered in a Datastore. Changes to the images will be discarded. A clone disk can be saved as other image.
- Volatile disks, created on-the-fly on the target hosts. Disks are disposed when the VM is shutdown and cannot be saved_as

Persistent and Clone Disks

Description	DISK Sub-Attribute	Xen	KVM	VMware
ID of the Image to use	IMAGE_ID	Mandatory (no IMAGE)	Mandatory (no IMAGE)	Mandatory (no IMAGE)
Name of the Image to use	IMAGE	Mandatory (no IMAGE_ID)	Mandatory (no IMAGE_ID)	Mandatory (no IMAGE_ID)
To select the IMAGE of a given user by her ID	IMAGE_UID	Optional	Optional	Optional
To select the IMAGE of a given user by her NAME	IMAGE_UNAME	Optional	Optional	Optional
Prefix for the emulated device this image will be mounted at. For instance, hd, sd, or vd for KVM virtio. If omitted, the dev_prefix attribute of the Image will be used	DEV_PREFIX	Optional	Optional	Optional
Device to map image disk. If set, it will overwrite the default device mapping.	TARGET	Optional	Optional	Optional
Specific image mapping driver	DRIVER	Optional e.g.: tap:aio:,file:	Optional e.g.: raw, qcow2	•
Selects the cache mechanism for the disk. Values are default, none, writethrough, writeback, directsync and unsafe. More info in the libvirt documentation	CACHE	•	Optional	•
Set how the image is exposed by the hypervisor	READONLY	Optional e.g.: yes, no. This attribute should only be used for special storage configurations	Optional e.g.: yes, no. This attribute should only be used for special storage configurations	Optional e.g.: yes, no. This attribute should only be used for special storage configurations
Set IO policy. Values are threads, native	IO	•	Optional (Needs qemu 1.1)	•
IO throttling attributes for the disk. They are specified in bytes or IOPS (IO Operations) and can be specified for the total (read+write) or specific for read or write. Total and read or write can not be used at the same time	TOTAL_BYTES_SEC, READ_BYTES_SEC, WRITE_BYTES_SEC TOTAL_IOPS_SEC, READ_IOPS_SEC, WRITE_IOPS_SEC	•	Optional (Needs qemu 1.1)	•
4.14 Virtual Machine Definition File By default these parameters are only allowed to be used by oneadmin.				75

Volatile DISKS

DISK Sub-Attribute	Description	XEN	KVM	VMWARE
TYPE	Type of the disk: swap, fs	Optional	Optional	Optional
SIZE	size in MB	Optional	Optional	Optional
FORMAT	filesystem for fs images: ext2, ext3... raw will not format the image.	Mandatory (for fs)	Mandatory (for fs)	Mandatory (for fs)
DEV_PREFIX	Prefix for the emulated device this image will be mounted at. For instance, hd, sd. If omitted, the default dev_prefix set in oned.conf will be used	Optional	Optional	Optional
TARGET	device to map disk	Optional	Optional	Optional
DRIVER	special disk mapping options. KVM: raw, qcow2. Xen: tap:aio:, file:	Optional	Optional	Optional
CACHE	Selects the cache mechanism for the disk. Values are default, none, writethrough, writeback, directsync and unsafe. More info in the libvirt documentation	•	Optional	•
READONLY	Set how the image is exposed by the hypervisor	Optional e.g.: yes, no. This attribute should only be used for special storage configurations	Optional e.g.: yes, no. This attribute should only be used for special storage configurations	Optional e.g.: yes, no. This attribute should only be used for special storage configurations
IO	Set IO policy. Values are threads, native	•	Optional	•
TOTAL_BYTES_SEC, READ_BYTES_SEC, WRITE_BYTES_SEC, TOTAL_IOPS_SEC, READ_IOPS_SEC, WRITE_BYTES_SEC	IO throttling attributes for the disk. They are specified in bytes or IOPS (IO Operations) and can be specified for the total (read+write) or specific for read or write. Total and read or write can not be used at the same time. By default these parameters are only allowed to be used by oneadmin.	•	Optional	•

Disks Device Mapping

If the TARGET attribute is not set for a disk, OpenNebula will automatically assign it using the following precedence, starting with `dev_prefix + a`:

- First **OS** type Image.
- Contextualization CDROM.
- **CDROM** type Images.
- The rest of **DATABLOCK** and **OS** Images, and **Volatile** disks.

Please visit the guide for *managing images* and the *image template reference* to learn more about the different image types.

You can find a complete description of the contextualization features in the *contextualization guide*.

The default device prefix `sd` can be changed to `hd` or other prefix that suits your virtualization hypervisor requirements. You can find more information in the daemon configuration guide.

An Example

This is a sample section for disks. There are four disks using the image repository, and two volatile ones. Note that `fs` and `swap` are generated on-the-fly:

```
# First OS image, will be mapped to sda. Use image with ID 2
DISK = [ IMAGE_ID = 2 ]

# First DATABLOCK image, mapped to sdb.
# Use the Image named Data, owned by the user named oneadmin.
DISK = [ IMAGE = "Data",
          IMAGE_UNAME = "oneadmin" ]

# Second DATABLOCK image, mapped to sdc
# Use the Image named Results owned by user with ID 7.
DISK = [ IMAGE = "Results",
          IMAGE_UID = 7 ]

# Third DATABLOCK image, mapped to sdd
# Use the Image named Experiments owned by user instantiating the VM.
DISK = [ IMAGE = "Experiments" ]

# Volatile filesystem disk, sde
DISK = [ TYPE = fs,
          SIZE = 4096,
          FORMAT = ext3 ]

# swap, sdf
DISK = [ TYPE = swap,
          SIZE = 1024 ]
```

Because this VM did not declare a CONTEXT or any disk using a CDROM Image, the first DATABLOCK found is placed right after the OS Image, in sdb. For more information on image management and moving please check the Storage guide.

4.1.8 Network Section

NIC Sub-Attribute	Description	Mandatory
NET-WORK_ID	ID of the network to attach this device, as defined by <code>onevnet</code> . Use if no NETWORK	Mandatory (No NETWORK)
NET-WORK	Name of the network to use (of those owned by user). Use if no NETWORK_ID	Mandatory (No NET-WORK_ID)
NET-WORK_UID	To select the NETWORK of a given user by her ID	Optional
NET-WORK_UNAME	To select the NETWORK of a given user by her NAME	Optional
IP	Request an specific IP from the NETWORK	Optional
MAC*	Request an specific HW address from the network interface	Optional
BRIDGE	Name of the bridge the network device is going to be attached to.	Optional
TARGET	name for the tun device created for the VM	Option for KVM and VMWare
SCRIPT	name of a shell script to be executed after creating the tun device for the VM	Optional
MODEL	hardware that will emulate this network interface. With Xen this is the type attribute of the vif. In KVM you can choose <code>virtio</code> to select its specific virtualization IO framework	Optional
WHITE_PORTS_TCP	range “: Permits access to the VM only through the specified ports in the TCP protocol. Supersedes BLACK_PORTS_TCP if defined.	Optional
BLACK_PORTS_TCP	range “: Doesn't permit access to the VM through the specified ports in the TCP protocol. Superseded by WHITE_PORTS_TCP if defined.	Optional
WHITE_PORTS_UDP	range “: Permits access to the VM only through the specified ports in the UDP protocol. Supersedes BLACK_PORTS_UDP if defined.	Optional
BLACK_PORTS_UDP	range “: Doesn't permit access to the VM through the specified ports in the UDP protocol. Superseded by WHITE_PORTS_UDP if defined.	Optional
ICMP	drop : Blocks ICMP connections to the VM. By default it's set to accept.	Optional

Warning: The PORTS and ICMP attributes require the firewalling functionality to be configured. Please read the firewall configuration guide.

Example, a VM with two NIC attached to two different networks:

```
NIC = [ NETWORK_ID = 1 ]

NIC = [ NETWORK      = "Blue",
        NETWORK_UID = 0 ]
```

For more information on setting up virtual networks please check the *Managing Virtual Networks guide*.

Network Defaults

You can define a `NIC_DEFAULT` attribute with values that will be copied to each new NIC. This is specially useful for an administrator to define configuration parameters, such as `MODEL`, that final users may not be aware of.

```
NIC_DEFAULT = [ MODEL = "virtio" ]
```

4.1.9 I/O Devices Section

The following I/O interfaces can be defined for a VM:

Note the hypervisor column states that the attribute is **Optional**, **Mandatory**, or – not supported for that hypervisor

Attribute	Description	XEN	KVM	VMWARE
INPUT	Define input devices, available sub-attributes: <ul style="list-style-type: none"> • TYPE: values are mouse or tablet • BUS: values are usb, ps2 or xen 	O (only usb tablet is supported)	O	•
GRAPHICS	Whether the VM should export its graphical display and how, available sub-attributes: <ul style="list-style-type: none"> • TYPE: values: vnc, sdl, spice • LISTEN: IP to listen on. • PORT: port for the VNC server • PASSWD: password for the VNC server • KEYMAP: keyboard configuration locale to use in the VNC display • RANDOM_PASSWD: if "YES", generate a random password for each VM 	O	O	•

Example:

```
GRAPHICS = [
  TYPE    = "vnc",
  LISTEN  = "0.0.0.0",
  PORT    = "5"]
```

Warning: For KVM hypervisor the port number is a real one, not the VNC port. So for VNC port 0 you should specify 5900, for port 1 is 5901 and so on.

Warning: If the user does not specify the port variable, OpenNebula will automatically assign `$VNC_BASE_PORT + $VMID`, allowing to generate different ports for VMs so they do not collide. The `VNC_BASE_PORT` is specified inside the `oned.conf` file.

4.1.10 Context Section

Context information is passed to the Virtual Machine via an ISO mounted as a partition. This information can be defined in the VM template in the optional section called Context, with the following attributes:

Attribute	Description	Mandatory
VARIABLE	Variables that store values related to this virtual machine or others. The name of the variable is arbitrary (in the example, we use <code>hostname</code>).	Optional
FILES*	space-separated list of paths to include in context device.	Optional
FILES_DS	space-separated list of File images to include in context device. (Not supported for vCenter)	Optional
TARGET	device to attach the context ISO.	Optional
TOKEN	YES to create a <code>token.txt</code> file for OneGate monitorization	Optional
NETWORK	YES to fill automatically the networking parameters for each NIC, used by the <i>Contextualization packages</i>	Optional

* only for users in `onedadmin` group

The values referred to by **VARIABLE** can be defined :

Hardcoded values:

```
HOSTNAME = "MAINHOST"
```

Using template variables

`<template_variable>`: any single value variable of the VM template, like for example:

```
IP_GEN = "10.0.0.$VMID"
```

`<template_variable>[<attribute>]`: Any single value contained in a multiple value variable in the VM template, like for example:

```
IP_PRIVATE = $NIC[IP]
```

`<template_variable>[<attribute>, <attribute2>=<value2>]`: Any single value contained in the variable of the VM template, setting one attribute to discern between multiple variables called the same way, like for example:

```
IP_PUBLIC = "$NIC[IP, NETWORK=\"Public\"]"
```

Using Virtual Network template variables

`$NETWORK[<vnet_attribute>, <NETWORK_ID|NETWORK|NIC_ID>=<vnet_id|vnet_name|nic_id>]`: Any single value variable in the Virtual Network template, like for example:

```
dns = "$NETWORK[DNS, NETWORK_ID=3]"
```

Note: The network **MUST** be in used by any of the NICs defined in the template. The `vnet_attribute` can be

TEMPLATE to include the whole vnet template in XML (base64 encoded).

Using Image template variables

`$IMAGE[<image_attribute>, <IMAGE_ID|IMAGE>=<img_id|img_name>]`: Any single value variable in the Image template, like for example:

```
root = "$IMAGE[ROOT_PASS, IMAGE_ID=0]"
```

Note: The image **MUST** be in used by any of the DISKS defined in the template. The image_attribute can be TEMPLATE to include the whole image template in XML (base64 encoded).

Using User template variables

`$USER[<user_attribute>]`: Any single value variable in the user (owner of the VM) template, like for example:

```
ssh_key = "$USER[SSH_KEY]"
```

Note: The user_attribute can be TEMPLATE to include the whole user template in XML (base64 encoded).

Pre-defined variables, apart from those defined in the template you can use:

- `$UID`, the uid of the VM owner
- `$UNAME`, the name of the VM owner
- `$GID`, the id of the VM owner's group
- `$GNAME`, the name of the VM owner's group
- `$TEMPLATE`, the whole template in XML format and encoded in base64

FILES_DS, each file must be registered in a FILE_DS datastore and has to be of type CONTEXT. Use the following to select files from Files Datastores:

- `$FILE[IMAGE=<image name>]`, to select own files
- `$FILE[IMAGE=<image name>, <IMAGE_UNAME|IMAGE_UID>=<owner name|owner id>]`, to select images owned by other users, by user name or uid.
- `$FILE[IMAGE_ID=<image id>]`, global file selection

Example:

```
CONTEXT = [  
  HOSTNAME = "MAINHOST",  
  IP_PRIVATE = "$NIC[IP]",  
  DNS = "$NETWORK[DNS, NAME=\"Public\"]",  
  IP_GEN = "10.0.0.$VMID",  
  FILES = "/service/init.sh /service/certificates /service/service.conf",  
  FILES_DS = "$FILE[IMAGE_ID=34] $FILE[IMAGE=\"kernel\"]",  
  TARGET = "sdc"  
]
```

4.1.11 Placement Section

The following attributes placement constraints and preferences for the VM:

Attribute	Description
SCHED_REQUIREMENTS	It is an expression that rules out provisioning hosts from list of machines suitable to run this VM.
SCHED_RANK	This field sets which attribute will be used to sort the suitable hosts for this VM. Basically, it defines which hosts are <i>more suitable</i> than others.
SCHED_DS_REQUIREMENTS	It is an expression that rules out entries from the pool of datastores suitable to run this VM.
SCHED_DS_RANK	States which attribute will be used to sort the suitable datastores for this VM. Basically, it defines which datastores are more suitable than others.

Example:

```
SCHED_REQUIREMENTS = "CPUSPEED > 1000"
SCHED_RANK         = "FREE_CPU"
SCHED_DS_REQUIREMENTS = "NAME=GoldenCephDS"
SCHED_DS_RANK      = FREE_MB
```

Predefined Host Attributes

There are some predefined Host attributes that can be used in the requirements and rank expressions:

Attribute	Meaning
NAME	Hostname.
MAX_CPU	Total CPU in the host, in (# cores * 100).
CPU_USAGE	Allocated used CPU in (# cores * 100). This value is the sum of all the CPU requested by VMs running on the host, and is updated instantly each time a VM is deployed or undeployed.
FREE_CPU	Real free CPU in (# cores * 100), as returned by the probes. This value is updated each monitorization cycle.
USED_CPU	Real used CPU in (# cores * 100), as returned by the probes. USED_CPU = MAX_CPU - FREE_CPU. This value is updated each monitorization cycle.
MAX_MEM	Total memory in the host, in KB.
MEM_USAGE	Allocated used memory in KB. This value is the sum of all the memory requested by VMs running on the host, and is updated instantly each time a VM is deployed or undeployed.
FREE_MEM	Real free memory in KB, as returned by the probes. This value is updated each monitorization cycle.
USED_MEM	Real used memory in KB, as returned by the probes. USED_MEM = MAX_MEM - FREE_MEM. This value is updated each monitorization cycle.
RUNNING_VMS	Number of VMs deployed on this host.
HYPERVISOR	Hypervisor name.

You can execute `onehost show <id> -x` to see all the attributes and their values.

Note: Check the Monitoring Subsystem guide to find out how to extend the information model and add any information probe to the Hosts.

Requirement Expression Syntax

The syntax of the requirement expressions is defined as:

```
stmt ::= expr ';'
expr ::= VARIABLE '=' NUMBER
       | VARIABLE '!=' NUMBER
       | VARIABLE '>' NUMBER
```

```

| VARIABLE '<' NUMBER
| VARIABLE '=' STRING
| VARIABLE '!=' STRING
| expr '&' expr
| expr '|' expr
| '!' expr
| '(' expr ')'

```

Each expression is evaluated to 1 (TRUE) or 0 (FALSE). Only those hosts for which the requirement expression is evaluated to TRUE will be considered to run the VM.

Logical operators work as expected (less '<', greater '>', '&' AND, '|' OR, '!' NOT), '=' means equals with numbers (floats and integers). When you use '=' operator with strings, it performs a shell wildcard pattern matching.

Any variable included in the Host template or its Cluster template can be used in the requirements. You may also use an XPath expression to refer to the attribute.

There is a special variable, CURRENT_VMS, that can be used to deploy VMs in a Host where other VMs are (not) running. It can be used only with the operators '=' and '!='

Examples:

```

# Only aquila hosts (aquila0, aquila1...), note the quotes
SCHED_REQUIREMENTS = "NAME = \"aquila*\""

# Only those resources with more than 60% of free CPU
SCHED_REQUIREMENTS = "FREE_CPU > 60"

# Deploy only in the Host where VM 5 is running
SCHED_REQUIREMENTS = "CURRENT_VMS = 5"

# Deploy in any Host, except the ones where VM 5 or VM 7 are running
SCHED_REQUIREMENTS = "(CURRENT_VMS != 5) & (CURRENT_VMS != 7)"

```

Warning: If using OpenNebula's default match-making scheduler in a hypervisor heterogeneous environment, it is a good idea to add an extra line like the following to the VM template to ensure its placement in a VMWare hypervisor enabled machine.

```
SCHED_REQUIREMENTS = "HYPERVISOR=\"vmware\""
```

Warning: Template variables can be used in the SCHED_REQUIREMENTS section.

- `$(template_variable)`: any single value variable of the VM template.
- `$(template_variable)[<attribute>]`: Any single value contained in a multiple value variable in the VM template.
- `$(template_variable)[<attribute>, <attribute2>=<value2>]`: Any single value contained in a multiple value variable in the VM template, setting one attribute to discern between multiple variables called the same way.

For example, if you have a custom probe that generates a MACS attribute for the hosts, you can do short of a MAC pinning, so only VMs with a given MAC runs in a given host.

```
SCHED_REQUIREMENTS = "MAC=\"${NIC[MAC]}\""
```


Rank Expression Syntax

The syntax of the rank expressions is defined as:

```
stmt ::= expr ';'
expr ::= VARIABLE
      | NUMBER
      | expr '+' expr
      | expr '-' expr
      | expr '*' expr
      | expr '/' expr
      | '-' expr
      | '(' expr ')'
```

Rank expressions are evaluated using each host information. '+', '-', '*', '/' and '-' are arithmetic operators. The rank expression is calculated using floating point arithmetics, and then round to an integer value.

Warning: The rank expression is evaluated for each host, those hosts with a higher rank are used first to start the VM. The rank policy must be implemented by the scheduler. Check the configuration guide to configure the scheduler.

Warning: Similar to the requirements attribute, any number (integer or float) attribute defined for the host can be used in the rank attribute

Examples:

```
# First those resources with a higher Free CPU
  SCHED_RANK = "FREE_CPU"

# Consider also the CPU temperature
  SCHED_RANK = "FREE_CPU * 100 - TEMPERATURE"
```

4.1.12 RAW Section

This optional section of the VM template is used whenever the need to pass special attributes to the underlying hypervisor arises. Anything placed in the data attribute gets passed straight to the hypervisor, unmodified.

RAW Sub-Attribute	Description	XEN	KVM	VMWARE
TYPE	Possible values are: kvm, xen, vmware	O	O	O
DATA	Raw data to be passed directly to the hypervisor	O	O	O
DATA_VMX	Raw data to be added directly to the .vmx file	.	.	O

Example:

Add a custom builder and bootloader to a Xen VM:

```
RAW      = [
  TYPE   = "xen",
  DATA  = "builder=\"linux\"
           bootloader=\"/usr/lib/xen/boot/domUloader.py\"
           bootargs=\"--entry=xvda2:/boot/vmlinuz-xenpae,/boot/vmlinuz-xenpae\" ]
```

Add a guest type and a specific scsi controller to a vmware VM:

```
RAW = [
  TYPE      = "vmware",
  DATA     = "<devices><controller type='scsi' index='0' model='lsilogic' /></devices>",
  DATA_VMX = "pciBridge0.present = \"TRUE\" \nguestOS=\"windows7srv-64\""
]
```

4.1.13 Restricted Attributes

All the **default** restricted attributes to users in the onedadmin group are summarized in the following list:

- CONTEXT/FILES
- NIC/MAC
- NIC/VLAN_ID
- NIC/BRIDGE
- NIC_DEFAULT/MAC
- NIC_DEFAULT/VLAN_ID
- NIC_DEFAULT/BRIDGE
- DISK/TOTAL_BYTES_SEC
- DISK/READ_BYTES_SEC
- DISK/WRITE_BYTES_SEC
- DISK/TOTAL_IOPS_SEC
- DISK/READ_IOPS_SEC
- DISK/WRITE_IOPS_SEC
- CPU_COST
- MEMORY_COST

These attributes can be configured in oned.conf.

4.1.14 User Inputs

USER_INPUTS provides the template creator with the possibility to dynamically ask the user instantiating the template for dynamic values that must be defined.

```
USER_INPUTS = [
  BLOG_TITLE="M|text|Blog Title",
  MYSQL_PASSWORD="M|password|MySQL Password",
  <VAR>="M|<type>|<desc>"
]

CONTEXT=[
  BLOG_TITLE="$BLOG_TITLE",
  MYSQL_PASSWORD="$MYSQL_PASSWORD" ]
```

Note that the CONTEXT references the variables defined in the USER_INPUTS so the value is injected into the VM.

Valid types are text and password.

4.2 Image Definition Template

This page describes how to define a new image template. An image template follows the same syntax as the *VM template*.

If you want to learn more about the image repository, you can do so [here](#).

Warning: There are some template attributes that can compromise the security of the system or the security of other VMs, and can be used **only** by users in the oneadmin group. These attributes can be configured in oned.conf, the default ones are labeled with * in the following tables. See the complete list in the *Restricted Attributes* section.

4.2.1 Template Attributes

The following attributes can be defined in the template.

Attribute	M / O	Value	Description
NAME	Mandatory	Any string	Name that the Image will get. Every image must have a unique name.
DESCRIPTION	Optional	Any string	Human readable description of the image for other users.
TYPE	Optional	OS, CDROM, DATABLOCK, KERNEL, RAMDISK, CONTEXT	Type of the image, explained in detail in the following section. If omitted, the default value is the one defined in oned.conf (install default is OS).
PERSISTENT	Optional	YES, NO	Persistence of the image. If omitted, the default value is NO.
PERSISTENT_TYPE	Optional	IMMUTABLE	An special persistent image, that will not be modified. This attribute should only be used for special storage configurations.
DEV_PREFIX	Optional	Any string	Prefix for the emulated device this image will be mounted at. For instance, hd, sd, or vd for KVM virtio. If omitted, the default value is the one defined in oned.conf (installation default is hd).
TARGET	Optional	Any string	Target for the emulated device this image will be mounted at. For instance, hdb, sdc. If omitted, it will be <i>assigned automatically</i> .
DRIVER	Optional	KVM: raw, qcow2 Xen: tap:aio:, file:	Specific image mapping driver. VMware is unsupported
PATH	Mandatory (if no SOURCE)	Any string	Path to the original file that will be copied to the image repository. If not specified for a DATABLOCK type image, an empty image will be created. Note that gzipped files are supported and OpenNebula will automatically decompress them. Bzip2 compressed files is also supported, but it's strongly discouraged since OpenNebula will not calculate it's size properly.
SOURCE	*Mandatory (if no PATH)	Any string	Source to be used in the DISK attribute. Useful for not file-based images.
DISK_TYPE	Optional	BLOCK, CDROM or FILE (default).	This is the type of the supporting media for the image: a block device (BLOCK) an ISO-9660 file or readonly block device (CDROM) or a plain file (FILE).
READ-ONLY	Optional	YES, NO.	This attribute should only be used for special storage configurations. It sets how the image is going to be exposed to the hypervisor. Images of type CDROM and those with PERSISTENT_TYPE set to IMMUTABLE will have READONLY set to YES. Otherwise, by default it is set to NO.
CLONE_FORMAT	Optional	thin, zeroedthick, eagerzeroedthick, thick, thin	Only for VMware images on vmdk datastores. Sets the format of the target image when cloning within the datastore. More information on types.
MD5	Optional	An md5 hash	MD5 hash to check for image integrity
SHA1	Optional	An sha1 hash	SHA1 hash to check for image integrity

Warning: Be careful when PATH points to a compressed bz2 image, since although it will work, OpenNebula will not calculate its size correctly.

Mandatory attributes for DATABLOCK images with no PATH set:

At-tribute	Value	Description
SIZE	An integer	Size in MB.
FSTYPE	String	Type of file system to be built. Plain. When the disk image is used directly by the hypervisor we can format the image, and so it is ready to be used by the guest OS. Values: ext2, ext3, ext4, ntfs, reiserfs, jfs, swap. Any other fs supported by mkfs will work if no special option is needed. Formatted. The disk image is stored in a hypervisor specific format VMDK or Qcow2. Then we cannot really make a filesystem on the image, just create the device and let the guest OS format the disk. Use raw to not to format the new image. Values: raw, qcow2, vmdk_*

4.2.2 Template Examples

Example of an OS image:

```
NAME           = "Ubuntu Web Development"
PATH           = /home/one_user/images/ubuntu_desktop.img
DESCRIPTION    = "Ubuntu 10.04 desktop for Web Development students.
Contains the pdf lessons and exercises as well as all the necessary
programming tools and testing frameworks."
```

Example of a CDROM image:

```
NAME           = "MATLAB install CD"
TYPE           = CDROM
PATH           = /home/one_user/images/matlab.iso
DESCRIPTION    = "Contains the MATLAB installation files. Mount it to install MATLAB on new OS images"
```

Example of a DATABLOCK image:

```
NAME           = "Experiment results"
TYPE           = DATABLOCK
# No PATH set, this image will start as a new empty disk
SIZE           = 3.08
FSTYPE        = ext3
DESCRIPTION    = "Storage for my Thesis experiments."
```

4.2.3 Restricted Attributes

All the **default** restricted attributes to users in the oneadmin group are summarized in the following list:

- SOURCE

4.3 Virtual Network Definition File

This page describes how to define a new Virtual Network template. A Virtual Network template follows the same syntax as the *VM template*.

If you want to learn more about the Virtual Network management, you can do so [here](#).

A VNET is defined by three parts: the physical network, the address space and the contextualization parameters. The following sections details each part.

4.3.1 Physical Network Attributes

It defines the **underlying networking infrastructure** that will support the VNET. These section includes the following values:

At-tribute	Description	Value	Mandatory
NAME	Name of the Virtual Network	String	YES
BRIDGE	Device to bind the virtual and physical network, depending on the network driver it may refer to different technologies or require host setups.	String	YES if no PHYDEV
VLAN	Set to YES to activate VLAN isolation when supported by the network drivers, see Virtual Network Manager drivers. Defaults to NO	YES/NO	NO
VLAN_ID	Identifier for the VLAN	Integer	NO
PHY-DEV	Name of the physical network device that will be attached to the bridge.	String	YES for 802.1Q

4.3.2 The Address Range

The addresses available in a VNET are defined by one or more Address Ranges (AR). Each address range is defined in an AR attribute to set a continuous address range.

Optionally, each AR can include configuration attributes (context or physical network attributes) that will override those provided by the VNET.

Note: When not manually set, MAC addresses are derived from the IPv4 address. A VNET lease for a given IP will have a MAC address in the form MAC_PREFIX:IP where MAC_PREFIX is defined in oned.conf.

IPv4 Address Range

At-tribute	Description	Mandatory
TYPE	IP4	YES
IP	First IP in the range in dot notation.	YES
MAC	First MAC, if not provided it will be generated using the IP and the MAC_PREFIX in oned.conf.	NO
SIZE	Number of addresses in this range.	YES

IPv6 Address Range

Attribute	Description	Mandatory
TYPE	IP6	YES
MAC	First MAC, if not provided it will be generated.	YES
GLOBAL_PREFIX	A /64 globally routable prefix	NO
ULA_PREFIX	A /64 unique local address (ULA) prefix corresponding to the fd00::/8 block	NO
SIZE	Number of addresses in this range.	YES

Dual IPv4-IPv6 Address Range

Attribute	Description	Mandatory
TYPE	IP4_6	YES
IP	First IP in the range in dot notation.	YES
MAC	First MAC, if not provided it will be generated using the IP and the MAC_PREFIX in oned.conf.	NO
GLOBAL_PREFIX	A /64 globally routable prefix	NO
ULA_PREFIX	A /64 unique local address (ULA) prefix corresponding to the fd00::/8 block	NO
SIZE	Number of addresses in this range.	YES

Ethernet Address Range

Attribute	Description	Mandatory
TYPE	ETHER	YES
MAC	First MAC, if not provided it will be generated randomly.	NO
SIZE	Number of addresses in this range.	YES

4.3.3 Contextualization Attributes

Attribute	Description
NETWORK_ADDRESS	Base network address
NETWORK_MASK	Network mask
GATEWAY	Router for this network, do not set when the network is not routable
DNS	Specific DNS for this network
GATEWAY6	IPv6 router for this network
CONTEXT_FORCE_IPV4	When a vnet is IPv6 the IPv4 is not configured unless this attribute is set
SEARCH_DOMAIN	Default search domains for DNS resolution

4.3.4 Examples

Sample IPv4 VNet:

```
# Configuration attributes (dummy driver)
NAME = "Private Network"
DESCRIPTION = "A private network for VM inter-communication"

BRIDGE = "bond-br0"

# Context attributes
NETWORK_ADDRESS = "10.0.0.0"
NETWORK_MASK = "255.255.255.0"
DNS = "10.0.0.1"
GATEWAY = "10.0.0.1"

#Address Ranges, only these addresses will be assigned to the VMs
AR=[TYPE = "IP4", IP = "10.0.0.10", SIZE = "100" ]

AR=[TYPE = "IP4", IP = "10.0.0.200", SIZE = "10" ]
```

Sample IPv4 VNet, using AR of just one IP:

```
# Configuration attributes (OpenvSwitch driver)
NAME          = "Public"
DESCRIPTION   = "Network with public IPs"

BRIDGE       = "br1"
VLAN         = "YES"
VLAN_ID      = 12

DNS           = "8.8.8.8"
GATEWAY      = "130.56.23.1"
LOAD_BALANCER = 130.56.23.2

AR=[ TYPE = "IP4", IP = "130.56.23.2", SIZE = "1" ]
AR=[ TYPE = "IP4", IP = "130.56.23.34", SIZE = "1" ]
AR=[ TYPE = "IP4", IP = "130.56.23.24", SIZE = "1" ]
AR=[ TYPE = "IP4", IP = "130.56.23.17", MAC= "50:20:20:20:20:21", SIZE = "1" ]
AR=[ TYPE = "IP4", IP = "130.56.23.12", SIZE = "1" ]
```

4.4 Command Line Interface

OpenNebula provides a set commands to interact with the system:

4.4.1 CLI

- `oneacct`: gets accounting data from OpenNebula
- `oneacl`: manages OpenNebula ACLs
- `onecluster`: manages OpenNebula clusters
- `onedatastore`: manages OpenNebula datastores
- `onedb`: OpenNebula database migration tool
- `onegroup`: manages OpenNebula groups
- `onehost`: manages OpenNebula hosts
- `oneimage`: manages OpenNebula images
- `onetemplate`: manages OpenNebula templates
- `oneuser`: manages OpenNebula users
- `onevdc`: manages OpenNebula Virtual DataCenters
- `onevm`: manages OpenNebula virtual machines
- `onevnet`: manages OpenNebula networks
- `onezone`: manages OpenNebula zones
- `onesecgroup`: manages OpenNebula security groups

The output of these commands can be customized by modifying the configuration files that can be found in `/etc/one/cli/`. They also can be customized on a per-user basis, in this case the configuration files should be placed in `$HOME/.one/cli`.

4.4.2 ECONE Commands

- `econe-upload`: Uploads an image to OpenNebula
- `econe-describe-images`: Lists all registered images belonging to one particular user.
- `econe-run-instances`: Runs an instance of a particular image (that needs to be referenced).
- `econe-describe-instances`: Outputs a list of launched images belonging to one particular user.
- `econe-terminate-instances`: Shutdowns a set of virtual machines (or cancel, depending on its state).
- `econe-reboot-instances`: Reboots a set of virtual machines.
- `econe-start-instances`: Starts a set of virtual machines.
- `econe-stop-instances`: Stops a set of virtual machines.
- `econe-create-volume`: Creates a new DATABLOCK in OpenNebula
- `econe-delete-volume`: Deletes an existing DATABLOCK.
- `econe-describe-volumes`: Describe all available DATABLOCKS for this user
- `econe-attach-volume`: Attaches a DATABLOCK to an instance
- `econe-detach-volume`: Detaches a DATABLOCK from an instance
- `econe-allocate-address`: Allocates a new elastic IP address for the user
- `econe-release-address`: Releases a public IP of the user
- `econe-describe-addresses`: Lists elastic IP addresses
- `econe-associate-address`: Associates a public IP of the user with a given instance
- `econe-disassociate-address`: Disassociate a public IP of the user currently associated with an instance
- `econe-create-keypair`: Creates the named keypair
- `econe-delete-keypair`: Deletes the named keypair, removes the associated keys
- `econe-describe-keypairs`: List and describe the key pairs available to the user
- `econe-register`: Registers an image

4.4.3 oneFlow Commands

- `oneflow`: oneFlow Service management
- `oneflow-template`: oneFlow Service Template management