# OpenNebula 5.8 Advanced Components Guide

*Release 5.8.3*

**OpenNebula Systems**

**Jun 19, 2019**

This document is being provided by OpenNebula Systems under the Creative Commons Attribution-NonCommercial-Share Alike License.

THE DOCUMENT IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE DOCUMENT.

# CONTENTS

# MULTI-VM APPLICATIONS AND AUTO-SCALING

## 1.1 Overview

Some applications require multiple VMs to implement their workflow. OpenNebula allows you to coordinate the deployment and resource usage of such applications through two components:

- VMGroup, to fine control the placement of related virtual machines.

- OneFlow, to define and manage multi-vm applications as single entities. OneFlow also let's you define dependencies and auto-scaling policies for the application components.

### 1.1.1 How Should I Read This Chapter

This chapter should be read after the infrastructure is properly setup, and contains working Virtual Machine templates.

Proceed to each section following these links:

- *VMGroup Management*
- *OneFlow Server Configuration*
- *OneFlow Services Management*
- *OneFlow Services Auto-scaling*

### 1.1.2 Hypervisor Compatibility

This chapter applies to all the hypervisors.

## 1.2 OneFlow Server Configuration

The OneFlow commands do not interact directly with the OpenNebula daemon, there is a server that takes the requests and manages the Service (multi-tiered application) life-cycle. This guide shows how to start OneFlow, and the different options that can be configured.

### 1.2.1 Installation

OneFlow server is shipped with the main distribution. The oneflow server is contained in the 'opennebula-flow' package, and the commands in the specific CLI package. Check the Installation guide for details of what packages you have to install depending on your distribution.

---

**Note:** Make sure you executed `install_gems` during the installation to install the required gems, in particular: `treetop`, `parse-cron`.

---

### 1.2.2 Configuration

The OneFlow configuration file can be found at `/etc/one/oneflow-server.conf`. It uses YAML syntax to define the following options:

| Option | Description |
|---|---|
| **Server Configuration** | |
| :one_xmlrpc | OpenNebula daemon host and port |
| :lcm_interval | Time in seconds between Life Cycle Manager steps |
| :host | Host where OneFlow will listen |
| :port | Port where OneFlow will listen |
| **Defaults** | |
| :default_cooldown | Default cooldown period after a scale operation, in seconds |
| :shutdown_action | Default shutdown action. Values: 'shutdown', 'shutdown-hard' |
| :action_number :action_period | Default number of virtual machines (action_number) that will receive the given call in each interval defined by action_period, when an action is performed on a Role. |
| :vm_name_template | Default name for the Virtual Machines created by oneflow. You can use any of the following placeholders: <ul><li>$SERVICE_ID</li><li>$SERVICE_NAME</li><li>$ROLE_NAME</li><li>$VM_NUMBER</li></ul> |
| **Auth** | |
| :core_auth | Authentication driver to communicate with OpenNebula core <ul><li>`cipher`: for symmetric cipher encryption of tokens</li><li>`x509`: for x509 certificate encryption of tokens</li></ul> For more information, visit the OpenNebula Cloud Auth documentation |
| **Log** | |
| :debug_level | Log debug level. 0 = ERROR, 1 = WARNING, 2 = INFO, 3 = DEBUG |

This is the default file

```
################################################################################
# Server Configuration
################################################################################

# OpenNebula daemon contact information
#
```

---

```
:one_xmlrpc: http://localhost:2633/RPC2

# Time in seconds between Life Cycle Manager steps
#
:lcm_interval: 30

# Host and port where OneFlow server will run
:host: 127.0.0.1
:port: 2474


################################################################################
# Defaults
################################################################################

# Default cooldown period after a scale operation, in seconds
:default_cooldown: 300

# Default shutdown action. Values: 'shutdown', 'shutdown-hard'
:shutdown_action: 'shutdown'

# Default oneflow action options when only one is supplied
:action_number: 1
:action_period: 60

# Default name for the Virtual Machines created by oneflow. You can use any
# of the following placeholders:
#    $SERVICE_ID
#    $SERVICE_NAME
#    $ROLE_NAME
#    $VM_NUMBER

:vm_name_template: '$ROLE_NAME_$VM_NUMBER_(service_$SERVICE_ID)'

##############################################################
# Auth
##############################################################

# Authentication driver to communicate with OpenNebula core
#   - cipher, for symmetric cipher encryption of tokens
#   - x509, for x509 certificate encryption of tokens
:core_auth: cipher

################################################################################
# Log
################################################################################

# Log debug level
#   0 = ERROR, 1 = WARNING, 2 = INFO, 3 = DEBUG
#
:debug_level: 2
```

### 1.2.3 Start OneFlow

To start and stop the server, use the `opennebula-flow` service:

```
# service opennebula-flow start
```

**Note:** By default, the server will only listen to requests coming from localhost. Change the `:host` attribute in `/etc/one/oneflow-server.conf` to your server public IP, or 0.0.0.0 so oneflow will listen on any interface.

Inside `/var/log/one/` you will find new log files for the server, and individual ones for each Service in `/var/log/one/oneflow/<id>.log`

```
/var/log/one/oneflow.error
/var/log/one/oneflow.log
```

### 1.2.4 Set the Environment Variables

By default the command line tools will use the `one_auth` file and the `http://localhost:2474` OneFlow URL. To change it, set the shell environment variables as explained in the Managing Users documentation.
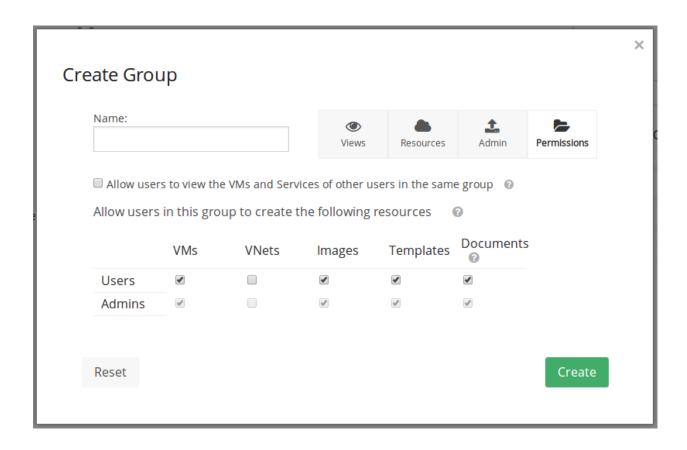
### 1.2.5 Enable the Sunstone Tabs

The OneFlow tabs (Services and Service Templates) are visible in Sunstone by default. To customize its visibility for each kind of user, visit the Sunstone views documentation

### 1.2.6 Advanced Setup

#### Permission to Create Services

By default, new groups are allowed to create Document resources. Documents are a special tool used by OneFlow to store Service Templates and instances. When a new Group is created, you can decide if you want to allow or deny its users to create OneFlow resources (Documents).
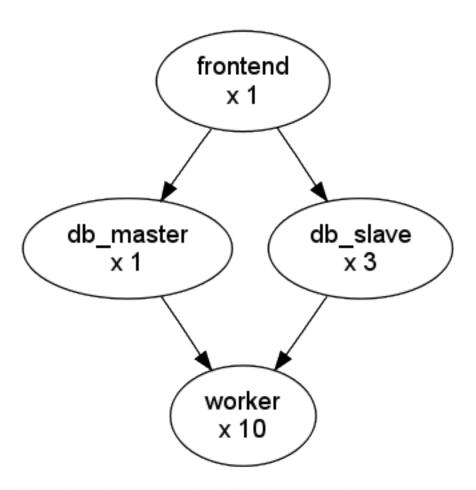
## 1.3 OneFlow Services Management

OneFlow allows users and administrators to define, execute and manage multi-tiered applications, which we call Services, composed of interconnected Virtual Machines with deployment dependencies between them. Each group of Virtual Machines is deployed and managed as a single entity, and is completely integrated with the advanced OpenNebula user and group management.

### 1.3.1 What Is a Service

The following diagram represents a multi-tier application. Each node represents a Role, and its cardinality (the number of VMs that will be deployed). The arrows indicate the deployment dependencies: each Role's VMs are deployed only when all its parent's VMs are running.

This Service can be represented with the following JSON template:

```json
{
  "name": "my_service",
  "deployment": "straight",
  "ready_status_gate": true|false,
  "roles": [
    {
      "name": "frontend",
      "vm_template": 0
    },
    {
      "name": "db_master",
      "parents": [
        "frontend"
      ],
      "vm_template": 1
    },
    {
      "name": "db_slave",
      "parents": [
        "frontend"
      ],
      "cardinality": 3,
      "vm_template": 2
    },
```

```
    {
      "name": "worker",
      "parents": [
        "db_master",
        "db_slave"
      ],
      "cardinality": 10,
      "vm_template": 3
    }
  ]
}
```

### 1.3.2 Managing Service Templates

OneFlow allows OpenNebula administrators and users to register Service Templates in OpenNebula, to be instantiated later as Services. These Templates can be instantiated several times, and also shared with other users.

Users can manage the Service Templates using the command `oneflow-template`, or Sunstone. For each user, the actual list of Service Templates available is determined by the ownership and permissions of the Templates.

#### Create and List Existing Service Templates

The command `oneflow-template create` registers a JSON template file. For example, if the previous example template is saved in /tmp/my_service.json, you can execute:

```
$ oneflow-template create /tmp/my_service.json
ID: 0
```

You can also create Service Templates from Sunstone:

To list the available Service Templates, use `oneflow-template list/show/top`:

```
$ oneflow-template list
       ID USER             GROUP           NAME
        0 oneadmin         oneadmin        my_service

$ oneflow-template show 0
SERVICE TEMPLATE 0 INFORMATION
ID                : 0
NAME              : my_service
```
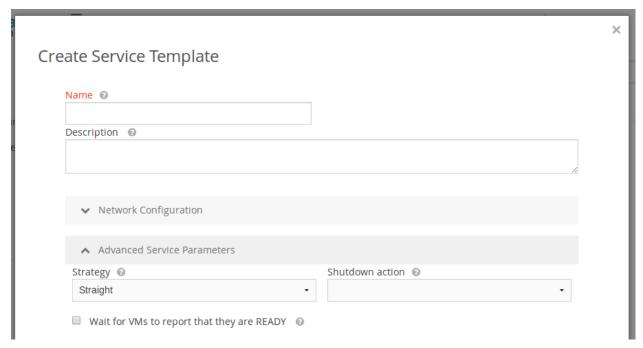
```
USER                    : oneadmin
GROUP                   : oneadmin

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

TEMPLATE CONTENTS
{
  "name": "my_service",
  "roles": [
    {

....
```

Templates can be deleted with `oneflow-template delete`.

### Determining when a VM is READY

Depending on the deployment strategy, OneFlow will wait until all the VMs in a specific Role are all in running state before deploying VMs that belong to a child Role. How OneFlow determines the running state of the VMs can be specified with the checkbox `Wait for VMs to report that the are READY` available in the Service creation dialog in Sunstone, or the attribute in `ready_status_gate` in the top level of the Service Template JSON.



If `ready_status_gate` is set to `true`, a VM will only be considered to be in running state the following points are true:

- VM is in running state for OpenNebula. Which specifically means that `LCM_STATE==3` and `STATE>=3`

- The VM has `READY=YES` in the user template.

The idea is to report via *OneGate* from inside the VM that it's running during the boot sequence:
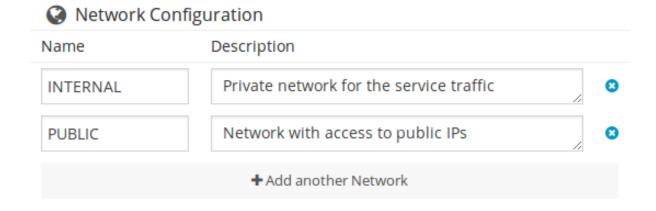
```
curl -X "PUT" http://<onegate>/vm \
  --header "X-ONEGATE-TOKEN: ..." \
  --header "X-ONEGATE-VMID: ..." \
  -d "READY = YES"
```

This can also be done directly using OpenNebula's interfaces: CLI, Sunstone or API.

If `ready_status_gate` is set to `false`, a VM will be considered to be in running state when it's in running state for OpenNebula (`LCM_STATE==3` and `STATE>=3`). Take into account that the VM will be considered RUNNING the very same moment the hypervisor boots the VM (before it loads the OS).

### Configure Dynamic Networks

Each Service Role has a Virtual Machine Template assigned. The VM Template will define the capacity, disks, and network interfaces. But instead of using the Virtual Networks set in the VM Template, the Service Template can define a set of dynamic networks.



Each Role can be attached to the dynamic networks individually.

## Roles



When a Service Template defines dynamic networks, the instantiate dialog will ask the user to select the networks to use for the new Service.

## Instantiate Service Template

Service Name ❓

Hadoop|

Number of instances ❓

1

## Network

### Private network for the service traffic

⟳                                               Search

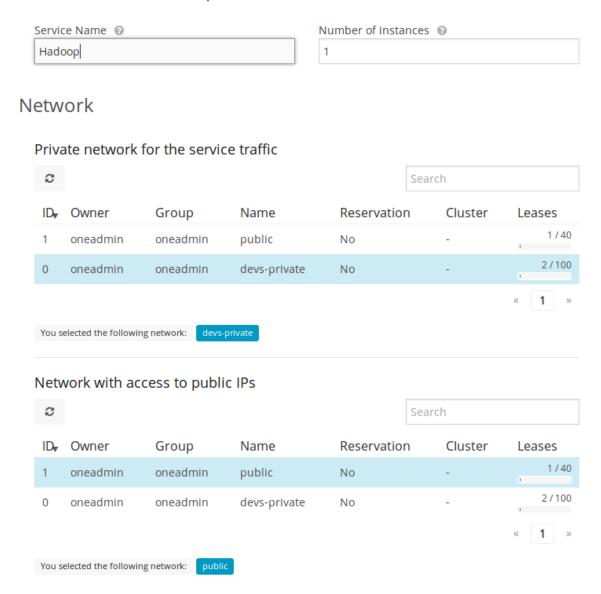| ID▾ | Owner | Group | Name | Reservation | Cluster | Leases |
|-----|-------|-------|------|-------------|---------|--------|
| 1 | oneadmin | oneadmin | public | No | - | 1 / 40 |
| 0 | oneadmin | oneadmin | devs-private | No | - | 2 / 100 |

«   1   »

You selected the following network:   devs-private

### Network with access to public IPs

⟳                                               Search

| ID▾ | Owner | Group | Name | Reservation | Cluster | Leases |
|-----|-------|-------|------|-------------|---------|--------|
| 1 | oneadmin | oneadmin | public | No | - | 1 / 40 |
| 0 | oneadmin | oneadmin | devs-private | No | - | 2 / 100 |

«   1   »

You selected the following network:   public

This allows you to create more generic Service Templates. For example, the same Service Template can be used by users of different groups that may have access to different Virtual Networks.

### 1.3.3 Managing Services

A Service Template can be instantiated as a Service. Each newly created Service will be deployed by OneFlow following its deployment strategy.
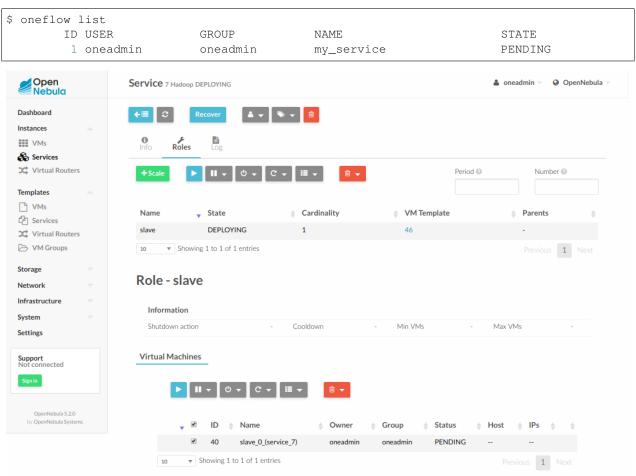
Each Service Role creates Virtual Machines in OpenNebula from VM Templates, that must be created beforehand.

#### Create and List Existing Services

New Services are created from Service Templates, using the `oneflow-template instantiate` command:

```
$ oneflow-template instantiate 0
ID: 1
```

To list the available Services, use `oneflow list/top`:

```
$ oneflow list
       ID USER             GROUP            NAME                     STATE
        1 oneadmin         oneadmin         my_service               PENDING
```



The Service will eventually change to `DEPLOYING`. You can see information for each Role and individual Virtual Machine using `oneflow show`

```
$ oneflow show 1
SERVICE 1 INFORMATION
ID                 : 1
NAME               : my_service
USER               : oneadmin
GROUP              : oneadmin
STRATEGY           : straight
SERVICE STATE      : DEPLOYING

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


ROLE frontend
ROLE STATE         : RUNNING
CARNIDALITY        : 1
```
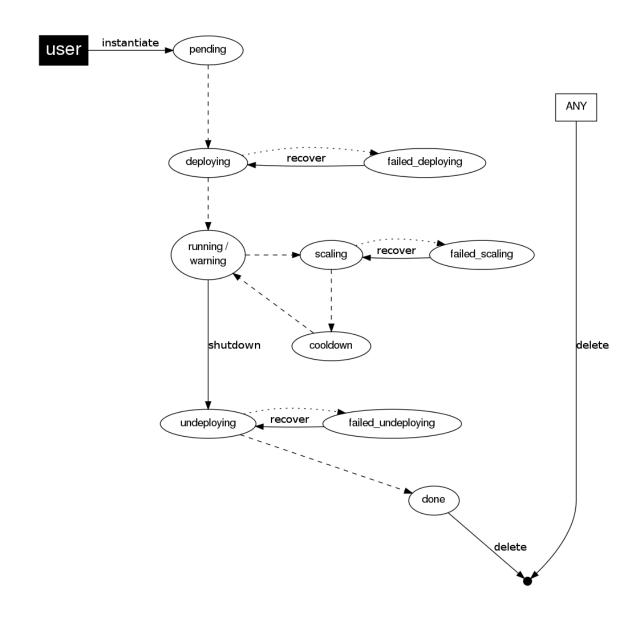
```
VM TEMPLATE        : 0
NODES INFORMATION
 VM_ID NAME                        STAT UCPU    UMEM HOST                   TIME
     0 frontend_0_(service_1)  runn   67  120.3M localhost            0d 00h01


ROLE db_master
ROLE STATE         : DEPLOYING
PARENTS            : frontend
CARNIDALITY        : 1
VM TEMPLATE        : 1
NODES INFORMATION
 VM_ID NAME                        STAT UCPU    UMEM HOST                   TIME
     1                            init          0K                   0d 00h00


ROLE db_slave
ROLE STATE         : DEPLOYING
PARENTS            : frontend
CARNIDALITY        : 3
VM TEMPLATE        : 2
NODES INFORMATION
 VM_ID NAME                        STAT UCPU    UMEM HOST                   TIME
     2                            init          0K                   0d 00h00
     3                            init          0K                   0d 00h00
     4                            init          0K                   0d 00h00


ROLE worker
ROLE STATE         : PENDING
PARENTS            : db_master, db_slave
CARNIDALITY        : 10
VM TEMPLATE        : 3
NODES INFORMATION
 VM_ID NAME                        STAT UCPU    UMEM HOST                   TIME




LOG MESSAGES
09/19/12 14:44 [I] New state: DEPLOYING
```

### Life-cycle

The `deployment` attribute defines the deployment strategy that the Life Cycle Manager (part of the *oneflow-server*)
will use. These two values can be used:

- **none**: All Roles are deployed at the same time.
- **straight**: Each Role is deployed when all its parent Roles are RUNNING.

Regardless of the strategy used, the Service will be RUNNING when all of the Roles are also RUNNING. Likewise, a
Role will enter this state only when all the VMs are running.

This table describes the Service states:

| Service State | Meaning |
|---|---|
| PENDING | The Service starts in this state, and will stay in it until the LCM decides to deploy it |
| DEPLOYING | Some Roles are being deployed |
| RUNNING | All Roles are deployed successfully |
| WARNING | A VM was found in a failure state |
| SCALING | A Role is scaling up or down |
| COOLDOWN | A Role is in the cooldown period after a scaling operation |
| UNDEPLOYING | Some Roles are being undeployed |
| DONE | The Service will stay in this state after a successful undeployment. It can be deleted |
| FAILED_DEPLOYING | An error occurred while deploying the Service |
| FAILED_UNDEPLOYING | An error occurred while undeploying the Service |
| FAILED_SCALING | An error occurred while scaling the Service |

Each Role has an individual state, described in the following table:

| Role State | Meaning |
| --- | --- |
| PENDING | The Role is waiting to be deployed |
| DEPLOYING | The VMs are being created, and will be monitored until all of them are `running` |
| RUNNING | All the VMs are `running` |
| WARNING | A VM was found in a failure state |
| SCALING | The Role is waiting for VMs to be deployed or to be shutdown |
| COOLDOWN | The Role is in the cooldown period after a scaling operation |
| UNDEPLOYING | The VMs are being shutdown. The Role will stay in this state until all VMs are `done` |
| DONE | All the VMs are `done` |
| FAILED_DEPLOYING | An error occurred while deploying the VMs |
| FAILED_UNDEPLOYING | An error occurred while undeploying the VMs |
| FAILED_SCALING | An error occurred while scaling the Role |

## Life-Cycle Operations

Services are deployed automatically by the Life Cycle Manager. To undeploy a running Service, users can use the commands `oneflow shutdown` and `oneflow delete`.

The command `oneflow shutdown` will perform a graceful a `terminate` on all the running VMs (see onevm terminate). If the `straight` deployment strategy is used, the Roles will be shutdown in the reverse order of the deployment.

After a successful shutdown, the Service will remain in the `DONE` state. If any of the VM terminate operations cannot be performed, the Service state will show `FAILED`, to indicate that manual intervention is required to complete the cleanup. In any case, the Service can be completely removed using the command `oneflow delete`.

If a Service and its VMs must be immediately undeployed, the command `oneflow delete` can be used from any Service state. This will execute a terminate operation for each VM and delete the Service. Please be aware that **this is not recommended**, because failed terminate actions may leave VMs in the system.

When a Service fails during a deployment, undeployment or scaling operation, the command `oneflow recover` can be used to retry the previous action once the problem has been solved.

## Elasticity

A Role's cardinality can be adjusted manually, based on metrics, or based on a schedule. To start the scalability immediately, use the command `oneflow scale`:

```
$ oneflow scale <serviceid> <role_name> <cardinality>
```

To define automatic elasticity policies, proceed to the *elasticity documentation guide*.

## Sharing Information between VMs

The Virtual Machines of a Service can share information with each other, using the *OneGate server*. OneGate allows Virtual Machine guests to push information to OpenNebula, and pull information about their own VM or Service.

From any VM, use the `PUT ${ONEGATE_ENDPOINT}/vm` action to store any information in the VM user template. This information will be in the form of attribute=vale, e.g. `ACTIVE_TASK = 13`. Other VMs in the Service can request that information using the `GET ${ONEGATE_ENDPOINT}/service` action.

You can read more details in the *OneGate API documentation*.

---

## 1.3.4 Managing Permissions

Both Services and Template resources are completely integrated with the OpenNebula user and group management. This means that each resource has an owner and group, and permissions. The VMs created by a Service are owned by the Service owner, so he can list and manage them.

For example, to change the owner and group of the Service 1, we can use `oneflow chown/chgrp`:

```
$ oneflow list
      ID USER            GROUP           NAME                       STATE
       1 oneadmin        oneadmin        my_service                 RUNNING

$ onevm list
   ID USER      GROUP     NAME           STAT UCPU    UMEM HOST          TIME
    0 oneadmin oneadmin frontend_0_(ser runn   17   43.5M localhost    0d 01h06
    1 oneadmin oneadmin db_master_0_(se runn   59  106.2M localhost    0d 01h06
...

$ oneflow chown my_service johndoe apptools

$ oneflow list
      ID USER            GROUP           NAME                       STATE
       1 johndoe         apptools        my_service                 RUNNING

$ onevm list
   ID USER      GROUP     NAME           STAT UCPU    UMEM HOST          TIME
    0 johndoe   apptools frontend_0_(ser runn   62   83.2M localhost    0d 01h16
    1 johndoe   apptools db_master_0_(se runn   74  115.2M localhost    0d 01h16
...
```

---

**Note:** The Service's VM ownership is also changed.

---

All Services and Templates have associated permissions for the **owner**, the users in its **group**, and **others**. For each one of these groups, there are three rights that can be set: **USE**, **MANAGE** and **ADMIN**. These permissions are very similar to those of UNIX file system, and can be modified with the command `chmod`.

For example, to allow all users in the `apptools` group to USE (list, show) and MANAGE (shutdown, delete) the Service 1:

```
$ oneflow show 1
SERVICE 1 INFORMATION
..

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

$ oneflow chmod my_service 660

$ oneflow show 1
SERVICE 1 INFORMATION
..

PERMISSIONS
```

```
OWNER             : um-
GROUP             : um-
OTHER             : ---
...
```

Another common scenario is having Service Templates created by oneadmin that can be instantiated by any user. To implement this scenario, execute:

```
$ oneflow-template show 0
SERVICE TEMPLATE 0 INFORMATION
ID                : 0
NAME              : my_service
USER              : oneadmin
GROUP             : oneadmin

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

$ oneflow-template chmod 0 604

$ oneflow-template show 0
SERVICE TEMPLATE 0 INFORMATION
ID                : 0
NAME              : my_service
USER              : oneadmin
GROUP             : oneadmin

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

Please refer to the OpenNebula documentation for more information about users & groups, and resource permissions.

### 1.3.5 Scheduling Actions on the Virtual Machines of a Role

You can use the `action` command to perform a VM action on all the Virtual Machines belonging to a Role. For example, if you want to suspend the Virtual Machines of the worker Role:

```
$ oneflow action <service_id> <role_name> <vm_action>
```

These are the commands that can be performed:

- `terminate`
- `terminate-hard`
- `undeploy`
- `undeploy-hard`
- `hold`
- `release`

---

- `stop`

- `suspend`

- `resume`

- `reboot`

- `reboot-hard`

- `poweroff`

- `poweroff-hard`

- `snapshot-create`

Instead of performing the action immediately on all the VMs, you can perform it on small groups of VMs with these options:

- `-p, --period x`: Seconds between each group of actions

- `-n, --number x`: Number of VMs to apply the action to each period

Let's say you need to reboot all the VMs of a Role, but you also need to avoid downtime. This command will reboot 2 VMs each 5 minutes:

```
$ oneflow action my-service my-role reboot --period 300 --number 2
```

The `/etc/one/oneflow-server.conf` file contains default values for `period` and `number` that are used if you omit one of them.

### 1.3.6 Recovering from Failures

Some common failures can be resolved without manual intervention, calling the `oneflow recover` command. This command has different effects depending on the Service state:

| State | New State | Recover action |
|---|---|---|
| FAILED_DEPLOYING | DEPLOYING | VMs in `DONE` or `FAILED` are terminated. |
| | | VMs in `UNKNOWN` are booted. |
| FAILED_UNDEPLOYING | UNDEPLOYING | The undeployment is resumed. |
| FAILED_SCALING | SCALING | VMs in `DONE` or `FAILED` are terminated. |
| | | VMs in `UNKNOWN` are booted. |
| | | For a scale-down, the shutdown actions are retried. |
| COOLDOWN | RUNNING | The Service is simply set to running before the cooldown period is over. |
| WARNING | WARNING | VMs in `DONE` or `FAILED` are terminated. |
| | | VMs in `UNKNOWN` are booted. |
| | | New VMs are instantiated to maintain the current cardinality. |

### 1.3.7 Service Template Reference

For more information on the resource representation, please check the API guide

Read the *elasticity policies documentation* for more information.

## 1.4 OneFlow Services Auto-scaling

A Service Role's cardinality can be adjusted manually, based on metrics, or based on a schedule.

### 1.4.1 Overview

When a scaling action starts, the Role and Service enter the `SCALING` state. In this state, the Role will `instantiate` or `terminate` a number of VMs to reach its new cardinality.

A Role with elasticity policies must define a minimum and maximum number of VMs:

```
"roles": [
    {
        "name": "frontend",
        "cardinality": 1,
        "vm_template": 0,

        "min_vms" : 1,
        "max_vms" : 5,
...
```

After the scaling, the Role and Service are in the `COOLDOWN` state for the configured duration. During a scale operation and the cooldown period, other scaling actions for the same or for other Roles are delayed until the Service is `RUNNING` again.

## 1.4.2 Set the Cardinality of a Role Manually

The command `oneflow scale` starts the scalability immediately.

```
$ oneflow scale <serviceid> <role_name> <cardinality>
```

You can force a cardinality outside the defined range with the `--force` option.

### 1.4.3 Maintain the Cardinality of a Role

The 'min_vms' attribute is a hard limit, enforced by the elasticity module. If the cardinality drops below this minimum, a scale-up operation will be triggered.

### 1.4.4 Set the Cardinality of a Role Automatically

#### Auto-scaling Types

Both elasticity_policies and scheduled_policies elements define an automatic adjustment of the Role cardinality. Three different adjustment types are supported:

- **CHANGE**: Add/subtract the given number of VMs
- **CARDINALITY**: Set the cardinality to the given number
- **PERCENTAGE_CHANGE**: Add/subtract the given percentage to the current cardinality

| Attribute | Type | Mandatory | Description |
| --- | --- | --- | --- |
| type | string | Yes | Type of adjustment. Values: CHANGE, CARDINALITY, PERCENTAGE_CHANGE |
| adjust | integer | Yes | Positive or negative adjustment. Its meaning depends on 'type' |
| min_adjust_step | integer | No | Optional parameter for PERCENTAGE_CHANGE adjustment type. If present, the policy will change the cardinality by at least the number of VMs set in this attribute. |

#### Auto-scaling Based on Metrics

Each Role can have an array of `elasticity_policies`. These policies define an expression that will trigger a cardinality adjustment.

These expressions can use performance data from

- The VM guest. Using the *OneGate server*, applications can send custom monitoring metrics to OpenNebula.
- The VM, at hypervisor level. The Virtualization Drivers return information about the VM, such as `CPU`, `NETTX` and `NETRX`.

```
"elasticity_policies" : [
  {
    "expression" : "ATT > 50",
    "type" : "CHANGE",
    "adjust" : 2,

    "period_number" : 3,
    "period" : 10
  },
  ...
]
```

The **expression** can use VM attribute names, float numbers, and logical operators (!, &, |). When an attribute is found, it will take the **average** value for all the **running VMs** that contain that attribute in the Role. If none of the VMs contain the attribute, the expression will evaluate to false.

The attribute will be looked for in `/VM/USER_TEMPLATE`, `/VM/MONITORING`, `/VM/TEMPLATE` and `/VM`, in that order. Logical operators have the usual precedence.

| Attribute | Type | Manda-tory | Description |
|---|---|---|---|
| expression | string | Yes | Expression to trigger the elasticity |
| pe-riod_number | inte-ger | No | Number of periods that the expression must be true before the elasticity is triggered |
| period | inte-ger | No | Duration, in seconds, of each period in period_number |

### Auto-scaling Based on a Schedule

Combined with the elasticity policies, each Role can have an array of `scheduled_policies`. These policies define a time, or a time recurrence, and a cardinality adjustment.

```
"scheduled_policies" : [
  {
    // Set cardinality to 2 each 10 minutes
    "recurrence" : "*/10 * * * *",

    "type" : "CARDINALITY",
    "adjust" : 2
  },
  {
    // +10 percent at the given date and time
    "start_time" : "2nd oct 2017 15:45",

    "type" : "PERCENTAGE_CHANGE",
    "adjust" : 10
  }
]
```

| Attribute | Type | Mandatory | Description |
|---|---|---|---|
| recurrence | string | No | Time for recurring adjustements. Time is specified with the Unix cron sytax |
| start_time | string | No | Exact time for the adjustement |

## 1.4.5 Visualize in the CLI

The `oneflow show / top` commands show the defined policies. When a Service is scaling, the VMs being created or terminated can be identified by an arrow next to their ID:

```
SERVICE 7 INFORMATION
...

ROLE frontend
ROLE STATE           : SCALING
CARNIDALITY          : 4
VM TEMPLATE          : 0
NODES INFORMATION
 VM_ID NAME                      STAT UCPU    UMEM HOST                        TIME
     4 frontend_0_(service_7)  runn    0   74.2M host03                   0d 00h04
```
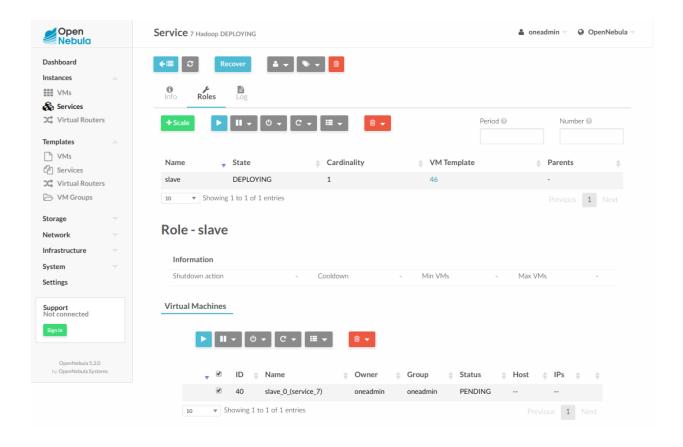
```
     5 frontend_1_(service_7)  runn    0  112.6M host02              0d 00h04
   | 6                          init         OK                      0d 00h00
   | 7                          init         OK                      0d 00h00


ELASTICITY RULES
MIN VMS             : 1
MAX VMS             : 5

ADJUST       EXPRESSION                                   EVALUATION PERIOD
+ 2          (ATT > 50) && !(OTHER_ATT = 5.5 || ABC <= 30)    0 / 3         10s
- 10 % (2)   ATT < 20                                         0 / 1          0s

ADJUST       TIME
= 6          0 9 * * mon,tue,wed,thu,fri
= 10         0 13 * * mon,tue,wed,thu,fri
= 2          30 22 * * mon,tue,wed,thu,fri


LOG MESSAGES
06/10/13 18:22 [I] New state: DEPLOYING
06/10/13 18:22 [I] New state: RUNNING
06/10/13 18:26 [I] Role frontend scaling up from 2 to 4 nodes
06/10/13 18:26 [I] New state: SCALING
```

## 1.4.6 Interaction with Individual VM Management

All the VMs created by a Service can be managed as regular VMs. When VMs are monitored in an unexpected state, this is what OneFlow interprets:

- VMs in a recoverable state ('suspend', 'poweroff', etc.) are considered healthy machines. The user will eventually decide to resume these VMs, so OneFlow will keep monitoring them. For the elasticity module, these VMs are just like 'running' VMs.

- VMs in the final 'done' state are cleaned from the Role. They do not appear in the nodes information table, and the cardinality is updated to reflect the new number of VMs. This can be seen as an manual scale-down action.

- VMs in 'unknown' or 'failed' are in an anomalous state, and the user must be notified. The Role and Service are set to the 'WARNING' state.

## 1.4.7 Examples

```
/*
Testing:

1) Update one VM template to contain
ATT = 40
and the other VM with
ATT = 60

Average will be 50, true evaluation periods will not increase in CLI output

2) Increase first VM ATT value to 45. True evaluations will increase each
10 seconds, the third time a new VM will be deployed.

3) True evaluations are reset. Since the new VM does not have ATT in its
template, the average will be still bigger than 50, and new VMs will be
deployed each 30s until the max of 5 is reached.

4) Update VM templates to trigger the scale down expression. The number of
VMs is adjusted -10 percent. Because 5 * 0.10 < 1, the adjustment is rounded to 1;
but the min_adjust_step is set to 2, so the final adjustment is -2 VMs.
*/
{
  "name": "Scalability1",
  "deployment": "none",
  "roles": [
```

```
    {
      "name": "frontend",
      "cardinality": 2,
      "vm_template": 0,

      "min_vms" : 1,
      "max_vms" : 5,

      "elasticity_policies" : [
        {
          // +2 VMs when the exp. is true for 3 times in a row,
          // separated by 10 seconds
          "expression" : "ATT > 50",

          "type" : "CHANGE",
          "adjust" : 2,

          "period_number" : 3,
          "period" : 10
        },
        {
          // -10 percent VMs when the exp. is true.
          // If 10 percent is less than 2, -2 VMs.
          "expression" : "ATT < 20",

          "type" : "PERCENTAGE_CHANGE",
          "adjust" : -10,
          "min_adjust_step" : 2
        }
      ]
    }
  ]
}
```

```
{
  "name": "Time_windows",
  "deployment": "none",
  "roles": [
    {
      "name": "frontend",
      "cardinality": 1,
      "vm_template": 0,

      "min_vms" : 1,
      "max_vms" : 15,

      // These policies set the cardinality to:
      //  6 from  9:00 to 13:00
      // 10 from 13:00 to 22:30
      //  2 from 22:30 to 09:00, and the weekend

      "scheduled_policies" : [
        {
          "type" : "CARDINALITY",
          "recurrence" : "0 9 * * mon,tue,wed,thu,fri",
          "adjust" : 6
```

---

**1.4. OneFlow Services Auto-scaling** 26

```
        },
        {
          "type" : "CARDINALITY",
          "recurrence" : "0 13 * * mon,tue,wed,thu,fri",
          "adjust" : 10
        },
        {
          "type" : "CARDINALITY",
          "recurrence" : "30 22 * * mon,tue,wed,thu,fri",
          "adjust" : 2
        }
      ]
    }
  ]
}
```

# 1.5 Virtual Machine Groups (VM Groups)

A VM Group defines a set of related VMs, and associated placement constraints for the VMs in the group. A VM Group allows you to place together (or separately) ceartain VMs (or VM classes, roles). VMGroups will help you to optimize the performance (e.g. not placing all the cpu bound VMs in the same host) or improve the fault tolerance (e.g. not placing all your front-ends in the same host) of your multi-VM applications.

## 1.5.1 Defining a VM Group

A VM Group consists of two parts: a set of roles, and a set of placement constraints for the roles. In a VM Group, a role defines a class of virtual machines that are subject to the same placement constraints and rules. Usually, you will put in the same role VMs implementing a given functionality of a multi-VM application, e.g. the front-ends or the database VMs. Additonally, you can define placement constraints for the VMs in the VM-Group, this placement rules can refer to the VMs within a role or VMs across roles.

A role is defined with the following attributes:

| Attribute | Type | Meaning |
|---|---|---|
| **NAME** | Manda-tory | The name of the role, it needs to be unique within the VM Group |
| **POLICY** | Optional | Placement policy for the VMs of the role. Possible values are: `AFFINED` and `ANTI_AFFINED` |
| **HOST_AFFINED** | Optional | Defines a set of hosts (by their ID) where the VMs of the role can be executed |
| **HOST_ANTI_AFFINED** | Optional | Defines a set of hosts (by their ID) where the VMs of the role cannot be executed |

Additional placement constraints can be imposed to the VMs of a role with the following attributes:

| Attribute | Type | Meaning |
|---|---|---|
| **AFFINED** | Optional | List of roles (comma separated) whose VMs has to be placed in the same host |
| **ANTI_AFFINED** | Optional | List of roles (comma separated) whose VMs cannot be placed in the same host |

To create a VM Group, use the Sunstone web interface, or create a template file following this example:

```
$ cat ./vmg.txt

NAME = "multi-tier server"

ROLE = [
    NAME    = "front-end",
    POLICY = "ANTI_AFFINED"
]

ROLE = [
    NAME         = "apps",
    HOST_AFFINED = "2,3,4"
]

ROLE = [ NAME = "db" ]

AFFINED = "db, apps"

$ onevmgroup create ./vmg.txt
ID: 0
```

**Note:** This guide focuses on the CLI command `onevmgroup`, but you can also manage VM Groups using Sunstone, through the VM Group tab.

### 1.5.2 Placement Policies

The following placement policies cane be applied to the VMs of a VM Group.

#### VM to Host Affinity

Specifies a set of hosts where the VMs of a role can be allocated. This policy is set in a role basis using the `HOST_AFFINED` and `HOST_ANTI_AFFINED` attributes. The host affinity rules are compatible with any other rules applied to the role VMs.

For example, if you want to place the VMs implementing the database for your application in high performance hosts you could use:

```
ROLE = [
    NAME         = "database",
    HOST_AFFINED = "1,2,3,4"
]
```

#### VM to VM Affinity

Specifies whether the VMs of a role have to be placed together in the same host (`AFFINED`) or scattered across different hosts (`ANTI_AFFINED`). The VM to VM affinity is set per role with the `POLICY` attribute.

For example, you may want to spread cpu-bound VMs across hosts to prevent contention

```
ROLE = [
    NAME   = "workers",
    POLICY = "ANTI_AFFINED"
]
```

#### Role to Role Affinity

Specifies whether the VMs of a role have to be placed together or separately with the VMs of other role. This useful to combine the Host-VM and VM-VM policies. Affinity rules for roles are set with the `AFFINED` and `ANTI_AFFINED` attributes.

For example, I want the VMs of a database to run together so they access the same storage, I want all the backup VMs to run in a separate hosts; and I want database and backups to be also in different hosts. Finally, I may have some constraints about where the database and backups may run:

```
ROLE = [
    NAME  = "apps",
    HOST_AFFINED = "1,2,3,4,5,6,7"
    POLICY = "AFFINED"
]

ROLE = [
    NAME = "backup",
```

(continues on next page)

```
    HOST_ANTI_AFFINED = "3,4"
    POLICY = "ANTI_AFFINED"
]

ANTI_AFFINED = "workers, backup"
```

> **Warning:** Note that a role policy has to be coherent with any role-role policy, i.e. a role with an `ANTI_AFFINED` policy cannot be included in any `AFFINED` role-role rule.

### Scheduler Configuration and Remarks

VMGroups are placed by dynamically generating the requirement (`SCHED_REQUIREMENTS`) of each VM an re-evaluating these expressions. Moreover, the following is also considered:

- The scheduler will look for a host with enough capacity for an affined set of VMs. If there is no such host all the affined VMs will remain pending.

- If new VMs are added to an affined role, it will pick one of the hosts where the VMs are running. By default, all should be running in the same host but if you manually migrate a VM to another host it will be considered feasible for the role.

- The scheduler does not have any synchronization point with the state of the VM group, it will start scheduling pending VMs as soon as they show up.

- Re-scheduling of VMs works as for any other VM, it will look for a different host considering the placement constraints.

## 1.5.3 Using a VM Group

Once you have defined your VM Group you can start adding VMs to it, by either picking a role and VM group at instantiation or by setting it in the VM Template. To apply a VM Group to your Virtual Machines either use the Sunstone wizard, or set the VM_GROUP attribute:

```
$ onetemplate update 0
...
VMGROUP = [ VMGROUP_NAME = "muilt-tier app", ROLE = "db" ]
```

You can also specify the `VM_GROUP` by its id (`VMGROUP_ID`), and in case of multiple groups with the same name you can select it by owner with `VMGROUP_UID`; as any other resource in OpenNebula.

> **Note:** You can also add the VMGROUP attribute when a VM is created (`onevm create`) or when the associated template is instantiated (`onetemplate instantiate`). This way the same VM template can be associated with different roles.

## 1.5.4 VM Group Management

VM Groups can be updated to edit or add new rules. Currently only role to role rules can be updated if there are no VMs in the roles. All base operations are supported for the VMGroup object: rename, chgrp, chown chmod, list, show and delete.

Note also that the same ACL/permission system is applied to VM Groups, so use access is required to place VMs in a group.

# HIGH AVAILABILITY

## 2.1 Overview

### 2.1.1 How Should I Read This Chapter

The *Front-end HA Setup* section will guide you through the process of setting an HA cluster.

If you want to enable an automatic Virtual Machine recovery in case of a Host failure, or if you want to learn how to manually recover a Virtual Machine in a failed state, please read the *Virtual Machines High Availability* section.

### 2.1.2 Hypervisor Compatibility

| Section | Compatibility |
|---|---|
| *Front-end HA Setup* | This Section applies to all Hypervisors. |
| *Virtual Machines High Availability* | This Section applies only to KVM and LXD. |

## 2.2 OpenNebula HA Setup

This guide walks you through the process of setting a high available cluster for OpenNebula core services: core (oned), scheduler (mm_sched).

OpenNebula uses a distributed consensus protocol to provide fault-tolerance and state consistency across OpenNebula services. In this section, you learn the basics of how to bootstrap and operate an OpenNebula distributed cluster.

> **Warning:** If you are interested in fail-over protection against hardware and operating system outages within your virtualized IT environment, check the *Virtual Machines High Availability Guide*.

### 2.2.1 Raft Overview

This section covers some internals on how OpenNebula implements Raft. You do not need to know these details to effectively operate OpenNebula on HA. These details are provided for those who wish to learn about them to fine tune their deployments.

A consensus algorithm is built around two concepts:

- **System State**, in OpenNebula the system state is the data stored in the database tables (users, ACLs, or the VMs in the system).

- **Log**, a sequence of SQL statements that are *consistently* applied to the OpenNebula DB in all servers to evolve the system state.

To preserve a consistent view of the system across servers, modifications to system state are performed through a special node, the *leader*. The servers in the OpenNebula cluster elects a single node to be the *leader*. The *leader* periodically sends heartbeats to the other servers, the *followers*, to keep its leadership. If a *leader* fails to send the heartbeat, *followers* promote to *candidates* and start a new election.

Whenever the system is modified (e.g. a new VM is added to the system), the *leader* updates the log and replicates the entry in a majority of *followers* before actually writing it to the database. The latency of DB operations are thus increased, but the system state is safely replicated, and the cluster can continue its operation in case of node failure.

In OpenNebula, read-only operations can be performed through any oned server in the cluster; this means that reads can be arbitrarily stale but generally within the round-trip time of the network.

### 2.2.2 Requirements and Architecture

The recommended deployment size is either 3 or 5 servers, which provides a fault-tolerance for 1 or 2 server failures, respectively. You can add, replace or remove servers once the cluster is up and running.

Every HA cluster requires:

- Odd number of servers (3 is recommended).

- Recommended identical servers capacity.

- Same software configuration of the servers (the sole difference would be the `SERVER_ID` field in `/etc/one/oned.conf`).

- Working database connection of the same type, MySQL is recommended.

- All the servers must share the credentials.

- Floating IP which will be assigned to the *leader*.

- Shared filesystem.

The servers should be configured in the following way:

- Sunstone (with or without Apache/Passenger) running on all the nodes.

- Shared datastores must be mounted on all the nodes.

### 2.2.3 Bootstrapping the HA cluster

This section shows on examples all steps required to deploy the HA Cluster.

> **Warning:** To maintain a healthy cluster during the procedure of adding servers to the clusters, make sure you add **only** one server at a time

---

**Important:** In the following, each configuration step starts with (initial) **Leader** or (future) **Follower** to indicate the server where the step must be performed.

---

### Configuration of the initial leader

We start with the first server, to perform the initial system bootstrapping.

- **Leader**: Start OpenNebula

- **Leader**: Add the server itself to the zone:

```
$ onezone list
C   ID NAME                     ENDPOINT
*    0 OpenNebula                http://localhost:2633/RPC2

# We are working on Zone 0

$ onezone server-add 0 --name server-0 --rpc http://192.168.150.1:2633/RPC2

# It's now available in the zone:

$ onezone show 0
ZONE 0 INFORMATION
ID                : 0
NAME              : OpenNebula


ZONE SERVERS
ID NAME             ENDPOINT
 0 server-0         http://192.168.150.1:2633/RPC2

HA & FEDERATION SYNC STATUS
ID NAME             STATE       TERM        INDEX       COMMIT      VOTE  FED_INDEX
 0 server-0         solo        0           -1          0           -1    -1

ZONE TEMPLATE
ENDPOINT="http://localhost:2633/RPC2"
```

---

**Important:** Floating IP should be used for **zone endpoints** and cluster private addresses for the zone **server endpoints**.

---

- **Leader**: Stop OpenNebula service and update `SERVER_ID` in `/etc/one/oned.conf`

```
FEDERATION = [
    MODE          = "STANDALONE",
    ZONE_ID       = 0,
    SERVER_ID     = 0, # changed from -1 to 0 (as 0 is the server id)
    MASTER_ONED   = ""
]
```

- **Leader**: [Optional] Enable the RAFT Hooks in `/etc/one/oned.conf`. This will add a floating IP to the system.

```
# Executed when a server transits from follower->leader
RAFT_LEADER_HOOK = [
    COMMAND = "raft/vip.sh",
    ARGUMENTS = "leader eth0 10.3.3.2/24"
]
```

```
# Executed when a server transits from leader->follower
RAFT_FOLLOWER_HOOK = [
    COMMAND = "raft/vip.sh",
    ARGUMENTS = "follower eth0 10.3.3.2/24"
]
```

- **Leader**: Start OpenNebula.

- **Leader**: Check the zone, the server is now the leader and has the floating IP:

```
$ onezone show 0
ZONE 0 INFORMATION
ID                : 0
NAME              : OpenNebula


ZONE SERVERS
ID NAME            ENDPOINT
 0 server-0        http://192.168.150.1:2633/RPC2


HA & FEDERATION SYNC STATUS
ID NAME            STATE       TERM        INDEX       COMMIT      VOTE  FED_INDEX
 0 server-0        leader      1           3           3           -1    -1


ZONE TEMPLATE
ENDPOINT="http://localhost:2633/RPC2"
$ ip -o a sh eth0|grep 10.3.3.2/24
2: eth0    inet 10.3.3.2/24 scope global secondary eth0\        valid_lft forever␣
→preferred_lft forever
```

### Adding more servers

**Warning:** This procedure will discard the OpenNebula database in the server you are adding and substitute it with the database of the initial leader.

**Warning:** Add only one host at a time. Repeat this process for every server you want to add.

- **Leader**: Create a DB backup in the initial leader and distribute it to the new server, along with the files in `/var/lib/one/.one/`:

```
$ onedb backup -u oneadmin -p oneadmin -d opennebula
MySQL dump stored in /var/lib/one/mysql_localhost_opennebula_2017-6-1_11:52:47.sql
Use 'onedb restore' or restore the DB using the mysql command:
mysql -u user -h server -P port db_name < backup_file

# Copy it to the other servers
$ scp /var/lib/one/mysql_localhost_opennebula_2017-6-1_11:52:47.sql <ip>:/tmp

# Copy the .one directory (make sure you preseve the owner: oneadmin)
$ ssh <ip> rm -rf /var/lib/one/.one
$ scp -r /var/lib/one/.one/ <ip>:/var/lib/one/
```

- **Follower**: Stop OpenNebula on the new server if it is running.

- **Follower**: Restore the database backup on the new server.

```
$ onedb restore -f -u oneadmin -p oneadmin -d opennebula /tmp/mysql_localhost_
↪opennebula_2017-6-1_11:52:47.sql
MySQL DB opennebula at localhost restored.
```

- **Leader**: Add the new server to OpenNebula (in the initial leader), and note the server id.

```
$ onezone server-add 0 --name server-1 --rpc http://192.168.150.2:2633/RPC2
```

- **Leader**: Check the zone, the new server is in error state, since OpenNebula on the new server is still not running. Make a note of the server id, in this case it is 1.

```
$ onezone show 0
ZONE 0 INFORMATION
ID                : 0
NAME              : OpenNebula


ZONE SERVERS
ID NAME            ENDPOINT
 0 server-0        http://192.168.150.1:2633/RPC2
 1 server-1        http://192.168.150.2:2633/RPC2

HA & FEDERATION SYNC STATUS
ID NAME            STATE      TERM      INDEX     COMMIT    VOTE  FED_INDEX
 0 server-0        leader     1         19        19        -1    -1
 1 server-1        error      -         -         -         -

ZONE TEMPLATE
ENDPOINT="http://localhost:2633/RPC2"
```

- **Follower**: Edit /etc/one/oned.conf on the new server to set the SERVER_ID for the new server. Make sure to enable the hooks as in the initial leader's configuration.

- **Follower**: Start OpenNebula service.

- **Leader**: Run *onezone show 0* to make sure that the new server is in follower state.

```
$ onezone show 0
ZONE 0 INFORMATION
ID                : 0
NAME              : OpenNebula


ZONE SERVERS
ID NAME            ENDPOINT
 0 server-0        http://192.168.150.1:2633/RPC2
 1 server-1        http://192.168.150.2:2633/RPC2

HA & FEDERATION SYNC STATUS
ID NAME            STATE      TERM      INDEX     COMMIT    VOTE  FED_INDEX
 0 server-0        leader     1         21        19        -1    -1
 1 server-1        follower   1         16        16        -1    -1

ZONE TEMPLATE
ENDPOINT="http://localhost:2633/RPC2"
```

---

**Note:** It may happen the **TERM/INDEX/COMMIT** does not match (like above). This is not important right now; it will sync automatically when the database is changed.

---

Repeat this section to add new servers. Make sure that you only add servers when the cluster is in a healthy state. That means there is 1 leader and the rest are in follower state. If there is one server in error state, fix it before proceeding.

## 2.2.4 Checking Cluster Health

Execute `onezone show <id>` to see if any of the servers are in error state. If they are in error state, check `/var/log/one/oned.log` in both the current leader (if any) and in the host that is in error state. All Raft messages will be logged in that file.

If there is no leader in the cluster please review `/var/log/one/oned.log` to make sure that there are no errors taking place.

## 2.2.5 Adding and Removing Servers

In order to add servers you need to use this command:

```
$ onezone server-add
Command server-add requires one parameter to run
## USAGE
server-add <zoneid>
        Add an OpenNebula server to this zone.
        valid options: server_name, server_rpc

## OPTIONS
    -n, --name                Zone server name
    -r, --rpc                 Zone server RPC endpoint
    -v, --verbose             Verbose mode
    -h, --help                Show this message
    -V, --version             Show version and copyright information
    --user name               User name used to connect to OpenNebula
    --password password       Password to authenticate with OpenNebula
    --endpoint endpoint       URL of OpenNebula xmlrpc frontend
```

Make sure that there is one leader (by running `onezone show <id>`), otherwise it will not work.

To remove a server, use the command:

```
$ onezone server-del
Command server-del requires 2 parameters to run.
## USAGE
server-del <zoneid> <serverid>
        Delete an OpenNebula server from this zone.

## OPTIONS
    -v, --verbose             Verbose mode
    -h, --help                Show this message
    -V, --version             Show version and copyright information
    --user name               User name used to connect to OpenNebula
    --password password       Password to authenticate with OpenNebula
    --endpoint endpoint       URL of OpenNebula xmlrpc frontend
```

The whole procedure is documented *above*.

---

## 2.2.6 Recovering servers

When a follower is down for some time it may fall out of the recovery window, i.e. the log may not include all the records needed to bring it up-to-date. In order to recover this server you need to:

- **Leader**: Create a DB backup and copy it to the failed follower. Note that you cannot reuse a previous backup.

- **Follower**: Stop OpenNebula if running.

- **Follower**: Restore the DB backup from the leader.

- **Follower**: Start OpenNebula.

- **Leader**: Reset the failing follower with:

```
$ onezone server-reset <zone_id> <server_id_of_failed_follower>
```

## 2.2.7 Shared data between HA nodes

HA deployment requires the filesystem view of most datastores (by default in `/var/lib/one/datastores/`) to be same on all front-ends. It is necessary to setup a shared filesystem over the datastore directories. This document does not cover configuration and deployment of the shared filesystem; it is left completely up to the cloud administrator.

OpenNebula stores virtual machine logs inside `/var/log/one/` as files named `${VMID}.log`. It is not recommended to share the whole log directory between the front-ends as there are also other OpenNebula logs which would be randomly overwritten. It is up to the cloud administrator to periodically backup the virtual machine logs on cluster leader and on fail-over to restore from the backup on a new leader (e.g. as part of the raft hook).

## 2.2.8 Sunstone

There are several types of the Sunstone deployment in HA environment. The basic one is Sunstone running on each OpenNebula front-end node configured with the local OpenNebula. Only one server, the leader with floating IP, is used by the clients.

It is possible to configure a load balancer (e.g. HAProxy, Pound, Apache or Nginx) over the front-ends to spread the load (read operations) among the nodes. In this case, the **Memcached** and shared `/var/tmp/` may be required, please see Configuring Sunstone for Large Deployments.

To easy scale out beyond the total number of core OpenNebula daemons, Sunstone can be running on separate machines. They should talk to the cluster floating IP (see `:one_xmlprc:` in `sunstone-server.conf`) and may also require **Memcached** and shared `/var/tmp/` between Sunstone and front-end nodes. Please check Configuring Sunstone for Large Deployments.

## 2.2.9 Raft Configuration Attributes

The Raft algorithm can be tuned by several parameters in the configuration file `/etc/one/oned.conf`. Following options are available:

| Raft: Algorithm Attributes | |
|---|---|
| `LIMIT_PURGE` | Number of DB log records that will be deleted on each purge. |
| `LOG_RETENTION` | Number of DB log records kept, it determines the synchronization window across servers and extra storage space needed. |
| `LOG_PURGE_TIMEOUT` | How often applied records are purged according the log retention value. (in seconds). |
| `ELECTION_TIMEOUT_MS` | Timeout to start a election process if no heartbeat or log is received from leader. |
| `BROADCAST_TIMEOUT_MS` | How often heartbeats are sent to followers. |
| `XMLRPC_TIMEOUT_MS` | To timeout raft related API calls. To set an infinite timeout set this value to 0. |

> **Warning:** Any change in these parameters can lead to the unexpected behavior during the fail-over and result in whole cluster malfunction. After any configuration change, always check the crash scenarios for the correct behavior.

### 2.2.10 Compatibility with the earlier HA

In OpenNebula <= 5.2, HA was configured using a classical active-passive approach, using Pacemaker and Corosync. While this still works for OpenNebula > 5.2, it is not the recommended way to set up a cluster. However, it is fine if you want to continue using that HA coming from earlier versions.

This is documented here: Front-end HA Setup.

## 2.3 Virtual Machines High Availability

This section's objective is to provide information in order to prepare for failures in the Virtual Machines or Hosts, and recover from them. These failures are categorized depending on whether they come from the physical infrastructure (Host failures) or from the virtualized infrastructure (VM crashes). In both scenarios, OpenNebula provides a cost-effective failover solution to minimize downtime from server and OS failures.

### 2.3.1 Host Failures

When OpenNebula detects that a host is down, a hook can be triggered to deal with the situation. OpenNebula comes with a script out-of-the-box that can act as a hook to be triggered when a host enters the ERROR state. This can very useful to limit the downtime of a service due to a hardware failure, since it can redeploy the VMs on another host.

Let's see how to configure `/etc/one/oned.conf` to set up this Host hook, to be triggered in the ERROR state. The following should be uncommented in the mentioned configuration file:

```
#-------------------------------------------------------------------------
HOST_HOOK = [
    name      = "error",
    on        = "ERROR",
    command   = "ft/host_error.rb",
    arguments = "$ID -m -p 5",
    remote    = "no" ]
#-------------------------------------------------------------------------
```

We are defining a host hook, named `error`, that will execute the script `ft/host_error.rb` locally with the following arguments:

| Argument | Description |
|---|---|
| **Host ID** | ID of the host containing the VMs to treat. It is compulsory and better left to **$ID**, that will be automatically filled by OpenNebula with the Host ID of the host that went down. |
| **Action** | This defines the action to be performed upon the VMs that were running in the host that went down. This can be: <ul><li>**-m** migrate VMs to another host. Only for images in shared storage</li><li>**-r** delete+recreate VMs running in the host. State will be lost.</li><li>**-d** delete VMs running in the host</li></ul> |
| **ForceSuspended** | [-f] force resubmission of suspended VMs |
| **AvoidTransient** | [-p <n>] avoid resubmission if host comes back after <n> monitoring cycles |

More information on hooks here.

> **Warning:** Note that spurious network errors may lead to a VM started twice in different hosts and possibly contend on shared resources. The previous script needs to fence the error host to prevent split brain VMs. You may use any fencing mechanism for the host and invoke it within the error hook.

## 2.3.2 Virtual Machine Failures

The overall state of a virtual machine in a failure condition will show as `failure` (or `fail` in the CLI). To find out the specific failure situation you need to check the `LCM_STATE` of the VM in the VM info tab (or `onevm show` in the CLI.). Moreover, a VM can be stuck in a transition (e.g. boot or save) because of a host or network failure. Typically these operations will eventually timeout and lead to a VM failure state.

The administrator has the ability to force a recovery action from Sunstone or from the CLI, with the `onevm recover` command. This command has the following options:

- `--success`: If the operation has been confirmed to succeed. For example, the administrator can see the VM properly running in the hypervisor, but the driver failed to inform OpenNebula of the successful boot.

- `--failure`: This will have the same effect as a driver reporting a failure. It is intended for VMs that get stuck in transient states. As an example, if a storage problem occurs and the administrator knows that a VM stuck in `prolog` is not going to finish the pending transfer, this action will manually move the VM to `prolog_failure`.

- `--retry`: To retry the previously failed action. Can be used, for instance, in case a VM is in `boot_failure` because the hypervisor crashed. The administrator can tell OpenNebula to retry the boot after the hypervisor is started again.

- `--retry --interactive`: In some scenarios where the failure was caused by an error in the Transfer Manager actions, each action can be rerun and debugged until it works. Once the commands are successful, a `success` should be sent. See the specific section below for more details.

- `--delete`: No recover action possible, delete the VM. This is equivalent to the deprecated OpenNebula < 5.0 command: `onevm delete`.

- `--recreate`: No recover action possible, delete and recreate the VM. This is equivalent to the deprecated OpenNebula < 5.0 command: `onevm delete --recreate`.

Note also that OpenNebula will try to automatically recover some failure situations using the monitor information. A specific example is that a VM in the `boot_failure` state will become `running` if the monitoring reports that the VM was found running in the hypervisor.

### Hypervisor Problems

The following list details failures states caused by errors related to the hypervisor.

- `BOOT_FAILURE`, The VM failed to boot but all the files needed by the VM are already in the host. Check the hypervisor logs to find out the problem, and once fixed recover the VM with the retry option.

- `BOOT_MIGRATE_FAILURE`, same as above but during a migration. Check the target hypervisor and retry the operation.

- `BOOT_UNDEPLOY_FAILURE`, same as above but during a resume after an undeploy. Check the target hypervisor and retry the operation.

- `BOOT_STOPPED_FAILURE`, same as above but during a resume after a stop. Check the target hypervisor and retry the operation.

### Transfer Manager / Storage Problems

The following list details failure states caused by errors in the Transfer Manager driver. These states can be recovered by checking the vm.log and looking for the specific error (disk space, permissions, mis-configured datastore, etc). You can execute `--retry` to relaunch the Transfer Manager actions after fixing the problem (freeing disk space, etc). You can execute `--retry --interactive` to launch a Transfer Manager Interactive Debug environment that will allow you to: (1) see all the TM actions in detail (2) relaunch each action until its successful (3) skip TM actions.

- `PROLOG_FAILURE`, there was a problem setting up the disk images needed by the VM.

- `PROLOG_MIGRATE_FAILURE`, problem setting up the disks in the target host.

- `EPILOG_FAILURE`, there was a problem processing the disk images (may be discard or save) after the VM execution.

- `EPILOG_STOP_FAILURE`, there was a problem moving the disk images after a stop.

- `EPILOG_UNDEPLOY_FAILURE`, there was a problem moving the disk images after an undeploy.

- `PROLOG_MIGRATE_POWEROFF_FAILURE`, problem restoring the disk images after a migration in a poweroff state.

- `PROLOG_MIGRATE_SUSPEND_FAILURE`, problem restoring the disk images after a migration in a suspend state.

- `PROLOG_RESUME_FAILURE`, problem restoring the disk images after a stop.

- `PROLOG_UNDEPLOY_FAILURE`, problem restoring the disk images after an undeploy.

Example of a Transfer Manager Interactive Debug environment (`onevm recover <id> --retry --interactive`):

```
$ onevm show 2|grep LCM_STATE
LCM_STATE            : PROLOG_UNDEPLOY_FAILURE

$ onevm recover 2 --retry --interactive
TM Debug Interactive Environment.
```
(continues on next page)

```
TM Action list:
(1) MV shared haddock:/var/lib/one//datastores/0/2/disk.0 localhost:/var/lib/one//
↪datastores/0/2/disk.0 2 1
(2) MV shared haddock:/var/lib/one//datastores/0/2 localhost:/var/lib/one//datastores/
↪0/2 2 0

Current action (1):
MV shared haddock:/var/lib/one//datastores/0/2/disk.0 localhost:/var/lib/one//
↪datastores/0/2/disk.0 2 1

Choose action:
(r) Run action
(n) Skip to next action
(a) Show all actions
(q) Quit
> r

LOG I  Command execution fail: /var/lib/one/remotes/tm/shared/mv haddock:/var/lib/one/
↪/datastores/0/2/disk.0 localhost:/var/lib/one//datastores/0/2/disk.0 2 1
LOG I  ExitCode: 1

FAILURE. Repeat command.

Current action (1):
MV shared haddock:/var/lib/one//datastores/0/2/disk.0 localhost:/var/lib/one//
↪datastores/0/2/disk.0 2 1

Choose action:
(r) Run action
(n) Skip to next action
(a) Show all actions
(q) Quit
> # FIX THE PROBLEM...

> r

SUCCESS

Current action (2):
MV shared haddock:/var/lib/one//datastores/0/2 localhost:/var/lib/one//datastores/0/2
↪2 0

Choose action:
(r) Run action
(n) Skip to next action
(a) Show all actions
(q) Quit
> r

SUCCESS

If all the TM actions have been successful and you want to
recover the Virtual Machine to the RUNNING state execute this command:
$ onevm recover 2 --success

$ onevm recover 2 --success
```

```
$ onevm show 2|grep LCM_STATE
LCM_STATE            : RUNNING
```

# DATA CENTER FEDERATION

## 3.1 Overview

Several OpenNebula instances can be configured as a **Federation**. Each instance of the Federation is called a **Zone**, and they are configured as one master and several slaves. Any other federation configurations (e.g. deeper master-slave hierarchy) aren't supported.

An OpenNebula Federation is a tightly coupled integration. All the Zones will share the same user accounts, groups, and permissions configuration. Moreover, you can define access policies federation-wide so users can be restricted to certain Zones, or to specific Clusters inside a Zone.

For the end users, a Federation allows them to use the resources no matter where they are. The integration is seamless, meaning that a user logged into the Sunstone web interface of a Zone just need to select the zone where she wants to work.

### 3.1.1 Architecture

The master Zone is responsible for updating the federated information and replicating the updates in the slaves. The federated information shared across zones includes users, groups, VDCs, ACL rules, marketplace, marketplace apps, and zones.

The slave Zones have read-only access to the local copy of the federated information. Write operations on the slaves are redirected to the master Zone. Note that you may suffer from stale reads while the data is replicating from the master to the slaves. However, this approach ensures sequential consistency across zones (each zone will replicate operations in the same order) without any impact on the speed of read-only actions.

The federated information replication is implemented with a log that includes a sequence of SQL statements applied to the shared tables. This log is replicated and applied in each Zone database. This replication model tolerates faulty connections, and even zone crashes without impacting the federation.

The administrators can share appliances across Zones deploying a private *OpenNebula Marketplace*.

### 3.1.2 Other Services

Although a single Sunstone server can connect to different Zones, all the other OpenNebula services will only work with the local Zone resources. This includes the Scheduler, the *Public Cloud Servers*, *OneFlow*, and *OneGate*.

### 3.1.3 How Should I Read This Chapter

Before reading this chapter make sure you have read the Deployment Guide.

Read the *Federation Configuration* section to learn how to setup a federation, and the *Federation Management* section to learn how to manage zones in OpenNebula.

After reading this chapter, you can continue configuring more *Advanced Components*.

### 3.1.4 Hypervisor Compatibility

This chapter applies both to KVM and vCenter.

## 3.2 OpenNebula Federation Configuration

This section will explain how to configure two (or more) OpenNebula zones to work as federation master and slave. The process described here can be applied to new installations or existing OpenNebula instances.

OpenNebula master zone server replicates database changes on slaves using a federated log. The log contains the SQL commands which should be applied in all zones.

---

**Important:** In the following, each configuration step starts with **Master** or **Slave** to indicate the server where the step must be performed.

---

**Important:** Master and slave servers need to talk to each other through their XML-RPC API. You may need to update the `LISTEN_ADDRESS`, and or `PORT` in `/etc/one/oned.conf` or any firewall rule blocking this communication. Note that by default this traffic is not secured, so if you are using public links you need to secure the communication.

---

**Important:** The federation can be setup with MySQL or SQLite backends, but you cannot mix them across zones. MySQL is recommended for production deployments.

---

### 3.2.1 Step 1. Configure the OpenNebula Federation Master Zone

Start by picking an OpenNebula to act as master of the federation. The master OpenNebula will be responsible for updating shared information across zones and replicating the updates to the slaves. You may start with an existing installation or with a new one (see the installation guide).

---

**Note:** When installing a new master from scratch be sure to start it at least once to properly bootstrap the database.

---

- **Master**: Edit the master zone endpoint. This can be done via Sunstone, or with the onezone command. Write down this endpoint to use it later when configuring the slaves.

```
$ onezone update 0
ENDPOINT = http://<master-ip>:2633/RPC2
```

---

**Note:** In the HA setup, the master-ip should be set to the **floating** IP address, see *the HA installation guide* for more details. In single server zones, just use the IP of the server.

---

- **Master**: Update `/etc/one/oned.conf` to change the mode to **master**.

---

```
FEDERATION = [
    MODE    = "MASTER",
    ZONE_ID = 0
]
```

- **Master**: Restart the OpenNebula.

You are now ready to add slave zones.

### 3.2.2 Step 2. Adding a New Federation Slave Zone

- **Slave**: Install OpenNebula on the slave as usual following the installation guide. Start OpenNebula at least once to bootstrap the zone database.

- **Slave**: Stop OpenNebula.

- **Master**: Create a zone for the slave, and write down the new Zone ID. This can be done via Sunstone, or with the onezone command.

```
$ vim /tmp/zone.tmpl
NAME     = slave-name
ENDPOINT = http://<slave-zone-ip>:2633/RPC2

$ onezone create /tmp/zone.tmpl
ID: 100

$ onezone list
   ID NAME
    0 OpenNebula
  100 slave-name
```

---

**Note:** In HA setups use the **floating** IP address for the slave-zone-ip, in single server zones just use the IP of the server.

---

- **Master**: Make a snapshot of the federated tables with the following command:

```
$ onedb backup --federated -s /var/lib/one/one.db
Sqlite database backup of federated tables stored in /var/lib/one/one.db_federated_
↪2017-6-15_8:52:51.bck
Use 'onedb restore' to restore the DB.
```

---

**Note:** This example shows how to make a database snapshot with SQLite. For MySQL just change the -s option with the corresponding MySQL options: -u <username> -p <password> -d <database_name>. For SQLite, you need to stop OpenNebula before taking the DB snapshot. This is not required for MySQL.

---

- **Master**: Copy the database snapshot to the slave.

- **Master**: Copy **only selected files** from the directory /var/lib/one/.one to the slave. This directory and its content must have **oneadmin as owner**. Replace only these files:

```
$ ls -1 /var/lib/one/.one
ec2_auth
one_auth
```

(continues on next page)

```
oneflow_auth
onegate_auth
sunstone_auth
```

- **Slave**: Update `/etc/one/oned.conf` to change the mode to **slave**, set the master's URL and the `ZONE_ID` obtained when the zone was created on master:

```
FEDERATION = [
    MODE        = "SLAVE",
    ZONE_ID     = 100,
    MASTER_ONED = "http://<master-ip>:2633/RPC2"
]
```

- **Slave**: Restore the database snapshot:

```
$ onedb restore --federated -s /var/lib/one/one.db /var/lib/one/one.db_federated_2017-
↪6-14_16:0:36.bck
Sqlite database backup restored in one.db
```

- **Slave**: Start OpenNebula.

The zone should be now configured and ready to use.

### 3.2.3 Step 3 [Optional]. Adding HA to a Federation Slave Zone

Now you can start adding more servers to the slave zone to provide it with HA capabilities. The procedure is the same as the one described for stand-alone zones in *the HA installation guide*. In this case, the replication works in a multi-tier fashion. The master replicates a database change to one of the zone servers. Then this server replicates the change across the zone servers.

---

**Important:** It is important to double check that the federation is working before adding HA servers to the zone, as you will be updating the zone metadata which is a federated information.

---

### 3.2.4 Importing Existing OpenNebula Zones

There is no automatic procedure to import existing users and groups into a running federation. However, you can preserve everything else like datastores, VMs, networks. . .

- **Slave**: Backup details of users, groups, and VDCs you want to recreate in the federated environment.
- **Slave**: Stop OpenNebula. If the zone was running an HA cluster, stop all servers and pick one of them to add the zone to the federation. Put this server in solo mode by setting `SERVER_ID` to `-1` in `/etc/one/oned.conf`.
- **Master, Slave**: Follow the procedure described in Step 2 to add a new zone.
- **Slave**: Recreate any user, group or VDC you need to preserve in the federated environment.

The Zone is now ready to use. If you want to add more HA servers, follow the standard procedure.

### 3.2.5 Updating a Federation

OpenNebula database has two different version numbers:

---

- federated (shared) tables version,

- local tables version.

---

**Important:** To federate OpenNebula zones, they must run the same version of the federated tables (which are pretty stable).

---

Upgrades to a version that does not increase the federated version can be done asynchronously in each zone. However, an update in the shared table version requires a coordinated update of all zones.

### 3.2.6 Administration account configuration

A Federation will have a unique oneadmin account. This is required to perform API calls across zones. It is recommended to not use this account directly in a production environment, and create an account in the 'oneadmin' group for each Zone administrator.

When additional access restrictions are needed, the Federation Administrator can create a special administrative group with total permissions for one zone only.

## 3.3 OpenNebula Federation Usage

A user will have access to all the Zones where at least one of her groups has VDC resources in. This access be can done through Sunstone or the CLI.

### 3.3.1 Using a Zone Through Sunstone

In the upper right corner of the Sunstone page, users will see a globe icon next to the name of the Zone you are currently using. If the user clicks on that, she will get a dropdown with all the Zones she has access to. Clicking on any of the Zones in the dropdown will get the user to that Zone.

What's happening behind the scenes is that the Sunstone server you are using is redirecting its requests to the OpenNebula oned process present in the other Zone. In the example above, if the user clicks on ZoneB, Sunstone contacts the OpenNebula listening at `http://zoneb.opennebula.front-end.server:2633/RPC2`.

> **Warning:** Uploading an image functionality is limited to the zone where the Sunstone instance the user is connecting to, even if it can switch to other federated zones.

### 3.3.2 Using a Zone Through CLI

Users can switch Zones through the command line using the onezone command. The following session can be examined to understand the Zone management through the CLI.

```
$ onezone list
C    ID NAME                      ENDPOINT
*     0 OpenNebula                http://localhost:2633/RPC2
    104 ZoneB                     http://ultron.c12g.com:2634/RPC2
```

We can see in the above command output that the user has access to both "OpenNebula" and "ZoneB", and it is currently in the "OpenNebula" Zone. The active Zone can be changed using the 'set' command of onezone:

```
$ onezone set 104
Endpoint changed to "http://ultron.c12g.com:2634/RPC2" in /home/<username>/.one/one_
↪endpoint

$ onezone list
C    ID NAME                      ENDPOINT
      0 OpenNebula                http://localhost:2633/RPC2
*   104 ZoneB                     http://ultron.c12g.com:2634/RPC2
```

All the subsequent CLI commands executed would connect to the OpenNebula listening at `http://zoneb.opennebula.front-end.server:2633/RPC2`.
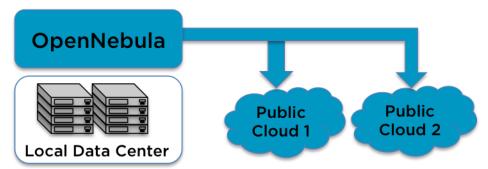
# CLOUD BURSTING

## 4.1 Overview

Cloud bursting is a model in which the local resources of a Private Cloud are combined with resources from remote Cloud providers. The remote provider could be a commercial Cloud service, such as Amazon EC2, Microsoft Azure or even OpenNebula based clouds. Such support for cloud bursting enables highly scalable hosting environments.



OpenNebula's approach to cloud bursting is based on the transparency to both end users and cloud administrators to use and maintain the cloud bursting functionality. The **transparency to cloud administrators** comes from the fact that a AWS EC2 region or an Azure location is modeled as any other host (albeit of potentially a much bigger capacity), so the scheduler can place VMs in the external cloud as it will do in any other local host.

On the other hand, the **transparency to end users** is offered through the hybrid template functionality: the same VM template in OpenNebula can describe the VM if it is deployed locally and also if it gets deployed in EC2, or Azure. Therefore users just have to instantiate the template, and OpenNebula will transparently choose if that is executed locally or remotely.

### 4.1.1 How Should I Read This Chapter

You should be reading this Chapter as part of the *Advanced Components Guide* review of that OpenNebula advanced functionality that you are interested in enabling and configuring.

Within this Chapter you can find a guide to configure and use the *Amazon EC2 driver*, the *Azure driver* and the *OpenNebula driver*. Additionally, there is support of IBM SoftLayer available as add-on (ie, not contained in the OpenNebula main distribution) here.

Once the cloud architecture has been designed the next step would be to learn how to install the OpenNebula front-end.

### 4.1.2 Hypervisor Compatibility

All the Sections in this Chapter applies to both KVM and vCentre based OpenNebula clouds.

## 4.2 Amazon EC2 Driver

### 4.2.1 Considerations & Limitations

You should take into account the following technical considerations when using the EC2 cloud with OpenNebula:

- There is no direct access to the hypervisor, so it cannot be monitored (we don't know where the VM is running on the EC2 cloud).
- The usual OpenNebula functionality for snapshotting, hot-plugging, or migration is not available with EC2.
- By default OpenNebula will always launch m1.small instances, unless otherwise specified.
- Monitoring of VMs in EC2 is done through CloudWatch. Only information related to the consumption of CPU and Networking (both inbound and outbound) is collected, since CloudWatch does not offer information of guest memory consumption.

Please refer to the EC2 documentation to obtain more information about Amazon instance types and image management:

- General information of instances

### 4.2.2 Prerequisites

- You must have a working account for AWS and signup for EC2 and S3 services.
- The aws-sdk ruby gem needs to be installed, version 2.5+. This gem is automatically installed as part of the installation process.

### 4.2.3 OpenNebula Configuration

Uncomment the EC2 IM and VMM drivers from `/etc/one/oned.conf` file in order to use the driver.

```
IM_MAD = [
    name       = "ec2",
    executable = "one_im_sh",
    arguments  = "-c -t 1 -r 0 ec2" ]

VM_MAD = [
    name       = "ec2",
    executable = "one_vmm_sh",
    arguments  = "-t 15 -r 0 ec2",
    type       = "xml" ]
```

Driver flags are the same as other drivers:

| FLAG | SETs |
|------|------|
| -t | Number of threads |
| -r | Number of retries |

First of all we take a look over our ec2_driver_conf file located in `/etc/one/ec2_driver.conf`:

```
proxy_uri:
state_wait_timeout_seconds: 300
instance_types:
    c1.medium:
    cpu: 2
    memory: 1.7
    ...
```

You can define an http proxy if the OpenNebula Frontend does not have access to the internet:

```
proxy_uri: http://...
```

Also, you can modify in the same file the default 300 seconds timeout that is waited for the VM to be in the EC2 running state in case you also want to attach to the instance a elastic ip:

```
state_wait_timeout_seconds: 300
```

> **Warning:** `instance_types` section shows us the machines that AWS is able to provide, the ec2 driver will retrieve this kind of information so it's better to not change it unless you are aware of your actions.

> **Warning:** If you were using OpenNebula before 5.4 you may have noticed that there are not AWS credentials in configuration file anymore, this is due security reasons. In 5.4 there is a new secure credentials storage for AWS so you do not need to store sensitive credential data inside your disk. OpenNebula daemon stores the data in an encrypted format.

After OpenNebula is restarted, create a new Host with AWS credentials that uses the ec2 drivers:

```
$ onehost create ec2 -t ec2 --im ec2 --vm ec2
```

> **Note:** `-t` is needed to specify what type of remote provider host we want to set up, if you've followed all the instruction properly your default editor should show in your screen asking for the credentials and other mandatory data that will allow you to communicate with AWS.

Once you have opened your editor you can look for additional help at the top of your screen, you have more information in *EC2 Specific Template Attributes* section. The basic three variables you have to set are: `EC2_ACCESS`, `EC2_SECRET` and `REGION_NAME`.

This can also be done creating a template file than can be used with the creation command:

```
$ echo 'EC2_ACCESS = "xXxXXxx"' >  ec2host.tpl
$ echo 'EC2_SECRET = "xXXxxXx"' >> ec2host.tpl
$ echo 'REGION_NAME= "xXXxxXx"' >> ec2host.tpl
$ onehost create ec2 -t ec2 ec2host.tpl --im ec2 --vm ec2
```

### 4.2.4 EC2 Specific Template Attributes

In order to deploy an instance in EC2 through OpenNebula you must include an EC2 section in the virtual machine template. This is an example of a virtual machine template that can be deployed in our local resources or in EC2.

---

```
CPU      = 0.5
MEMORY   = 128

# KVM template machine, this will be use when submitting this VM to local resources
DISK     = [ IMAGE_ID = 3 ]
NIC      = [ NETWORK_ID = 7 ]

# PUBLIC_CLOUD template, this will be use wen submitting this VM to EC2
PUBLIC_CLOUD = [ TYPE="EC2",
                 AMI="ami-00bafcb5",
                 KEYPAIR="gsg-keypair",
                 INSTANCETYPE=m1.small]

#Add this if you want to use only EC2 cloud
#SCHED_REQUIREMENTS = 'HOSTNAME = "ec2"'
```

Check an exhaustive list of attributes in the Virtual Machine Definition File Reference Section.

Default values for all these attributes can be defined in the `/etc/one/ec2_driver.default` file.

```
<!--
 Default configuration attributes for the EC2 driver
 (all domains will use these values as defaults)
 Valid attributes are: AKI AMI CLIENTTOKEN INSTANCETYPE KEYPAIR LICENSEPOOL
    PLACEMENTGROUP PRIVATEIP RAMDISK SUBNETID TENANCY USERDATA SECURITYGROUPS
    AVAILABILITYZONE EBS_OPTIMIZED ELASTICIP TAGS
 Use XML syntax to specify defaults, note elements are UPCASE
 Example:
 <TEMPLATE>
   <PUBLIC_CLOUD>
     <KEYPAIR>gsg-keypair</KEYPAIR>
     <INSTANCETYPE>m1.small</INSTANCETYPE>
   </PUBLIC_CLOUD>
 </TEMPLATE>
-->

<TEMPLATE>
  <PUBLIC_CLOUD>
    <INSTANCETYPE>m1.small</INSTANCETYPE>
  </PUBLIC_CLOUD>
</TEMPLATE>
```

---

**Note:** The PUBLIC_CLOUD sections allow for substitutions from template and virtual network variables, the same way as the CONTEXT section allows.

---

These values can furthermore be asked to the user using user inputs. A common scenario is to delegate the User Data to the end user. For that, a new User Input named USERDATA can be created of text64 (the User Data needs to be encoded on base64) and a placeholder added to the PUBLIC_CLOUD section:

```
PUBLIC_CLOUD = [ TYPE="EC2",
                 AMI="ami-00bafcb5",
                 KEYPAIR="gsg-keypair",
                 INSTANCETYPE=m1.small,
                 USERDATA="$USERDATA"]
```

### Auth Attributes

After successfully executing onehost create with -t option, your default editor will open. An example follows of how you can complete this area:

```
EC2_ACCESS = "this_is_my_ec2_access_key_identifier"
EC2_SECRET = "this_is_my_ec2_secret_key"
REGION_NAME = "us-east-1"
CAPACITY = [
    M1_SMALL = "3",
    M1_LARGE = "1" ]
```

The first two attributes have the authentication info required by AWS:

- **EC2_ACCESS**: Amazon AWS Access Key

- **EC2_SECRET**: Amazon AWS Secret Access Key

This information will be encrypted as soon as the host is created. In the host template the values of the `EC2_ACCESS` and `EC2_SECRET` attributes will be encrypted.

- **REGION_NAME**: it's the name of AWS region that your account uses to deploy machines.

In the example the region is set to *us-east-1*, you can get this information at the EC2 web console.

- **CAPACITY**: This attribute sets the size and number of EC2 machines that your OpenNebula host will handle, you can see `instance_types` section in `ec2_driver.conf` file to know the supported names. Dot ('.') is not permitted, you have to change it to underscores (_) and capitalize the names (`m1.small` => `M1_SMALL`).

### Context Support

If a CONTEXT section is defined in the template, it will be available as USERDATA inside the VM and can be retrieved by running the following command:

```
$ curl http://169.254.169.254/latest/user-data
ONEGATE_ENDPOINT="https://onegate...
SSH_PUBLIC_KEY="ssh-rsa ABAABeqzaC1y...
```

If the linux context packages for EC2 are installed in the VM, these parameters will be used to configure the VM. This is the list of the supported parameters for EC2.

For example, if you want to enable SSH access to the VM, an existing EC2 keypair name can be provided in the EC2 template section or the SSH public key of the user can be included in the CONTEXT section of the template.

---

**Note:** If a value for the USERDATA attribute is provided in the EC2 section of the template, the CONTEXT section will be ignored and the value provided as USERDATA will be available instead of the CONTEXT information.

---

## 4.2.5 Hybrid VM Templates

A powerful use of cloud bursting in OpenNebula is the ability to use hybrid templates, defining a VM if OpenNebula decides to launch it locally, and also defining it if it is going to be outsourced to Amazon EC2. The idea behind this is to reference the same kind of VM even if it is incarnated by different images (the local image and the remote AMI).

An example of a hybrid template:

---

```
## Local Template section
NAME=MNyWebServer

CPU=1
MEMORY=256

DISK=[IMAGE="nginx-golden"]
NIC=[NETWORK="public"]

EC2=[
  AMI="ami-xxxxx" ]
```

OpenNebula will use the first portion (from NAME to NIC) in the above template when the VM is scheduled to a local virtualization node, and the EC2 section when the VM is scheduled to an EC2 node (ie, when the VM is going to be launched in Amazon EC2).

## 4.2.6 Testing

You must create a template file containing the information of the AMIs you want to launch. Additionally if you have an elastic IP address you want to use with your EC2 instances, you can specify it as an optional parameter.

```
CPU       = 1
MEMORY    = 1700

# KVM template machine, this will be use when submitting this VM to local resources
DISK      = [ IMAGE_ID = 3 ]
NIC       = [ NETWORK_ID = 7 ]

#EC2 template machine, this will be use wen submitting this VM to EC2

PUBLIC_CLOUD = [ TYPE="EC2",
                 AMI="ami-00bafcb5",
                 KEYPAIR="gsg-keypair",
                 INSTANCETYPE=m1.small]

#Add this if you want to use only EC2 cloud
#SCHED_REQUIREMENTS = 'HOSTNAME = "ec2"'
```

You only can submit and control the template using the OpenNebula interface:

```
$ onetemplate create ec2template
$ onetemplate instantiate ec2template
```

Now you can monitor the state of the VM with

```
$ onevm list
    ID USER       GROUP      NAME         STAT CPU      MEM      HOSTNAME        TIME
     0 oneadmin oneadmin one-0          runn   0       0K           ec2    0d 07:03
```

Also you can see information (like IP address) related to the amazon instance launched via the command. The attributes available are:

- AWS_DNS_NAME

- AWS_PRIVATE_DNS_NAME

- AWS_KEY_NAME

- AWS_AVAILABILITY_ZONE

- AWS_PLATFORM

- AWS_VPC_ID

- AWS_PRIVATE_IP_ADDRESS

- AWS_IP_ADDRESS

- AWS_SUBNET_ID

- AWS_SECURITY_GROUPS

- AWS_INSTANCE_TYPE

```
$ onevm show 0
VIRTUAL MACHINE 0 INFORMATION
ID                  : 0
NAME                : pepe
USER                : oneadmin
GROUP               : oneadmin
STATE               : ACTIVE
LCM_STATE           : RUNNING
RESCHED             : No
HOST                : ec2
CLUSTER ID          : -1
START TIME          : 11/15 14:15:16
END TIME            : -
DEPLOY ID           : i-a0c5a2dd

VIRTUAL MACHINE MONITORING
USED MEMORY         : 0K
NET_RX              : 208K
NET_TX              : 4K
USED CPU            : 0.2

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

VIRTUAL MACHINE HISTORY
SEQ HOST            ACTION            DS        START         TIME     PROLOG
  0 ec2             none               0  11/15 14:15:37   2d 21h48m  0h00m00s

USER TEMPLATE
PUBLIC_CLOUD=[
  TYPE="EC2",
  AMI="ami-6f5f1206",
  INSTANCETYPE="m1.small",
  KEYPAIR="gsg-keypair" ]
SCHED_REQUIREMENTS="ID=4"

VIRTUAL MACHINE TEMPLATE
AWS_AVAILABILITY_ZONE="us-east-1d"
AWS_DNS_NAME="ec2-54-205-155-229.compute-1.amazonaws.com"
AWS_INSTANCE_TYPE="m1.small"
AWS_IP_ADDRESS="54.205.155.229"
AWS_KEY_NAME="gsg-keypair"
AWS_PRIVATE_DNS_NAME="ip-10-12-101-169.ec2.internal"
```

(continues on next page)

```
AWS_PRIVATE_IP_ADDRESS="10.12.101.169"
AWS_SECURITY_GROUPS="sg-8e45a3e7"
```

### 4.2.7 Scheduler Configuration

Since ec2 Hosts are treated by the scheduler like any other host, VMs will be automatically deployed in them. But you probably want to lower their priority and start using them only when the local infrastructure is full.

#### Configure the Priority

The ec2 drivers return a probe with the value PRIORITY = -1. This can be used by the scheduler, configuring the 'fixed' policy in `sched.conf`:

```
DEFAULT_SCHED = [
    policy = 4
]
```

The local hosts will have a priority of 0 by default, but you could set any value manually with the 'onehost/onecluster update' command.

There are two other parameters that you may want to adjust in sched.conf:

```
-  MAX_DISPATCH: Maximum number of Virtual Machines actually dispatched to a host in
→each scheduling action
-  MAX_HOST: Maximum number of Virtual Machines dispatched to a given host in each
→scheduling action
```

In a scheduling cycle, when MAX_HOST number of VMs have been deployed to a host, it is discarded for the next pending VMs.

For example, having this configuration:

* MAX_HOST = 1

* MAX_DISPATCH = 30

* 2 Hosts: 1 in the local infrastructure, and 1 using the ec2 drivers

* 2 pending VMs

The first VM will be deployed in the local host. The second VM will have also sort the local host with higher priority, but because 1 VMs was already deployed, the second VM will be launched in ec2.

A quick way to ensure that your local infrastructure will be always used before the ec2 hosts is to **set MAX_DISPATH to the number of local hosts**.

#### Force a Local or Remote Deployment

The ec2 drivers report the host attribute PUBLIC_CLOUD = YES. Knowing this, you can use that attribute in your VM requirements.

To force a VM deployment in a local host, use:

```
SCHED_REQUIREMENTS = "!(PUBLIC_CLOUD = YES)"
```

To force a VM deployment in an ec2 host, use:

```
SCHED_REQUIREMENTS = "PUBLIC_CLOUD = YES"
```

### 4.2.8 Importing VMs

VMs running on EC2 that were not launched through OpenNebula can be imported in OpenNebula.

### 4.2.9 Permissions requirement

If the user account that is going to be used does not have full permissions here is a table that summarizes the privileges required by ec2 driver.

#### Service STS

| Privileges | Resources |
|---|---|
| Write/DecodeAuthorizationMessage | Support all resources |

#### Service EC2

| Privileges | Resources |
|---|---|
| List/DescribeInstances | Support all resources |
| Read/DescribeTags | Support all resources |
| Write/AssociateAddress | Support all resources |
| Write/RunInstances | image, instance, key-pair, network-interface, security-group, subnet, volume |
| Write/StartInstances | image, instance, key-pair, network-interface, security-group, subnet, volume |
| Write/StopInstances | image, instance, key-pair, network-interface, security-group, subnet, volume |
| Write/TerminateInstances | image, instance, key-pair, network-interface, security-group, subnet, volume |
| Write/CreateTags | instance |

## 4.3 Azure Driver

### 4.3.1 Considerations & Limitations

You should take into account the following technical considerations when using the Microsoft Azure (AZ) cloud with OpenNebula:

- There is no direct access to the hypervisor, so it cannot be monitored (we don't know where the VM is running on the Azure cloud).

- The usual OpenNebula functionality for snapshotting, hot-plugging, or migration is not available with Azure.

- By default OpenNebula will always launch Small (1 CPU, 1792 MB RAM) instances, unless otherwise specified. The following table is an excerpt of all the instance types available in Azure, a more exhaustive list can be found (and edited) in `/etc/one/az_driver.conf`.

| Name | CPU Capacity | Memory Capacity |
|------|--------------|-----------------|
| ExtraSmall | 0.1 Cores | 768 MB |
| Small | 1 Cores | 1792 MB |
| Medium | 2 Cores | 3584 MB |
| Large | 4 Cores | 7168 MB |
| ExtraLarge | 8 Cores | 14336 MB |
| A5 | 2 Cores | 14336 MB |
| A6 | 4 Cores | 28672 MB |
| A7 | 8 Cores | 57344 MB |
| A8 | 8 Cores | 57344 MB |
| A9 | 16 Cores | 114688 MB |

### 4.3.2 Prerequisites

- You must have a working account for Azure.

- First, the Subscription ID, that can be uploaded and retrieved from All services -> Subscriptions.

- Second, the Management Certificate file, that can be created with the following steps- We need the .pem file (for the ruby gem) and the .cer file (to upload to Azure):

```
## Install openssl

## CentOS
$ sudo yum install openssl

## Ubuntu
$ sudo apt-get install openssl

# Move to a folder (wherever you prefer, it's better if you choose a private folder␣
→to store all yours keys)
$ mkdir ~/.ssh/azure && cd ~/.ssh/azure

## Create certificate
$ openssl req -x509 -nodes -days 365 -newkey rsa:2048 -keyout myPrivateKey.key -out␣
→myCert.pem
$ chmod 600 myPrivateKey.key

## Generate .cer file for Azure
$ openssl x509 -outform der -in myCert.pem -out myCert.cer

## You should have now your .pem certificate and your private key
$ find .
==>
    ./myCert.pem
    ./myPrivateKey.key
```

- Third, the certificate file (.cer) has to be uploaded to All services -> Subscriptions -> Management certificates

- In order to allow azure driver to properly authenticate with our Azure account, you need to sign your .pem file:

```
## Concatenate key and pem certificate (sign with private key)
$ cat myCert.pem myPrivateKey.key > azureOne.pem
```

azureOne.pem is the result of your signed cert, OpenNebula does not need you to store this cert in any certain location of the OpenNebula front-end filesystem, please keep it safe. Remember that you will need to read it in the host creation

process. We'll talk about how perform this action later.

- The following gem is required: `azure`. This gem is automatically installed as part of the installation process. Otherwise, run the `install_gems` script as root:

```
# /usr/share/one/install_gems cloud
```

### 4.3.3 OpenNebula Configuration

Uncomment the Azure AZ IM and VMM drivers from `/etc/one/oned.conf` file in order to use the driver.

```
IM_MAD = [
     name       = "az",
     executable = "one_im_sh",
     arguments  = "-c -t 1 -r 0 az" ]

VM_MAD = [
    name       = "az",
    executable = "one_vmm_sh",
    arguments  = "-t 15 -r 0 az",
    type       = "xml" ]
```

Driver flags are the same as other drivers:

| FLAG | SETs |
|------|------|
| -t | Number of threads, i.e. number of actions performed at the same time |
| -r | Number of retries when contacting Azure service |

Azure driver has his own configuration file with a few options ready to customize, take a look inside your opennebula etc folder, edit the file `/etc/one/az_driver.conf`:

```
proxy_uri:
instance_types:
    ExtraSmall:
        cpu: 1
        memory: 0.768
    Small:
        cpu: 1
        memory: 1.75
    Medium:
        cpu: 2
        memory: 3.5
    Large:
        cpu: 4
        memory: 7.0
    ExtraLarge:
        cpu: 8
        memory: 14.0
    ...
```

In the above file, each instance_type represents the physical resources that Azure will serve.

If the OpenNebula frontend needs to use a proxy to connect to the public Internet you also need to configure the proxy in that file. The parameter is called `proxy_uri`. Authenticated proxies are not supported, that is, the ones that require user name and password. For example, if the proxy is in `10.0.0.1` and its port is `8080` the configuration line should read:

```
proxy_uri: http://10.0.0.1:8080
```

> **Warning:** `instance_types` section shows us the machines that Azure is able to provide, the azure driver will retrieve this kind of information so it's better to not change it unless you are aware of your actions.

> **Warning:** If you were using OpenNebula before 5.4 you may have noticed that there are not Microsoft credentials in configuration file anymore, this is due security reasons. In 5.4 there is a new secure credentials storage for Microsoft's accounts so you do not need to store sensitive credential data inside your disk. OpenNebula daemon stores the data in an encrypted format.

Once the file is saved, OpenNebula needs to be restarted, create a new Host with Microsoft's credentials that uses the AZ drivers:

```
$ onehost create azure_host -t az -i az -v az
```

> **Note:** `-t` is needed to specify what type of remote provider host we want to set up, if you've followed all the instruction properly your default editor should show in your screen asking for the credentials and other mandatory data that will allow you to communicate with Azure.

Once you have opened your editor you can look for additional help at the top of your screen, you have more information in *Azure Auth template Attributes* section. The basic three variables you have to set are: `AZ_ID`, `AZ_CERT` and `REGION_NAME`.

### 4.3.4 Azure Specific Template Attributes

In order to deploy an instance in Azure through OpenNebula you must include an PUBLIC_CLOUD section in the virtual machine template. This is an example of a virtual machine template that can be deployed in our local resources or in Azure.

```
CPU       = 0.5
MEMORY    = 128

# KVM template machine, this will be use when submitting this VM to local resources
DISK      = [ IMAGE_ID = 3 ]
NIC       = [ NETWORK_ID = 7 ]

# Azure template machine, this will be use wen submitting this VM to Azure
PUBLIC_CLOUD = [
  TYPE=AZURE,
  INSTANCE_TYPE=ExtraSmall,
  IMAGE=b39f27a8b8c64d52b05eac6a62ebad85__Ubuntu-14_04-LTS-amd64-server-20140606.1-en-
→us-30GB,
  VM_USER="azuser",
  VM_PASSWORD="myr@nd0mPass9",
  WIN_RM="https",
  TCP_ENDPOINTS="80",
  SSHPORT=2222
]
```

```
#Add this if you want this VM to only go to the West EuropeAzure cloud
#SCHED_REQUIREMENTS = 'HOSTNAME = "west-europe"'
```

These are the attributes that can be used in the PUBLIC_CLOUD section of the template for TYPE "AZURE":

Check an exhaustive list of attributes in the Virtual Machine Definition File Reference Section.

---

**Note:** The PUBLIC_CLOUD sections allow for substitutions from template and virtual network variables, the same way as the CONTEXT section allows.

---

Default values for all these attributes can be defined in the `/etc/one/az_driver.default` file.

```
<!--
 Default configuration attributes for the Azure driver
 (all domains will use these values as defaults)
 Valid attributes are: INSTANCE_TYPE, IMAGE, VM_USER, VM_PASSWORD, LOCATION,
 STORAGE_ACCOUNT, WIN_RM, CLOUD_SERVICE, TCP_ENDPOINTS, SSHPORT, AFFINITY_GROUP,
 VIRTUAL_NETWORK_NAME, SUBNET and AVAILABILITY_SET
 Use XML syntax to specify defaults, note elements are UPCASE
 Example:
 <TEMPLATE>
   <AZURE>
     <LOCATION>west-europe</LOCATION>
     <INSTANCE_TYPE>Small</INSTANCE_TYPE>
     <CLOUD_SERVICE>MyDefaultCloudService</CLOUD_SERVICE>
     <IMAGE>0b11de9248dd4d87b18621318e037d37__RightImage-Ubuntu-12.04-x64-v13.4</
↪IMAGE>
     <VM_USER>MyUser</VM_USER>
     <VM_PASSWORD>MyPassword</VM_PASSWORD>
     <STORAGE_ACCOUNT>MyStorageAccountName</STORAGE_ACCOUNT>
     <WIN_RM>http</WIN_RM>
     <CLOUD_SERVICE>MyCloudServiceName</CLOUD_SERVICE>
     <TCP_ENDPOINTS>80,3389:3390</TCP_ENDPOINTS>
     <SSHPORT>2222</SSHPORT>
     <AFFINITY_GROUP>MyAffinityGroup</AFFINITY_GROUP>
     <VIRTUAL_NETWORK_NAME>MyVirtualNetwork</VIRTUAL_NETWORK_NAME>
     <SUBNET>MySubNet<SUBNET>
     <AVAILABILITY_SET>MyAvailabilitySetName<AVAILABILITY_SET>
   </AZURE>
 </TEMPLATE>
-->

<TEMPLATE>
  <AZURE>
     <LOCATION>west-europe</LOCATION>
     <INSTANCE_TYPE>Small</INSTANCE_TYPE>
  </AZURE>
</TEMPLATE>
```

---

**Note:** Valid Azure images to set in the IMAGE atribute of the PUBLIC_CLOUD section can be extracted with the following ruby snippet:

---

```ruby
#!/usr/bin/env ruby

require "azure"

# Get a list of available virtual machine images
def get_image_names
    vm_image_management = Azure.vm_image_management
    vm_image_management.list_os_images.each do |image|
        puts "#{image.os_type}"
        puts "      locations: #{image.locations}"
        puts "      name     : #{image.name}"
        puts
    end
end

Azure.configure do |config|
        config.management_certificate = '/path-to/azureOne.pem'
        config.subscription_id        = 'XXXXXXXX-XXXX-XXXX-XXXX-XXXXXXXXXXXX'
end

get_image_names
```

### Azure Auth Attributes

After successfully executing onehost create with -t option, your default editor will open. An example follows of how you can complete this area:

```
AZ_ID = "this-is-my-azure-identifier"
AZ_CERT = "-----BEGIN CERTIFICATE-----
        xxxxxxxxxxxxxxxxxxxxxxxxxxxxxx
        -----END CERTIFICATE-----
        -----BEGIN PRIVATE KEY-----
        xxxxxxxxxxxxxxxxxxxxxxxxxxxxxx
        -----END PRIVATE KEY-----"

REGION_NAME = "West Europe"
CAPACITY = [
    SMALL = "3",
    MEDIUM = "1" ]
```

The first two attributes have the authentication info required by Azure:

- **AZ_ID**: your Microsoft Azure account identifier, found in All services -> Subscriptions.

- **AZ_CERT**: The certificate that you signed before, in our example this file is called 'azureOne.pem' you only need to read this file one time to set this attribute and start using Azure:

```
$ cat ~/.ssh/azure/azureOne.pem
```

- Copy the content into your system clipboard without any mistake selecting all the text (ctrl + Shift + c if you are under normal linux terminal).

- Paste into AZ_CERT value, make sure you are inside quotes without leaving any blankspace.

This information will be encrypted as soon as the host is created. In the host template the values of the `AZ_ID` and `AZ_CERT` attributes will be encrypted to maintain a secure way in your future communication with azure.

- **REGION_NAME**: it's the name of Azure region that your account uses to deploy machines. You can check Microsoft's site to know more about the region availability Regions Azure page.

In the example the region is set to *West Europe*.

- **CAPACITY**: This attribute sets the size and number of Azure machines that your OpenNebula host will handle, you can see `instance_types` section in `azure_driver.conf` file to know the supported names. Remember that its mandatory to capitalize the names (`Small` => `SMALL`).

## 4.3.5 Multi Azure Location/Account Support

It is possible to define various Azure hosts to allow OpenNebula the managing of different Azure locations or different Azure accounts. OpenNebula choses the datacenter in which to launch the VM in the following way:

- if the VM description contains the LOCATION attribute, then OpenNebula knows that the VM needs to be launch in this Azure location

- if the name of the host matches the region name (remember, this is the same as an Azure location), then OpenNebula knows that the VMs sent to this host needs to be launched in that Azure datacenter

- if the VM doesn't have a LOCATION attribute, and the host name doesn't match any of the defined regions, then the default region is picked.

When you create a new host the credentials and endpoint for that host are retrieved from the `/etc/one/az_driver.conf` file using the host name. Therefore, if you want to add a new host to manage a different datacenter, i.e. `west-europe`, just add your credentials and the capacity limits to the the `west-europe` section in the conf file, and specify that name (west-europe) when creating the new host.

```
regions:
    ...
    west-europe:
        region_name: "West Europe"
        pem_management_cert: "<path-to-your-vonecloud-pem-certificate-here>"
        subscription_id: "your-subscription-id"
        management_endpoint:
        capacity:
            Small: 5
            Medium: 1
            Large: 0
```

After that, create a new Host with the `west-europe` name:

```
$ onehost create west-europe -i az -v az
```

If the Host name does not match any regions key, the `default` will be used.

You can define a different Azure section in your template for each Azure host, so with one template you can define different VMs depending on which host it is scheduled, just include a LOCATION attribute in each PUBLIC_CLOUD section:

```
PUBLIC_CLOUD = [ TYPE=AZURE,
                 INSTANCE_TYPE=Small,
                 IMAGE=b39f27a8b8c64d52b05eac6a62ebad85__Ubuntu-14_04-LTS-amd64-
→server-20140606.1-en-us-30GB,
                 VM_USER="MyUserName",
                 VM_PASSWORD="MyPassword",
                 LOCATION="brazil-south"
]
```

(continues on next page)

```
PUBLIC_CLOUD = [ TYPE=AZURE,
                 INSTANCE_TYPE=Medium,
                 IMAGE=0b11de9248dd4d87b18621318e037d37__RightImage-Ubuntu-12.04-x64-
→v13.4,
                 VM_USER="MyUserName",
                 VM_PASSWORD="MyPassword",
                 LOCATION="west-europe"
]
```

You will have a small Ubuntu 14.04 VM launched when this VM template is sent to host *brazil-south* and a medium Ubuntu 13.04 VM launched whenever the VM template is sent to host *west-europe*.

> **Warning:** If only one Azure host is defined, the Azure driver will deploy all Azure templates onto it, not paying attention to the **LOCATION** attribute.

### 4.3.6 Hybrid VM Templates

A powerful use of cloud bursting in OpenNebula is the ability to use hybrid templates, defining a VM if OpenNebula decides to launch it locally, and also defining it if it is going to be outsourced to Azure. The idea behind this is to reference the same kind of VM even if it is incarnated by different images (the local image and the Azure image).

An example of a hybrid template:

```
## Local Template section
NAME=MNyWebServer

CPU=1
MEMORY=256

DISK=[IMAGE="nginx-golden"]
NIC=[NETWORK="public"]

PUBLIC_CLOUD = [ TYPE=AZURE,
                 INSTANCE_TYPE=Medium,
                 IMAGE=0b11de9248dd4d87b18621318e037d37__RightImage-Ubuntu-12.04-x64-
→v13.4,
                 VM_USER="MyUserName",
                 VM_PASSWORD="MyPassword",
                 LOCATION="west-europe"
]
```

OpenNebula will use the first portion (from NAME to NIC) in the above template when the VM is scheduled to a local virtualization node, and the PUBLIC_CLOUD section of TYPE="AZURE" when the VM is scheduled to an Azure node (ie, when the VM is going to be launched in Azure).

### 4.3.7 Testing

You must create a template file containing the information of the VMs you want to launch.

```
CPU       = 1
MEMORY    = 1700
```

```
# KVM template machine, this will be use when submitting this VM to local resources
DISK      = [ IMAGE_ID = 3 ]
NIC       = [ NETWORK_ID = 7 ]

# Azure template machine, this will be use when submitting this VM to Azure

PUBLIC_CLOUD = [ TYPE=AZURE,
                 INSTANCE_TYPE=Medium,
                 IMAGE=0b11de9248dd4d87b18621318e037d37__RightImage-Ubuntu-12.04-x64-
→v13.4,
                 VM_USER="MyUserName",
                 VM_PASSWORD="MyPassword",
                 LOCATION="west-europe"
]

# Add this if you want to use only Azure cloud
#SCHED_REQUIREMENTS = 'HYPERVISOR = "AZURE"'
```

You can submit and control the template using the OpenNebula interface:

```
$ onetemplate create aztemplate
$ onetemplate instantiate aztemplate
```

Now you can monitor the state of the VM with

```
$ onevm list
    ID USER       GROUP      NAME        STAT CPU     MEM       HOSTNAME       TIME
     0 oneadmin oneadmin one-0        runn   0      0K    west-europe    0d 07:03
```

Also you can see information (like IP address) related to the Azure instance launched via the command. The attributes available are:

- AZ_AVAILABILITY_SET_NAME
- AZ_CLOUD_SERVICE_NAME,
- AZ_DATA_DISKS,
- AZ_DEPLOYMENT_NAME,
- AZ_DISK_NAME,
- AZ_HOSTNAME,
- AZ_IMAGE,
- AZ_IPADDRESS,
- AZ_MEDIA_LINK,
- AZ_OS_TYPE,
- AZ_ROLE_SIZE,
- AZ_TCP_ENDPOINTS,
- AZ_UDP_ENDPOINTS,
- AZ_VIRTUAL_NETWORK_NAME

```
$ onevm show 0
VIRTUAL MACHINE 0 INFORMATION
ID                  : 0
NAME                : one-0
USER                : oneadmin
GROUP               : oneadmin
STATE               : ACTIVE
LCM_STATE           : RUNNING
RESCHED             : No
START TIME          : 06/25 13:05:29
END TIME            : -
HOST                : west-europe
CLUSTER ID          : -1
DEPLOY ID           : one-0_opennebuladefaultcloudservicename-0


VIRTUAL MACHINE MONITORING
USED MEMORY         : 0K
USED CPU            : 0
NET_TX              : 0K
NET_RX              : 0K

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

VIRTUAL MACHINE HISTORY
SEQ HOST            ACTION              DS          START         TIME      PROLOG
  0 west-europe    none                -1  06/25 13:06:25   0d 00h06m   0h00m00s



USER TEMPLATE
PUBLIC_CLOUD=[
  IMAGE="b39f27a8b8c64d52b05eac6a62ebad85__Ubuntu-14_04-LTS-amd64-server-20140606.1-
→en-us-30GB",
  INSTANCE_TYPE="ExtraSmall",
  SSH_PORT="2222",
  TCP_ENDPOINTS="80",
  TYPE="AZURE",
  VM_PASSWORD="MyVMPassword",
  VM_USER="MyUserName",
  WIN_RM="https" ]
VIRTUAL MACHINE TEMPLATE
AUTOMATIC_REQUIREMENTS="!(PUBLIC_CLOUD = YES) | (PUBLIC_CLOUD = YES & (HYPERVISOR =
→AZURE | HYPERVISOR = AZURE))"
AZ_CLOUD_SERVICE_NAME="opennebuladefaultcloudservicename-0"
AZ_DEPLOYMENT_NAME="OpenNebulaDefaultCloudServiceName-0"
AZ_DISK_NAME="OpenNebulaDefaultCloudServiceName-0-one-0_
→OpenNebulaDefaultCloudServiceName-0-0-201406251107210062"
AZ_HOSTNAME="ubuntu"
AZ_IMAGE="b39f27a8b8c64d52b05eac6a62ebad85__Ubuntu-14_04-LTS-amd64-server-20140606.1-
→en-us-30GB"
AZ_IPADDRESS="191.233.70.93"
AZ_MEDIA_LINK="http://one0opennebuladefaultclo.blob.core.windows.net/vhds/disk_2014_
→06_25_13_07.vhd"
AZ_OS_TYPE="Linux"
```

(continues on next page)

```
AZ_ROLE_SIZE="ExtraSmall"
AZ_TCP_ENDPOINTS="name=SSH,vip=23.97.101.202,publicport=2222,local_
↪port=22,local_
↪port=tcp;name=TCP-PORT-80,vip=23.97.101.202,publicport=80,local_port=80,local_
↪port=tcp"
CPU="1"
MEMORY="1024"
VMID="0"
```

### 4.3.8 Scheduler Configuration

Since Azure Hosts are treated by the scheduler like any other host, VMs will be automatically deployed in them. But you probably want to lower their priority and start using them only when the local infrastructure is full.

#### Configure the Priority

The Azure drivers return a probe with the value PRIORITY = -1. This can be used by the scheduler, configuring the 'fixed' policy in `sched.conf`:

```
DEFAULT_SCHED = [
    policy = 4
]
```

The local hosts will have a priority of 0 by default, but you could set any value manually with the 'onehost/onecluster update' command.

There are two other parameters that you may want to adjust in sched.conf:

- MAX_DISPATCH: Maximum number of Virtual Machines actually dispatched to a host in each scheduling action
- MAX_HOST: Maximum number of Virtual Machines dispatched to a given host in each scheduling action

In a scheduling cycle, when MAX_HOST number of VMs have been deployed to a host, it is discarded for the next pending VMs.

For example, having this configuration:

- MAX_HOST = 1
- MAX_DISPATCH = 30
- 2 Hosts: 1 in the local infrastructure, and 1 using the Azure drivers
- 2 pending VMs

The first VM will be deployed in the local host. The second VM will have also sort the local host with higher priority, but because 1 VMs was already deployed, the second VM will be launched in Azure.

A quick way to ensure that your local infrastructure will be always used before the Azure hosts is to **set MAX_DISPATH to the number of local hosts**.

#### Force a Local or Remote Deployment

The Azure drivers report the host attribute PUBLIC_CLOUD = YES. Knowing this, you can use that attribute in your VM requirements.

To force a VM deployment in a local host, use:

```
SCHED_REQUIREMENTS = "!(PUBLIC_CLOUD = YES)"
```

To force a VM deployment in a Azure host, use:

```
SCHED_REQUIREMENTS = "PUBLIC_CLOUD = YES"
```

### 4.3.9 Importing VMs

VMs running on Azure that were not launched through OpenNebula can be imported in OpenNebula.

## 4.4 One-to-One hybrid Driver

### 4.4.1 Considerations & Limitations

- You need to setup an user account in the remote OpenNebula cloud. All operations will be mapped to this account in the remote OpenNebula cloud, you may need to consider this when setting up quotas or access rules for this remote account.

- The following operations are not supported for VMs in the remote cloud: disk operations (attach/detach, snapshots or save), nic operations (attach/detach), migrations or reconfigurations.

- You need to refer the remote VM Templates by their ID, see below for hints to do so.

- Context attribute can be added to the local VM Template, these values will be copied to the remote machine.

> **Warning:** The attribute FILES is not supported.

- All states of local virtual machines mirror those of the remote virtual machines except for running, that it may also include pending or any other active states.

### 4.4.2 Prerequisites

An user account needs to be setup in remote OpenNebula. This user should have access to the VM Templates that you are going to expose for hybrid access.

> **Note:** You can check if the user has access to the remote template with the command `onetemplate list --endpoint <REMOTE_ENDPOINT> --user <REMOTE_USER> --password <REMOTE_PASS>`.

### 4.4.3 OpenNebula Configuration

Uncomment the OpenNebula IM and VMM drivers from `/etc/one/oned.conf` file in order to use the driver.

```
IM_MAD = [
    NAME          = "one",
    SUNSTONE_NAME = "OpenNebula",
    EXECUTABLE    = "one_im_sh",
    ARGUMENTS     = "-c -t 1 -r 0 one" ]
```

(continues on next page)

```
VM_MAD = [
    NAME           = "one",
    SUNSTONE_NAME  = "OpenNebula",
    EXECUTABLE     = "one_vmm_sh",
    ARGUMENTS      = "-t 15 -r 0 one",
    TYPE           = "xml",
    KEEP_SNAPSHOTS = "no"
]
```

Driver flags are the same as other drivers:

| FLAG | SETs |
| --- | --- |
| -t | Number of threads |
| -r | Number of retries |

### Defining a Hybrid OpenNebula Cloud

First create a new Host with *im* and *vm* drivers set to *one*.

```
$ onehost create <name> -i one -v one
```

And, now add the user attributes to connect to the hybrid OpenNebula within host template:

```
$ onehost update <hostid>

ONE_USER = <remote_username>

ONE_PASSWORD = <remote_password>
ONE_ENDPOINT = <remote_endpoint>
ONE_CAPACITY = [
    CPU    = 0,
    MEMORY = 0
]
```

The following table describes each supported attribute to configure the hybrid OpenNebula cloud.

**Note:** The attribute ONE_CAPACITY includes two attributes, CPU and MEMORY, these attributes limit the usage of the remote resources.

| ATTRIBUTE | DESCRIPTION |
| --- | --- |
| ONE_USER | Remote username |
| ONE_PASSWORD | Remote password for the username |
| ONE_ENDPOINT | Remote endpoint url to access |
| ONE_CAPACITY | 0 indicate that the quotas are taken from the user and group quotas of the remote OpenNebula user. Alternatively, you can set a hard limit here |

### 4.4.4 OpenNebula to OpenNebula Specific Template Attributes

The following section describes how to create a new local template and map it to the remote template.

### Local VM Template

Firstly, find out the remote templates you have access to:

```
$ onetemplate list --endpoint http://<hybrid_OpenNebula_cloud>:2633/RPC2 --user
↪<username> --password <pass>
```

Now, create a new local template for each remote template you want to use. It is recommended to set the same CPU and MEMORY as the remote Template. For example:
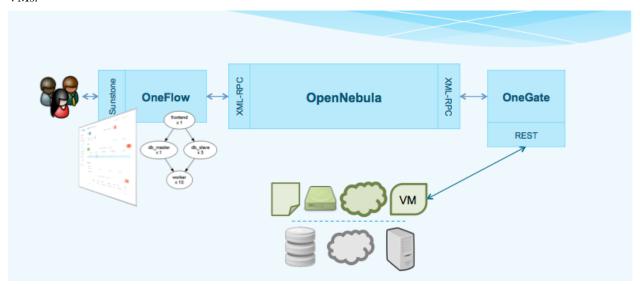
```
$ cat template.txt
NAME = "hybrid-template"

CPU    = 0.1
MEMORY = 128

PUBLIC_CLOUD = [
    TEMPLATE_ID = "0",
    TYPE        = "opennebula"
]

CONTEXT=[
    NETWORK="yes"
]

$ onetemplate create template.txt
ID: 0
```

**Note:** Your hybrid VM Template must set TEMPLATE_ID to the target VM Template ID in the **remote OpenNebula**.

# APPLICATION INSIGHT

## 5.1 Overview

The OneGate component allows Virtual Machine guests to pull and push VM information from OpenNebula. Users and administrators can use it to gather metrics, detect problems in their applications, and trigger OneFlow elasticity rules from inside the VM.

For Virtual Machines that are part of a Multi-VM Application (*OneFlow Service*), they can also retrieve the Service information directly from OneGate and trigger actions to reconfigure the Service or pass information among different VMs.



### 5.1.1 How Should I Read This Chapter

This chapter should be read after the infrastructure is properly setup, and contains working Virtual Machine templates.

Proceed to each section following these links:

- *OneGate Server Configuration*
- *Application Monitoring*

### 5.1.2 Hypervisor Compatibility

This chapter applies to all the hypervisors.

## 5.2 OneGate Server Configuration

The OneGate service allows Virtual Machine guests to pull and push VM information from OpenNebula. Although it is installed by default, its use is completely optional.

### 5.2.1 Requirements

Check the Installation guide for details of what package you have to install depending on your distribution

OneGate can be used with VMs running on KVM, vCenter, and *EC2*.

Currently, OneGate is not supported for VMs instantiated in Azure, since the authentication token is not available inside these VMs. OneGate support for these drivers will be include in upcoming releases.

### 5.2.2 Configuration

The OneGate configuration file can be found at `/etc/one/onegate-server.conf`. It uses YAML syntax to define the following options:

**Server Configuration**

- `one_xmlrpc`: OpenNebula daemon host and port
- `host`: Host where OneGate will listen
- `port`: Port where OneGate will listen
- `ssl_server`: SSL proxy URL that serves the API (set if is being used)

**Log**

- `debug_level`: Log debug level. 0 = ERROR, 1 = WARNING, 2 = INFO, 3 = DEBUG

**Auth**

- `auth`: Authentication driver for incomming requests.
  - `onegate`: based on token provided in the context
- `core_auth`: Authentication driver to communicate with OpenNebula core.
  - `cipher` for symmetric cipher encryption of tokens
  - `x509` for x509 certificate encryption of tokens. For more information, visit the OpenNebula Cloud Auth documentation.
- `oneflow_server` Endpoint where the OneFlow server is listening.
- `permissions` By default OneGate exposes all the available API calls, each of the actions can be enabled/disabled in the server configuration.
- `restricted_attrs` Attrs that cannot be modified when updating a VM template
- `restricted_actions` Actions that cannot be performed on a VM

This is the default file

```
################################################################################
# Server Configuration
################################################################################
```

(continues on next page)

```yaml
# OpenNebula sever contact information
#
:one_xmlrpc: http://localhost:2633/RPC2

# Server Configuration
#
:host: 127.0.0.1
:port: 5030

# SSL proxy URL that serves the API (set if is being used)
#:ssl_server: https://service.endpoint.fqdn:port/


################################################################################
# Log
################################################################################


# Log debug level
#   0 = ERROR, 1 = WARNING, 2 = INFO, 3 = DEBUG
#
:debug_level: 3


################################################################################
# Auth
################################################################################


# Authentication driver for incomming requests
#   onegate, based on token provided in the context
#
:auth: onegate

# Authentication driver to communicate with OpenNebula core
#   cipher, for symmetric cipher encryption of tokens
#   x509, for x509 certificate encryption of tokens
#
:core_auth: cipher


################################################################################
# OneFlow Endpoint
################################################################################


:oneflow_server: http://localhost:2474


################################################################################
# Permissions
################################################################################

:permissions:
  :vm:
    :show: true
    :show_by_id: true
    :update: true
    :update_by_id: true
    :action_by_id: true
  :service:
    :show: true
    :change_cardinality: true
```

```
# Attrs that cannot be modified when updating a VM template
:restricted_attrs
  - SCHED_REQUIREMENTS
  - SERVICE_ID
  - ROLE_NAME

# Actions that cannot be performed on a VM
:restricted_actions
  #- deploy
  #- delete
  #- hold
  ...
```

## 5.2.3 Start OneGate

To start and stop the server, use the `opennebula-gate` service:

```
# systemctl start opennebula-gate
```

Or use service in older Linux systems:

```
# service opennebula-gate start
```

> **Warning:** By default, the server will only listen to requests coming from localhost. Change the `:host` attribute in `/etc/one/onegate-server.conf` to your server public IP, or 0.0.0.0 so onegate will listen on any interface.

Inside `/var/log/one/` you will find new log files for the server:

```
/var/log/one/onegate.error
/var/log/one/onegate.log
```

## 5.2.4 Use OneGate

Before your VMs can communicate with OneGate, you need to edit `/etc/one/oned.conf` and set the OneGate endpoint. This IP must be reachable from your VMs.

```
ONEGATE_ENDPOINT = "http://192.168.0.5:5030"
```

At this point the service is ready, you can continue to the *OneGate usage documentation*.

## 5.2.5 Configuring a SSL Proxy

This is an example on how to configure Nginx as a ssl proxy for Onegate in Ubuntu.

Update your package lists and install Nginx:

```
$ sudo apt-get update
$ sudo apt-get install nginx
```

You should get an official signed certificate, but for the purpose of this example we will generate a self-signed SSL certificate:

```
$ cd /etc/one
$ sudo openssl req -x509 -nodes -days 365 -newkey rsa:2048 -keyout /etc/one/cert.key -
↪out /etc/one/cert.crt
```

Next you will need to edit the default Nginx configuration file or generate a new one. Change the ONEGATE_ENDPOINT variable with your own domain name.

```
server {
  listen 80;
  return 301 https://$host$request_uri;
}

server {
  listen 443;
  server_name ONEGATE_ENDPOINT;

  ssl_certificate         /etc/one/cert.crt;
  ssl_certificate_key     /etc/one/cert.key;

  ssl on;
  ssl_session_cache  builtin:1000  shared:SSL:10m;
  ssl_protocols  TLSv1 TLSv1.1 TLSv1.2;
  ssl_ciphers HIGH:!aNULL:!eNULL:!EXPORT:!CAMELLIA:!DES:!MD5:!PSK:!RC4;
  ssl_prefer_server_ciphers on;

  access_log              /var/log/nginx/onegate.access.log;

  location / {

    proxy_set_header        Host $host;
    proxy_set_header        X-Real-IP $remote_addr;
    proxy_set_header        X-Forwarded-For $proxy_add_x_forwarded_for;
    proxy_set_header        X-Forwarded-Proto $scheme;

    # Fix the "It appears that your reverse proxy set up is broken" error.
    proxy_pass          http://localhost:5030;
    proxy_read_timeout  90;

    proxy_redirect      http://localhost:5030 https://ONEGATE_ENDPOINT;
  }
}
```

Update /etc/one/oned.conf with the new OneGate endpoint

```
ONEGATE_ENDPOINT = "https://ONEGATE_ENDPOINT"
```

Update /etc/one/onegate-server.conf with the new OneGate endpoint and uncomment the ssl_server parameter

```
:ssl_server: https://ONEGATE_ENDPOINT
```

Then restart oned, onegate-server and Nginx:

```
$ sudo service nginx restart
$ sudo service opennebula restart
$ sudo service opennebula-gate restart
```
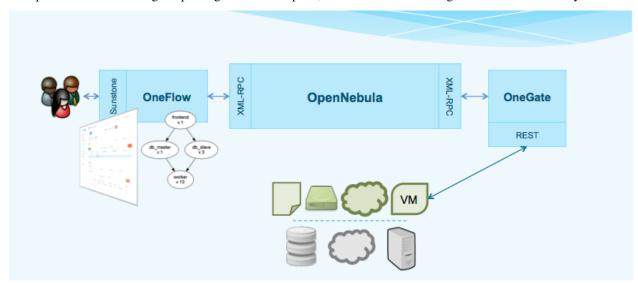
## 5.3 OneGate Usage

The OneGate component allows Virtual Machine guests to pull and push VM information from OpenNebula. Users and administrators can use it to gather metrics, detect problems in their applications, and trigger OneFlow elasticity rules from inside the VM.

For Virtual Machines that are part of a Multi-VM Application (*OneFlow Service*), they can also retrieve the Service information directly from OneGate and trigger actions to reconfigure the Service or pass information among different VMs.

### 5.3.1 OneGate Workflow Explained

OneGate is a server that listens to http connections from the Virtual Machines. OpenNebula assigns an individual token to each VM instance, and Applications running inside the VM use this token to interact with the OneGate API. This token is generated using VM information and signed with the owner User template attribute TOKEN_PASSWORD. This password can be changed updating the User template, but tokens from existing VMs will not work anymore.



### 5.3.2 OneGate Usage

First, the cloud administrator must configure and start the *OneGate server*.

#### Setup the VM Template

Your VM Template must set the `CONTEXT/TOKEN` attribute to `YES`.

```
CPU     = "0.5"
MEMORY  = "1024"

DISK = [
  IMAGE_ID = "0" ]
NIC = [
  NETWORK_ID = "0" ]
```
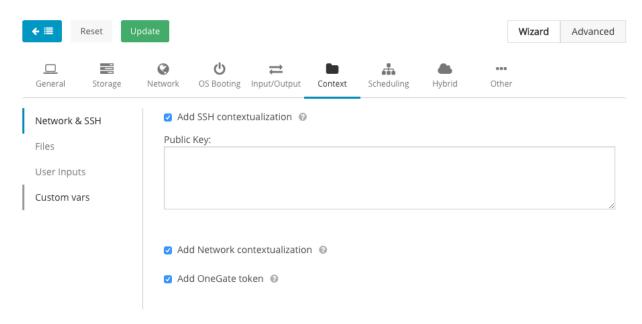
<div align="right">(continues on next page)</div>

```
CONTEXT = [
  TOKEN = "YES" ]
```

or check the OneGate checkbox in Sunstone:



When this Template is instantiated, OpenNebula will automatically add the ONEGATE_ENDPOINT context variable, and a token.txt will be placed in the context cdrom. This token.txt file is only accessible from inside the VM.

```
...

CONTEXT=[
  DISK_ID="1",
  ONEGATE_ENDPOINT="http://192.168.0.1:5030",
  TARGET="hdb",
  TOKEN="YES" ]
```

In vCenter this information is available in the extraConfig section of the VM metadata, available in the guest OS through the VMware tools as explained in the contextualization guide.

### Using the OneGate Client inside the Guest VM

A ruby client that implements the OneGate API is included in the official OpenNebula context packages. This is a simple command line interface to interact with the OneGate server, it will handle the authentication and requests complexity.

### OneGate Client Usage

Available commands and usage are shown with `onegate -h`.

---

With the appropriate policies implemented in the Service, these mechanisms allow Services to be self-managed, enabling self-configuration, self-healing, self-optimization and self-protection.

### Self-Awareness

There are several actions available to retrieve information of the Virtual Machine and the Service it belongs to. A Virtual Machine can also retrieve information of other Virtual Machines that are part of the Service.

### Retrieving Information of the VM

Using the `onegate vm show` command the information of the Virtual Machine will be retrieved. For a detailed version use the `--json` option and all the information will be returned in JSON format.

If no argument is provided, the information of the current Virtual Machine will be retrieved. Alternatively, a VM ID can be provided to retrieve the information of a specific Virtual Machine.

```
$ onegate vm show
VM 8
NAME               : master_0_(service_1)
STATE              : RUNNING
IP                 : 192.168.122.23
```

### Retrieving information of the Service

Using the `onegate service show` command the information of the Service will be retrieved. For a detailed version use the `--json` option and all the information will be returned in JSON format.

```
$ onegate service show
SERVICE 1
NAME               : PANACEA service
STATE              : RUNNING

ROLE master
VM 8
NAME               : master_0_(service_1)
STATE              : RUNNING
IP                 : 192.168.122.23

ROLE slave
VM 9
NAME               : slave_0_(service_1)
STATE              : RUNNING
```

### Updating the VM Information

The Virtual Machine can update the information of itself or other Virtual Machine of the Service. This information can be retrieved from any of the Virtual Machines.

For example, the master Virtual Machine can change the `ACTIVE` attribute from one Virtual Machine to another one. Then, this information can be used to trigger any kind of action in the other Virtual Machine.

```
$ onegate vm update 9 --data ACTIVE=YES
$ onegate vm show 9 --json
{
  "VM": {
    "NAME": "slave_0_(service_1)",
    "ID": "9",
    "STATE": "3",
    "LCM_STATE": "3",
    "USER_TEMPLATE": {
      "ACTIVE": "YES",
      "FROM_APP": "4fc76a938fb81d3517000003",
      "FROM_APP_NAME": "ttylinux - kvm",
      "LOGO": "images/logos/linux.png",
      "ROLE_NAME": "slave",
      "SERVICE_ID": "1"
    },
    "TEMPLATE": {
      "NIC": [

      ]
    }
  }
}
```

### Self-Configuration

There are several actions to adapt the Service to a given situation. Actions on any of the Virtual Machines can be performed individually. Also, the size of the Service can be customized just specifying a cardinality for each of the roles.

### Performing actions on a VM

The following actions can be performed in any of the Virtual Machines of the Service.

- `onegate vm resume`: Resumes the execution of the a saved VM. Valid states: STOPPED, SUSPENDED, UNDEPLOYED, POWEROFF

- `onegate vm stop`: Stops a running VM. The VM state is saved and transferred back to the front-end along with the disk files. Valid states: RUNNING

- `onegate vm suspend`: Saves a running VM. It is the same as `onegate vm stop`, but the files are left in the remote machine to later restart the VM there (i.e. the resources are not freed and there is no need to re-schedule the VM). Valid states: RUNNING

- `onegate vm terminate`: Terminates the given VM. The VM life cycle will end. With –hard it unplugs the VM. Valid states: any except those with a pending driver response

- `onegate vm reboot`: Reboots the given VM, this is equivalent to execute the reboot command from the VM console. The VM will be ungracefully rebooted if –hard is used. Valid states: RUNNING

- `onegate vm poweroff`: Powers off the given VM. The VM will remain in the poweroff state, and can be powered on with the `onegate vm resume` command. Valid states: RUNNING

- `onegate vm resched`: Sets the rescheduling flag for the VM. The VM will be moved to a different host based on the scheduling policies. Valid states: RUNNING

- `onegate vm unresched`: Unsets the rescheduling flag for the VM. Valid states: RUNNING

- `onegate vm hold`: Sets the given VM on hold. A VM on hold is not scheduled until it is released. Valid states: PENDING

- `onegate vm release`: Releases a VM on hold. See *onegate vm hold* Valid states: HOLD

```
$ onegate vm terminate --hard 9
```

### Change Service cardinality

The number of Virtual Machines of a Service can be also modified from any of the Virtual Machines that have access to the OneGate Server. The Virtual Machines of Services are grouped in Roles and each Role has a cardinality (number of Virtual Machines). This cardinality can be increased or decreased, in case the given cardinality is lower than the current one, Virtual Machines will be terminated to meet the given number. If the cardinality is greater than the current one, new Virtual Machines will be instantiated using the VM Template associated to the Role.

```
$ onegate service scale --role slave --cardinality 2
$ onegate service show
SERVICE 1
NAME                    : PANACEA service
STATE                   : SCALING

ROLE master
VM 8
NAME                    : master_0_(service_1)
STATE                   : RUNNING
IP                      : 192.168.122.23

ROLE slave
VM 9
NAME                    : slave_0_(service_1)
STATE                   : RUNNING
VM 10
NAME                    : slave_1_(service_1)
STATE                   : PENDING
```

### 5.3.3 OneGate API

OneGate provides a REST API. To use this API you will need to get some data from the CONTEXT file.

The contextualization cdrom should contain the `context.sh` and `token.txt` files.

```
# mkdir /mnt/context
# mount /dev/hdb /mnt/context
# cd /mnt/context
# ls
context.sh  token.txt
# cat context.sh
# Context variables generated by OpenNebula
DISK_ID='1'
ONEGATE_ENDPOINT='http://192.168.0.1:5030'
VMID='0'
TARGET='hdb'
TOKEN='yes'
```

```
# cat token.txt
yCxieDUS7kra7Vn9ILA0+g==
```

With that data, you can obtain the headers required for all the ONEGATE API methods:

- **Headers**:
    - `X-ONEGATE-TOKEN: token.txt contents`
    - `X-ONEGATE-VMID: <vmid>`

OneGate supports these actions:

### Self-awareness

- `GET ${ONEGATE_ENDPOINT}/vm`: To request information about the current Virtual Machine.

- `GET ${ONEGATE_ENDPOINT}/vms/${VM_ID}`: To request information about a specific Virtual Machine of the Service. The information is returned in JSON format and is ready for public cloud usage:

```
$ curl -X "GET" "${ONEGATE_ENDPOINT}/vm" \
    --header "X-ONEGATE-TOKEN: `cat token.txt`" \
    --header "X-ONEGATE-VMID: $VMID"
{
    "VM": {
        "ID": ...,
        "NAME": ...,
        "TEMPLATE": {
            "NIC": [
                {
                    "IP": ...,
                    "IP6_LINK": ...,
                    "MAC": ...,
                    "NETWORK": ...,
                },
                // more nics ...
            ]
        },
        "USER_TEMPLATE": {
            "ROLE_NAME": ...,
            "SERVICE_ID": ...,
            // more user template attributes
        }
    }
}
```

- `PUT ${ONEGATE_ENDPOINT}/vm`: To add information to the template of the current VM. The new information is placed inside the VM's user template section. This means that the application metrics are visible from the command line, Sunstone, or the APIs, and can be used to trigger OneFlow elasticity rules.

- `PUT ${ONEGATE_ENDPOINT}/vms/${VM_ID}`: To add information to the template of a specific VM of the Service.

```
$ curl -X "PUT" "${ONEGATE_ENDPOINT}/vm" \
    --header "X-ONEGATE-TOKEN: `cat token.txt`" \
    --header "X-ONEGATE-VMID: $VMID" \
    -d "APP_LOAD = 9.7"
```

The new metric is stored in the user template section of the VM:

```
$ onevm show 0
...
USER TEMPLATE
APP_LOAD="9.7"
```

• `GET ${ONEGATE_ENDPOINT}/service`: To request information about the Service. The information is returned in JSON format and is ready for public cloud usage. By pushing data `PUT /vm` from one VM and pulling the Service data from another VM `GET /service`, nodes that are part of a OneFlow Service can pass values from one to another.

```
$ curl -X "GET" "${ONEGATE_ENDPOINT}/service" \
    --header "X-ONEGATE-TOKEN: `cat token.txt`" \
    --header "X-ONEGATE-VMID: $VMID"

{
    "SERVICE": {
        "id": ...,
        "name": ...,
        "roles": [
            {
                "name": ...,
                "cardinality": ...,
                "state": ...,
                "nodes": [
                    {
                        "deploy_id": ...,
                        "running": true|false,
                        "vm_info": {
                            // VM template as return by GET /VM
                        }

                    },
                    // more nodes ...
                ]
            },
            // more roles ...
        ]
    }
}
```

• `GET ${ONEGATE_ENDPOINT}`: returns information endpoints:

```
$ curl -X "GET" "${ONEGATE_ENDPOINT}/service" \
    --header "X-ONEGATE-TOKEN: `cat token.txt`" \
    --header "X-ONEGATE-VMID: $VMID"

{
    "vm_info": "http://<onegate_endpoint>/vm",
    "service_info": "http://<onegate_endpoint>/service"
}
```

### Self-configuration

• `PUT ${ONEGATE_ENDPOINT}/service/role/${ROLE_NAME}`: To change the cardinality of a specific role of the Service:

---

```
$ curl -X "PUT" "${ONEGATE_ENDPOINT}/service/role/worker" \
    --header "X-ONEGATE-TOKEN: `cat token.txt`" \
    --header "X-ONEGATE-VMID: $VMID" \
    -d "{'cardinality' : 10}"
```

- POST ${ONEGATE_ENDPOINT}/vms/${VM_ID}/action: To perform an action on a specific VM of the Service. Supported actions (resume, stop, suspend, terminate, reboot, poweroff, resched, unresched, hold, release)

```
$ curl -X "POST" "${ONEGATE_ENDPOINT}/vms/18/action" \
    --header "X-ONEGATE-TOKEN: `cat token.txt`" \
    --header "X-ONEGATE-VMID: $VMID" \
    -d "{'action' : {'perform': 'resched'}}"
```

### 5.3.4 Sample Application Monitoring Script

```bash
1   #!/bin/bash
2
3   # -------------------------------------------------------------------------- #
4   # Copyright 2002-2016, OpenNebula Project, OpenNebula Systems          #
5   #                                                                      #
6   # Licensed under the Apache License, Version 2.0 (the "License"); you may   #
7   # not use this file except in compliance with the License. You may obtain   #
8   # a copy of the License at                                             #
9   #                                                                      #
10  # http://www.apache.org/licenses/LICENSE-2.0                           #
11  #                                                                      #
12  # Unless required by applicable law or agreed to in writing, software   #
13  # distributed under the License is distributed on an "AS IS" BASIS,     #
14  # WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied.  #
15  # See the License for the specific language governing permissions and   #
16  # limitations under the License.                                       #
17  #-------------------------------------------------------------------------- #
18
19  ############################################################################
20  # Initialization
21  ############################################################################
22
23  ERROR=0
24
25  if [ -z $ONEGATE_TOKEN ]; then
26      echo "ONEGATE_TOKEN env variable must point to the token.txt file"
27      ERROR=1
28  fi
29
30  if [ -z $ONEGATE_ENDPOINT ]; then
31      echo "ONEGATE_ENDPOINT env variable must be set"
32      ERROR=1
33  fi
34
35  if [ $ERROR = 1 ]; then
36      exit -1
37  fi
38
39  TMP_DIR=`mktemp -d`
```

(continues on next page)

```
40    echo "" > $TMP_DIR/metrics
41
42    ############################################################################
43    # Memory metrics
44    ############################################################################
45
46    MEM_TOTAL=`grep MemTotal: /proc/meminfo | awk '{print $2}'`
47    MEM_FREE=`grep MemFree: /proc/meminfo | awk '{print $2}'`
48    MEM_USED=$(($MEM_TOTAL-$MEM_FREE))
49
50    MEM_USED_PERC="0"
51
52    if ! [ -z $MEM_TOTAL ] && [ $MEM_TOTAL -gt 0 ]; then
53        MEM_USED_PERC=`echo "$MEM_USED $MEM_TOTAL" | \
54            awk '{ printf "%.2f", 100 * $1 / $2 }'`
55    fi
56
57    SWAP_TOTAL=`grep SwapTotal: /proc/meminfo | awk '{print $2}'`
58    SWAP_FREE=`grep SwapFree: /proc/meminfo | awk '{print $2}'`
59    SWAP_USED=$(($SWAP_TOTAL - $SWAP_FREE))
60
61    SWAP_USED_PERC="0"
62
63    if ! [ -z $SWAP_TOTAL ] && [ $SWAP_TOTAL -gt 0 ]; then
64        SWAP_USED_PERC=`echo "$SWAP_USED $SWAP_TOTAL" | \
65            awk '{ printf "%.2f", 100 * $1 / $2 }'`
66    fi
67
68
69    #echo "MEM_TOTAL = $MEM_TOTAL" >> $TMP_DIR/metrics
70    #echo "MEM_FREE = $MEM_FREE" >> $TMP_DIR/metrics
71    #echo "MEM_USED = $MEM_USED" >> $TMP_DIR/metrics
72    echo "MEM_USED_PERC = $MEM_USED_PERC" >> $TMP_DIR/metrics
73
74    #echo "SWAP_TOTAL = $SWAP_TOTAL" >> $TMP_DIR/metrics
75    #echo "SWAP_FREE = $SWAP_FREE" >> $TMP_DIR/metrics
76    #echo "SWAP_USED = $SWAP_USED" >> $TMP_DIR/metrics
77    echo "SWAP_USED_PERC = $SWAP_USED_PERC" >> $TMP_DIR/metrics
78
79    ############################################################################
80    # Disk metrics
81    ############################################################################
82
83    /bin/df -k -P | grep '^/dev' > $TMP_DIR/df
84
85    cat $TMP_DIR/df | while read line; do
86        NAME=`echo $line | awk '{print $1}' | awk -F '/' '{print $NF}'`
87
88        DISK_TOTAL=`echo $line | awk '{print $2}'`
89        DISK_USED=`echo $line | awk '{print $3}'`
90        DISK_FREE=`echo $line | awk '{print $4}'`
91
92        DISK_USED_PERC="0"
93
94        if ! [ -z $DISK_TOTAL ] && [ $DISK_TOTAL -gt 0 ]; then
95            DISK_USED_PERC=`echo "$DISK_USED $DISK_TOTAL" | \
96                awk '{ printf "%.2f", 100 * $1 / $2 }'`
```

```bash
 97        fi
 98
 99        #echo "DISK_TOTAL_$NAME = $DISK_TOTAL" >> $TMP_DIR/metrics
100        #echo "DISK_FREE_$NAME = $DISK_FREE" >> $TMP_DIR/metrics
101        #echo "DISK_USED_$NAME = $DISK_USED" >> $TMP_DIR/metrics
102        echo "DISK_USED_PERC_$NAME = $DISK_USED_PERC" >> $TMP_DIR/metrics
103    done
104
105    ###############################################################################
106    # PUT command
107    ###############################################################################
108
109    VMID=$(source /mnt/context.sh; echo $VMID)
110
111    curl -X "PUT" $ONEGATE_ENDPOINT/vm \
112        --header "X-ONEGATE-TOKEN: `cat $ONEGATE_TOKEN`" \
113        --header "X-ONEGATE-VMID: $VMID" \
114        --data-binary @$TMP_DIR/metrics
```

**PUBLIC CLOUD**

## 6.1 Overview

A Public Cloud is an **extension of a Private Cloud to expose RESTful Cloud interfaces**. Cloud interfaces can be added to your Private or Hybrid Cloud if you want to provide partners or external users with access to your infrastructure, or to sell your overcapacity. Obviously, a local cloud solution is the natural back-end for any public cloud.

The *EC2 Query subset* interfaces provide a simple and remote management of cloud (virtual) resources at a high abstraction level. There is no modification in the operation of OpenNebula to expose Cloud interfaces. Users can interface the infrastructure using any Private or Public Cloud interface.

### 6.1.1 How Should I Read This Chapter

Before reading this chapter make sure you have read the Deployment Guide.

Read the *EC2 Server Configuration* to understand how to start the EC2 API for OpenNebula. *OpenNebula EC2 User Guide* contains a reference of the supported commands and their usage.

After reading this chapter you can continue configuring more *Advanced Components*.

### 6.1.2 Hypervisor Compatibility

This Chapter applies both to KVM and vCenter.

## 6.2 EC2 Server Configuration

The OpenNebula EC2 Query is a web service that enables you to launch and manage virtual machines in your OpenNebula installation through the Amazon EC2 Query Interface. In this way, you can use any EC2 Query tool or utility to access your Private Cloud. The EC2 Query web service is implemented upon the **OpenNebula Cloud API** (OCA) layer that exposes the full capabilities of an OpenNebula private cloud; and Sinatra, a widely used light web framework.

The current implementation includes the basic routines to use a Cloud, namely: image upload and registration, and the VM run, describe and terminate operations. The following sections explain you how to install and configure the EC2 Query web service on top of a running OpenNebula cloud.

---

**Note:** The OpenNebula EC2 Query service provides a Amazon EC2 Query API compatible interface to your cloud, that can be used alongside the native OpenNebula CLI or OpenNebula Sunstone. The OpenNebula distribution includes the tools needed to use the EC2 Query service.

---

### 6.2.1 Requirements & Installation

You must have an OpenNebula site properly configured and running, be sure to check the OpenNebula Installation and Configuration Guides to set up your private cloud first. This guide also assumes that you are familiar with the configuration and use of OpenNebula.

The OpenNebula EC2 Query service was installed during the OpenNebula installation, and the dependencies of this service are installed when using the install_gems tool as explained in the installation guide

### 6.2.2 Configuration

The service is configured through the `/etc/one/econe.conf` file, where you can set up the basic operational parameters for the EC2 Query web service. The available options are:

**Server configuration**

- `tmpdir`: Directory to store temp files when uploading images
- `one_xmlrpc`: oned xmlrpc service, http://localhost:2633/RPC2
- `host`: Host where econe server will run
- `port`: Port where econe server will run
- `ssl_server`: URL for the EC2 service endpoint, when configured through a proxy

**Log**

- `debug_level`: Log debug level, `0 = ERROR`, `1 = WARNING`, `2 = INFO`, `3 = DEBUG`.

---

**Auth**

- `auth`: Authentication driver for incomming requests

- `core_auth`: Authentication driver to communicate with OpenNebula core. Check this guide for more information about the core_auth system

**File based templates**

- `use_file_templates`: Use former file based templates for instance types instead of OpenNebula templates

- `instance_types`: DEPRECATED The VM types for your cloud

**Resources**

- `describe_with_terminated_instances`: Include terminated instances in the describe_instances xml. When this parameter is enabled all the VMs in DONE state will be retrieved in each describe_instances action and then filtered. This can cause performance issues when the pool of VMs in DONE state is huge

- `terminated_instances_expiration_time`: Terminated VMs will be included in the list till the termination date + terminated_instances_expiration_time is reached

- `datastore_id`: Datastore in which the Images uploaded through EC2 will be allocated, by default 1

- `cluster_id`: Cluster associated with the EC2 resources, by default no Cluster is defined

**Elastic IP**

- `elasticips_vnet_id`: VirtualNetwork containing the elastic ips to be used with EC2. If no defined the Elastic IP functionality is disabled

- `associate_script`: Script to associate a public IP with a private IP arguments: elastic_ip private_ip vnet_template(base64_encoded)

- `disassociate_script`: Script to disassociate a public IP arguments: elastic_ip

**EBS**

- `ebs_fstype`: FSTYPE that will be used when creating new volumes (DATABLOCKs)

> **Warning:** The `:host` **must** be a FQDN, do not use IP's here.

### Cloud Users

The cloud users have to be created in OpenNebula by `oneadmin` using the `oneuser` utility. Once a user is registered in the system, using the same procedure as to create private cloud users, they can start using the system.

The users will authenticate using the Amazon EC2 procedure with `AWSAccessKeyId` their OpenNebula's user name and `AWSSecretAccessKey` their OpenNebula's hashed password.

The cloud administrator can limit the interfaces that these users can use to interact with OpenNebula by setting the driver `public` for them. Using that driver cloud users will not be able to interact with OpenNebula through Sunstone, CLI nor XML-RPC.

```
$ oneuser chauth cloud_user public
```

### Defining VM Types

You can define as many Virtual Machine types as you want, just:

- Create a new OpenNebula template for the new type and make it available for the users group. You can use restricted attributes and set permissions like any other OpenNebula resource. **You must include the EC2_INSTANCE_TYPE parameter inside the template definition**, otherwise the template will not be available to be used as an instance type in EC2.

```
# This is the content of the /tmp/m1.small file
NAME = "m1.small"
EC2_INSTANCE_TYPE = "m1.small"
CPU = 1
MEMORY = 1700
...
```

```
$ onetemplate create /tmp/m1.small
$ onetemplate chgrp m1.small users
$ onetemplate chmod m1.small 640
```

The template must include all the required information to instantiate a new virtual machine, such as network configuration, capacity, placement requirements, etc. This information will be used as a base template and will be merged with the information provided by the user.

The user will select an instance type along with the ami id, keypair and user data when creating a new instance. Therefore, **the template should not include the OS**, since it will be specified by the user with the selected AMI.

---

**Note:** The templates are processed by the EC2 server to include specific data for the instance.

---

## 6.2.3 Starting the Cloud Service

To start the EC2 Query service just issue the following command

```
$ econe-server start
```

You can find the econe server log file in `/var/log/one/econe-server.log`.

To stop the EC2 Query service:

```
$ econe-server stop
```

## 6.2.4 Advanced Configuration

### Enabling Keypair

In order to benefit from the Keypair functionality, the images that will be used by the econe users must be prepared to read the EC2_PUBLIC_KEY and EC2_USER_DATA from the CONTEXT disk. This can be easiliy achieved with the new contextualization packages, generating a new custom contextualization package like this one:

```
#!/bin/bash
echo "$EC2_PUBLIC_KEY" > /root/.ssh/authorized_keys
```

### Enabling Elastic IP Functionality

An Elastic IP address is associated with the user, not a particular instance, and the user controls that address until she chooses to release it. This way the user can remap his public IP addresses to any of his instances.

---

In order to enable this functionality you have to follow the following steps in order to create a VNET containing the elastic IPS

- Create a new Virtual Network as oneadmin containing the public IPs that will be controlled by the EC2 users. Each IP **must be placed in its own AR**:

```
NAME    = "ElasticIPs"

PHYDEV  = "eth0"
VLAN    = "YES"
VLAN_ID = 50
BRIDGE  = "brhm"

AR  = [IP=10.0.0.1, TYPE=IP4, SIZE=1]
AR  = [IP=10.0.0.2, TYPE=IP4, SIZE=1]
AR  = [IP=10.0.0.3, TYPE=IP4, SIZE=1]
AR  = [IP=10.0.0.4, TYPE=IP4, SIZE=1]

# Custom Attributes to be used in Context
GATEWAY = 130.10.0.1
```

```
$ onevnet create /tmp/fixed.vnet
ID: 8
```

This VNET will be managed by the oneadmin user, therefore `USE` permission for the ec2 users is not required

- Update the econe.conf file with the VNET ID:

```
:elastic_ips_vnet: 8
```

- Provide associate and disassociate scripts

The interaction with the infrastructure has been abstracted, therefore two scripts have to be provided by the cloud administrator in order to interact with each specific network configuration. These two scripts enable us to adapt this feature to different configurations and data centers.

These scripts are language agnostic and their path has to be specified in the econe configuration file:

```
:associate_script: /usr/bin/associate_ip.sh
:disassociate_script: /usr/bin/disassociate_ip.sh
```

The associate script will receive three arguments: **elastic_ip** to be associated; **private_ip** of the instance; **Virtual Network template** base64 encoded.

The disassociate script will receive three arguments: **elastic_ip** to be disassociated.

Scripts to interact with OpenFlow can be found in the following ecosystem project

### Using a Specific Group for EC2

It is recommended to create a new group to handle the ec2 cloud users:

```
$ onegroup create ec2
ID: 100
```

Create and add the users to the ec2 group (ID:100):

```
$ oneuser create clouduser my_password
ID: 12
$ oneuser chgrp 12 100
```

Also, you will have to create ACL rules so that the cloud users are able to deploy their VMs in the allowed hosts.

```
$ onehost list
  ID NAME            CLUSTER   RVM     ALLOCATED_CPU      ALLOCATED_MEM   STAT
   1 kvm1            -           2     110 / 200 (55%)   640M / 3.6G (17%)   on
   1 kvm2            -           2     110 / 200 (55%)   640M / 3.6G (17%)   on
   1 kvm3            -           2     110 / 200 (55%)   640M / 3.6G (17%)   on
```

These rules will allow users inside the ec2 group (ID:100) to deploy VMs in the hosts kvm01 (ID:0) and kvm03 (ID:3)

```
$ oneacl create "@100 HOST/#1 MANAGE"
$ oneacl create "@100 HOST/#3 MANAGE"
```

You **have to create a VNet network** using the `onevnet utility` with the IP's you want to lease to the VMs created with the EC2 Query service.

```
$ onevnet create /tmp/templates/vnet
ID: 12
```

Remember that you will have to add this VNet (ID:12) to the users group (ID:100) and give USE (640) permissions to the group in order to get leases from it.

```
$ onevnet chgrp 12 100
$ onevnet chmod 12 640
```

> **Warning:** You will have to update the NIC template, inside the `/etc/one/ec2query_templates` directory, in order to use this VNet ID

### Configuring a SSL Proxy

OpenNebula EC2 Query Service runs natively just on normal HTTP connections. If the extra security provided by SSL is needed, a proxy can be set up to handle the SSL connection that forwards the petition to the EC2 Query Service and takes back the answer to the client.

This set up needs:

- A server certificate for the SSL connections
- An HTTP proxy that understands SSL
- EC2Query Service configuration to accept petitions from the proxy

If you want to try out the SSL setup easily, you can find in the following lines an example to set a self-signed certificate to be used by a lighttpd configured to act as an HTTP proxy to a correctly configured EC2 Query Service.

Let's assume the server were the lighttpd proxy is going to be started is called `cloudserver.org`. Therefore, the steps are:

### 1. Snakeoil Server Certificate

We are going to generate a snakeoil certificate. If using an Ubuntu system follow the next steps (otherwise your mileage may vary, but not a lot):

- Install the `ssl-cert` package

```
$ sudo apt-get install ssl-cert
```

- Generate the certificate

```
$ sudo /usr/sbin/make-ssl-cert generate-default-snakeoil
```

- As we are using lighttpd, we need to append the private key with the certificate to obtain a server certificate valid to lighttpd

```
$ sudo cat /etc/ssl/private/ssl-cert-snakeoil.key /etc/ssl/certs/ssl-cert-snakeoil.
→pem > /etc/lighttpd/server.pem
```

### 2. lighttpd as a SSL HTTP Proxy

You will need to edit the `/etc/lighttpd/lighttpd.conf` configuration file and:

- Add the following modules (if not present already)

  - mod_access

  - mod_alias

  - mod_proxy

  - mod_accesslog

  - mod_compress

- Change the server port to 443 if you are going to run lighttpd as root, or any number above 1024 otherwise:

```
server.port                = 8443
```

- Add the proxy module section:

```
#### proxy module
## read proxy.txt for more info
proxy.server               = ( "" =>
                                 ("" =>
                                  (
                                    "host" => "127.0.0.1",
                                    "port" => 4567
                                  )
                                 )
                              )


#### SSL engine
ssl.engine                 = "enable"
ssl.pemfile                = "/etc/lighttpd/server.pem"
```

The host must be the server hostname of the computer running the EC2Query Service, and the port the one that the EC2Query Service is running on.

### 3. EC2Query Service Configuration

The `econe.conf` needs to define the following:

```
# Host and port where econe server will run
:host: localhost
:port: 4567

#SSL proxy URL that serves the API (set if is being used)
:ssl_server: https://cloudserver.org:8443/
```

Once the lighttpd server is started, EC2Query petitions using HTTPS uris can be directed to `https://cloudserver.org:8443`, that will then be unencrypted, passed to localhost, port 4567, satisfied (hopefully), encrypted again and then passed back to the client.

> **Warning:** Note that `:ssl_server` **must** be an URL that may contain a custom path.

## 6.3 OpenNebula EC2 Usage

The EC2 Query API offers the functionality exposed by Amazon EC2: upload images, register them, run, monitor and terminate instances, etc. In short, Query requests are HTTP or HTTPS requests that use the HTTP verb GET or POST and a Query parameter.

OpenNebula implements a subset of the EC2 Query interface, enabling the creation of public clouds managed by OpenNebula.

### 6.3.1 AMIs

- **upload image**: Uploads an image to OpenNebula
- **register image**: Register an image into OpenNebula
- **describe images**: Lists all registered images belonging to one particular user.

### 6.3.2 Instances

- **run instances**: Runs an instance of a particular image (that needs to be referenced).
- **describe instances**: Outputs a list of launched images belonging to one particular user.
- **terminate instances**: Shutdowns a set of virtual machines (or cancel, depending on its state).
- **reboot instances**: Reboots a set of virtual machines.
- **start instances**: Starts a set of virtual machines.
- **stop instances**: Stops a set of virtual machines.

### 6.3.3 EBS

- **create volume**: Creates a new DATABLOCK in OpenNebula
- **delete volume**: Deletes an existing DATABLOCK.

- **describe volumes**: Describe all available DATABLOCKs for this user

- **attach volume**: Attaches a DATABLOCK to an instance

- **detach volume**: Detaches a DATABLOCK from an instance

- **create snapshot**:

- **delete snapshot**:

- **describe snapshot**:

### 6.3.4 Elastic IPs

- **allocate address**: Allocates a new elastic IP address for the user

- **release address**: Releases a publicIP of the user

- **describe addresses**: Lists elastic IP addresses

- **associate address**: Associates a publicIP of the user with a given instance

- **disassociate address**: Disassociate a publicIP of the user currently associated with an instance

### 6.3.5 Keypairs

- **create keypair**: Creates the named keypair

- **delete keypair**: Deletes the named keypair, removes the associated keys

- **describe keypairs**: List and describe the key pairs available to the user

### 6.3.6 Tags

- **create-tags**

- **describe-tags**

- **remove-tags**

Commands description can be accessed from the Command Line Reference.

#### User Account Configuration

An account is needed in order to use the OpenNebula cloud. The cloud administrator will be responsible for assigning these accounts, which have a one to one correspondence with OpenNebula accounts, so all the cloud administrator has to do is check the *configuration guide to setup accounts*, and automatically the OpenNebula cloud account will be created.

In order to use such an account, the end user can make use of clients programmed to access the services described in the previous section. For this, she has to set up his environment, particularly the following aspects:

- **Authentication**: This can be achieved in three different ways, here listed in order of priority (i.e. values specified in the argument line supersede environmental variables)

    - Using the **commands arguments**. All the commands accept an **Access Key** (as the OpenNebula username) and a **Secret Key** (as the OpenNebula hashed password)

    - Using **EC2_ACCESS_KEY** and **EC2_SECRET_KEY** environment variables the same way as the arguments

– If none of the above is available, the **ONE_AUTH** variable will be checked for authentication (with the same used for OpenNebula CLI).

• **Server location**: The command need to know where the OpenNebula cloud service is running. That information needs to be stored within the **EC2_URL** environment variable (in the form of a http URL, including the port if it is not the standard 80).

> **Warning:** The `EC2_URL` has to use the FQDN of the EC2-Query Server

### Hello Cloud!

Lets take a walk through a typical usage scenario. In this brief scenario it will be shown how to upload an image to the OpenNebula image repository, how to register it in the OpenNebula cloud and perform operations upon it.

• **upload_image**

Assuming we have a working Gentoo installation residing in an **.img** file, we can upload it into the OpenNebula cloud using the **econe-upload** command:

```
$ econe-upload /images/gentoo.img
Success: ImageId ami-00000001
$ econe-register ami-00000001
Success: ImageId ami-00000001
```

• **describe_images**

We will need the **ImageId** to launch the image, so in case we forgotten we can list registered images using the **econe-describe-images** command:

```
$ econe-describe-images -H
Owner       ImageId       Status        Visibility   Location
--------------------------------------------------------------------------
helen       ami-00000001  available     private        ↳
→19ead5de585f43282acab4060bfb7a07
```

• **run_instance**

Once we recall the ImageId, we will need to use the **econe-run-instances** command to launch an Virtual Machine instance of our image:

```
$ econe-run-instances -H ami-00000001
Owner       ImageId               InstanceId InstanceType
--------------------------------------------------------------------------
helen       ami-00000001          i-15       m1.small
```

We will need the **InstanceId** to monitor and shutdown our instance, so we better write down that `i-15`.

• **describe_instances**

If we have too many instances launched and we don't remember everyone of them, we can ask **econe-describe-instances** to show us which instances we have submitted.

```
$ econe-describe-instances  -H
Owner       Id   ImageId       State         IP              Type
--------------------------------------------------------------------------
    ↳----------------------
helen       i-15  ami-00000001 pending       147.96.80.33    m1.small
```

We can see that the instances with Id i-15 has been launched, but it is still pending, i.e., it still needs to be deployed into a physical host. If we try the same command again after a short while, we should be seeing it running as in the following excerpt:

```
$ econe-describe-instances  -H
Owner       Id   ImageId      State        IP               Type
------------------------------------------------------------------------------
↪---------------------
helen       i-15  ami-00000001 running      147.96.80.33     m1.small
```

- **terminate_instances**

After we put the Virtual Machine to a good use, it is time to shut it down to make space for other Virtual Machines (and, presumably, to stop being billed for it). For that we can use the **econe-terminate-instances** passing to it as an argument the **InstanceId** that identifies our Virtual Machine:

```
$ econe-terminate-instances i-15
Success: Terminating i-15 in running state
```

---

**Note:** You can obtain more information on how to use the above commands accessing their Usage help passing them the **-h** flag

---

# MARKETPLACE

## 7.1 Overview

Sharing, provisioning and consuming cloud images is one of the main concerns when using Cloud. OpenNebula provides a simple way to create and integrate with a cloud image provider, called MarketPlaces. Think of them as external datastores.

A MarketPlace can be:

- **Public**: accessible universally by all OpenNebula's.

- **Private**: local within an organization and specific for a single OpenNebula (a single zone) or shared by a federation (a collection of zones).

A MarketPlace is a repository of MarketPlaceApps. A MarketPlaceApp can be thought of as an external Image optionally associated to a Virtual Machine Template.

Using MarketPlaces is very convenient, as it will allow you to move images across different kinds of datastores (using the MarketPlace as an exchange point), it is a way to share OpenNebula images in a Federation, as these resources are federated. In an OpenNebula deployment where the different VDCs don't share any resources, a MarketPlace will act like a shared datastore for all the users.

### 7.1.1 Supported Actions

MarketPlaces support various actions:

| Action | Description |
|---|---|
| create | Create a new MarketPlace. |
| monitor | This automatic action, discovers the available MarketPlaceApps and monitors the available space of the MarketPlace. |
| delete | Removes a MarketPlace from OpenNebula. For the Publick MarketPlace, it will also remove the MarketPlaceApps, but for any other type of MarketPlace this will not remove the MarketPlaceApps, and will only work if the MarketPlace is empty. |
| *other* | Generic actions common to OpenNebula resources are also available: update, chgrp, chown, chmod and rename. |

As for the MarketPlaceApps, they support these actions:

| Action | Description |
|--------|-------------|
| `create` | Upload a local image into the the MarketPlace. **NOTE:** This action can only be done with marketplaces associated with the current zone. |
| `export` | Export the MarketPlaceApp and download it into a local Datastore. |
| `delete` | Removes a MarketPlaceApp. |
| `download` | Downloads a MarketPlaceApp to a file. |
| *other* | Generic actions common to OpenNebula resources are also available: update, chgrp, chown, chmod, rename, enable and disable. |

> **Warning:** In order to use the `download` functionality make sure you read the Sunstone Advanced Guide.

### 7.1.2 Backends

MarketPlaces store the actual MarketPlaceApp images. How they do so depends on the backend driver. Currently these drivers are shipped with OpenNebula:

| Driver | Up-load | Description |
|--------|---------|-------------|
| *one* | No | This driver allows read access to the official public OpenNebula Systems Marketplace, as well as to the OpenNebula AppMarket Add-on. |
| *http* | Yes | When an image is uploaded to a MarketPlace of this kind, the image is written into a file in a specified folder, which is in turn available via a web-server. |
| *S3* | Yes | Images are stored into a S3 API-capable service. This means it can be stored in the official AWS S3 service , or in services that implement that API like Ceph Object Gateway S3 |
| *LXD* | No | Enables creating base images from the public LXD image repository |

OpenNebula ships with the OpenNebula Systems MarketPlace pre-registered, so users can access it directly.

### 7.1.3 Use Cases

Using the MarketPlace is recommended in many scenarios, to name a few:

- When starting with an empty OpenNebula, the public OpenNebula Systems Marketplace contains a catalog of OpenNebula-ready cloud images, allowing you to get on your feet very quickly.

- You can upload an image into a MarketPlace, and download it later on to another Datastores even if the source and target Datastores are of a different type, thus enabling image cloning from any datastore to any other datastore.

- In a federation, it is almost essential to have a shared MarketPlace in order to share MarketPlaceApps across zones.

- MarketPlaces are a great way to provide content for the users in VDCs with no initial virtual resources.

### 7.1.4 How Should I Read This Chapter

Before reading this chapter make sure you have read the Deployment Guide.

Read the *OpenNebula Systems MarketPlace* as it's global for all the OpenNebula installations. Then read the specific guide for the MarketPlace flavor you are interested in. Finally, read the *Managing MarketPlaceApps* to understand what operations you can perform on MarketPlaceApps.

After reading this chapter you can continue configuring more *Advanced Components*.
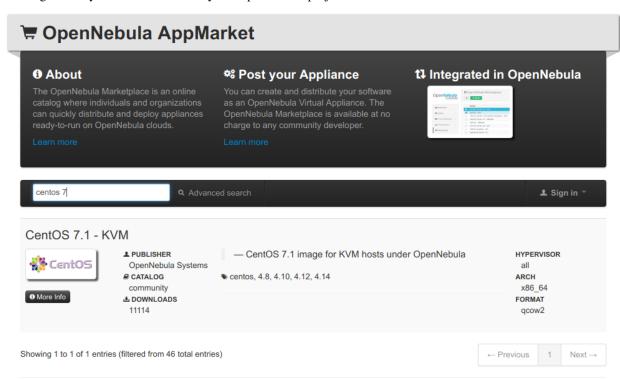
### 7.1.5 Hypervisor Compatibility

This chapter applies to all Hypervisors.

## 7.2 OpenNebula Systems MarketPlace

### 7.2.1 Overview

OpenNebula Systems provides a public and official MarketPlace, universally available to all the OpenNebula's. The OpenNebula Marketplace is a catalog of third party virtual appliances ready to run in OpenNebula environments. This MarketPlace is available here: http://marketplace.opennebula.systems. Anyone can request an account and upload their appliances and share them with other OpenNebula's, however, as opposed to other MarketPlaces, MarketPlaceApp creation is not done through OpenNebula, but by following the instructions in http://marketplace.opennebula.systems. Deletion of MarketPlaceApps is likewise limited.

The MarketPlaceApps included in the official MarketPlace are third-party contributions of other OpenNebula users, meaning that they are not certificated by the OpenNebula project.



You can also connect to MarketPlaces deployed with the OpenNebula Add-on AppMarket. The already deployed AppMarkets can still be used, but they are now deprecated in favor of the *HTTP MarketPlaces*.

## 7.2.2 Requirements

The url http://marketplace.opennebula.systems must be reachable from the OpenNebula Frontend.

## 7.2.3 Configuration

The Official OpenNebula Systems Marketplace is pre-registered in OpenNebula:

```
$ onemarket list
ID NAME                                     SIZE AVAIL      APPS MAD     ZONE
 0 OpenNebula Public                          0M -            46 one        0
```

Therefore it does not require any additional action from the administrator.

However, to connect to OpenNebula Add-on AppMarkets, it is possible to do so by creating a new MarketPlace template with the following attributes:

| Attribute | Description |
|---|---|
| NAME | Required |
| MARKET_MAD | Must be one. |
| ENDPOINT | (**Required to connect to AppMarket**) The URL of AppMarket. |

For example, the following examples illustrates the creation of a MarketPlace:

```
$ cat market.conf
NAME = PrivateMarket
MARKET_MAD = one
ENDPOINT = "http://privatemarket.opennebula.org"

$ onemarket create market.conf
ID: 100
```

## 7.2.4 Tuning & Extending

System administrators and integrators are encouraged to modify these drivers in order to integrate them with their datacenter. Please refer to the Market Driver Development guide to learn about the driver details.

## 7.3 HTTP MarketPlace

### 7.3.1 Overview

MarketPlaces of *HTTP* make use of a conventional HTTP server to expose the images (MarketPlaceApps) uploaded to a MarketPlace of this kind. The image will be placed in a specific directory that must be configured to be exposed by HTTP.

This is a fully supported MarketPlace with all the implemented features.

### 7.3.2 Requirements

A web-server should be deployed either in the Frontend or in a node reachable by the Frontend. A directory that will be used to store the uploaded images (MarketPlaceApps) should be configured to have the desired available space, and the web-server must be configured in order to grant HTTP access to that directory.

It is recommended to use either Apache or NGINX as they are known to work properly with the potentially large size of the MarketPlaceApp files. However, other web servers may work as long as they're capable to handle the load.

The web-server should be deployed by the administrator before registering the MarketPlace.

### 7.3.3 Configuration

These are the configuration attributes of a MarketPlace template of the *HTTP* kind.

| Attribute | Description |
|---|---|
| NAME | Required |
| MARKET_MAD | Must be `http` |
| PUBLIC_DIR | (**Required**) Absolute directory path to place images, the document root for http server, in the Frontend or in the hosts pointed at by the `BRIDGE_LIST` directive. |
| BASE_URL | (**Required**) URL base to generate MarketPlaceApp endpoints. |
| BRIDGE_LIST | (Optional) Comma separated list of servers to access the public directory. If not defined, public directory will be local to the Frontend. |

For example, the following examples illustrates the creation of an MarketPlace:

```
$ cat market.conf
NAME = PrivateMarket
MARKET_MAD = http
BASE_URL = "http://frontend.opennebula.org/"
PUBLIC_DIR = "/var/loca/market-http"
BRIDGE_LIST = "web-server.opennebula.org"

$ onemarket create market.conf
ID: 100
```

### 7.3.4 Tuning & Extending

System administrators and integrators are encouraged to modify these drivers in order to integrate them with their datacenter. Please refer to the Market Driver Development guide to learn about the driver details.

## 7.4 S3 MarketPlace

### 7.4.1 Overview

This MarketPlace uses an S3 API-capable service as the backend. This means MarketPlaceApp images will be stored in the official AWS S3 service , or in services that implement that API like Ceph Object Gateway S3.

### 7.4.2 Limitations

Since the S3 API does not provide the available space, this space is hard-coded into the driver file, limiting it to 1TB. See below to learn how to change the default value.

### 7.4.3 Requirements

To use this driver you require access to an S3 API-capable service.

- If you want to use AWS Amazon S3, you can start with the Getting Started guide.
- For Ceph S3 you must follow the Configuring Ceph Object Gateway guide.

Make sure you obtain both an *access_key* and a *secret_key* of a user that has access to a bucket with the exclusive purpose of storing MarketPlaceApp images.

### 7.4.4 Configuration

These are the configuration attributes of a MarketPlace template of the *S3* kind:

| Attribute | Description |
|---|---|
| NAME | Required |
| MARKET_MAD | Must be `s3` |
| ACCESS_KEY_ | (**Required**) The access key of the S3 user. |
| SECRET_ACCESS | (**Required**) The secret key of the S3 user. |
| BUCKET | (**Required**) The bucket where the files will be stored. |
| REGION | (**Required**) The region to connect to. If you are using Ceph-S3 any value here will work. |
| ENDPOINT | (Optional) This is only required if you are connecting to a service other than Amazon AWS S3 service. Preferably don't use an endpoint that includes the bucket as leading part of the host's url. |
| SIGNATURE_V | (Optional) Leave blank for Amazon AWS S3 service. If connecting to Ceph S3 it **must** be `s3`. |
| FORCE_PATH_ | (Optional) Leave blank for Amazon AWS S3 service. If connecting to Ceph S3 it **must** be YES. |
| TOTAL_MB | (Optional) This parameter defines the Total size of the MarketPlace in MB. It defaults to `1024` GB. |
| READ_LENGTH | (Optional) Split the file into chunks of this size (in MB). You should **never** user a quantity larger than *100*. Defaults to *32* (MB). |

For example, the following examples illustrates the creation of an MarketPlace:

```
$ cat market.conf
NAME=S3CephMarket
ACCESS_KEY_ID="I0PJDPCIYZ665MW88W9R"
SECRET_ACCESS_KEY="dxaXZ8U90SXydYzyS5ivamEP20hkLSUViiaR"
BUCKET="opennebula-market"
ENDPOINT="http://ceph-gw.opennebula.org"
FORCE_PATH_STYLE="YES"
MARKET_MAD=s3
REGION="default"
SIGNATURE_VERSION=s3

$ onemarket create market.conf
ID: 100
```

### 7.4.5 Tuning & Extending

In order to change the available size of the MarketPlace from 1TB to your desired value, you can modify */var/lib/one/remotes/market/s3/monitor* and change:

```
TOTAL_MB_DEFAULT = 1048576 # Default maximum 1TB
```

System administrators and integrators are encouraged to modify these drivers in order to integrate them with their datacenter. Please refer to the Market Driver Development guide to learn about the driver details.

## 7.5 Linux Containers MarketPlace

### 7.5.1 Overview

Linux Containers image server hosts a public image server for use by LXC and LXD. It is the default image server on LXD.

OpenNebula's linuxcontainers marketplace enables users to easily download, contextualize and add Linux Container's images to an OpenNebula's image datastore. Linux Containers images are compressed **.tar.xz** files. In order to use them, this marketplace creates an image, where it dumps the content and later uploads it to OpenNebula. The marketplace will automatically take care of downloading the correct context package for your image and installing it inside the container. The marketplace also creates a VM template with a set of required values and usability ones.

### 7.5.2 Requirements

- OpenNebula's frontend needs an Internet connection to https://images.linuxcontainers.org and https://github.com
- Aproximately 6GB of storage plus the container image size configured on your frontend

### 7.5.3 Configuration

Several parameters can be specified on the marketplace's template:

| Attribute | Description |
|---|---|
| NAME | Required |
| MARKET_MAD | Must be `linuxcontainers`. |
| ENDPOINT | The URL of the Market. Default: https://images.linuxcontainers.org |
| IMAGE_SIZE_MB | Size in MB for the image holding the rootfs. Default: 1024 |
| FILESYSTEM | Filesystem used for the image. Default: ext4 |
| FORMAT | Image format. Default: raw |

The OpenNebula frontend already ships with a configured LXD marketplace.

The following template will create a working marketplace with the default values:

```
NAME="Linux Containers"
MARKET_MAD="linuxcontainers"
DESCRIPTION="MarketPlace for the public image server fo LXC at linuxcontainers.org"
```

Save this contents on a file (e.g. lxcmarket) and create the market with:

```
$ onemarket create lxcmarket
```

You can also use Sunstone and input the previous values through the UI.

### 7.5.4 Tuning & Extending

System administrators and integrators are encouraged to modify these drivers in order to integrate them with their datacenter. Please refer to the Market Driver Development guide to learn about the driver details.

## 7.6 MarketPlaceApps Usage

### 7.6.1 Overview

As mentioned in the *MarketPlace Overview*, a MarketPlaceApp is a resource that can be either a single Image associated with a template, a template associated with one or more images, or a flow composed of one or more templates associated with images. In short, MarketPlaceApps are composed of metadata (a template) and binary data (the images). This guide addresses how you can manage these MarketPlaceApps in OpenNebula, considering both the metadata and the images.

MarketPlaceApps can be managed either using the CLI with the `onemarketapp` command or with the Sunstone GUI. In this section we will detail the actions available for MarketPlaceApps in both interfaces. MarketPlaceApps are common OpenNebula objects and respond the the common actions shared by all the OpenNebula objets: *list*, *show*, *create*, *delete*, etc, plus some specific ones.

### 7.6.2 Requirements
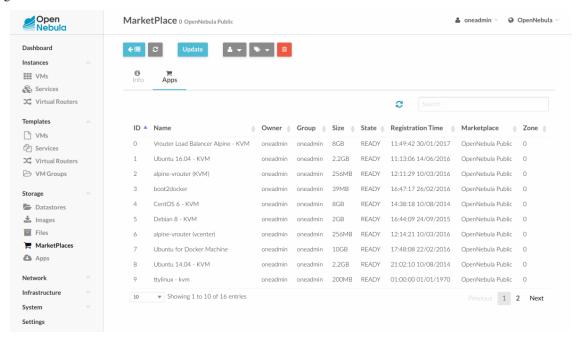
This section only apply to vCenter apps.

In order to use a vCenter app it is needed to attach the image to one vCenter VM Template which had been previously imported. An existing VM Template can be cloned and its disks replaced with the image from the marketplace. Once the VM Template its ready the appliance can be instantiated.

### 7.6.3 Listing MarketPlaceApps

Using the CLI:

```
$ onemarketapp list
 ID NAME                    VERSION  SIZE STAT TYPE  REGTIME          MARKET
  0 ttylinux - kvm              1.0   40M  rdy  img 06/08/46   OpenNebula Public
  1 ttylinux - VMware           1.1  102M  rdy  img 08/08/16   OpenNebula Public
  2 Carina Environment Manage   1.0  1.2G  rdy  img 05/14/26   OpenNebula Public
  3 Testing gUSE installation   1.0   16G  rdy  img 07/22/86   OpenNebula Public
  4 gUse v3.5.2               3.5.2   16G  rdy  img 04/02/43   OpenNebula Public
  5 Vyatta Core 6.5R1 - kvm     1.0    2G  rdy  img 07/22/86   OpenNebula Public
  6 gUSE CloudBroker Wrapper    1.0   16G  rdy  img 04/08/43   OpenNebula Public
  7 debian-7.1-amd64-kvm        1.3    5G  rdy  img 07/22/86   OpenNebula Public
  8 Hadoop 1.2 Master           1.0  1.3G  rdy  img 04/07/43   OpenNebula Public
  9 Hadoop 1.2 Slave            1.0  1.3G  rdy  img 05/18/14   OpenNebula Public
```

Using Sunstone:



### 7.6.4  Show a MarketPlaceApp

Using the CLI:

```
$ onemarketapp show 0
MARKETPLACE APP 0 INFORMATION
ID              : 0
NAME            : ttylinux - kvm
TYPE            : IMAGE
USER            : oneadmin
GROUP           : oneadmin
MARKETPLACE     : OpenNebula Public
STATE           : rdy


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


DETAILS
SOURCE          : http://marketplace.opennebula.systems//appliance/
→4fc76a938fb81d3517000003/download/0
MD5             : 04c7d00e88fa66d9aaa34d9cf8ad6aaa
PUBLISHER       : OpenNebula.org
PUB. DATE       : Wed Jun  8 22:17:19 137435166546
VERSION         : 1.0
DESCRIPTION     : This is a very small image that works with OpenNebula. It's already
→contextualized. The purpose of this image is to test OpenNebula deployments,
→without wasting network bandwith thanks to the tiny footprint of this image
(40MB).
SIZE            : 40M
ORIGIN_ID       : -1
```

```
FORMAT          : raw

IMPORT TEMPLATE


MARKETPLACE APP TEMPLATE
IMPORTED="YES"
IMPORT_ID="4fc76a938fb81d3517000003"
TAGS="linux, ttylinux,  4.8,  4.10"
VMTEMPLATE64=
↪"Q09OVEVYVCA9IFsgTkVUV09SSyAgPSJZRVMiLFNTSF9QVUJMSUNfS0VZICA9IiRVU0VSW1NTSF9QVUJMSUNfS0VZXSJdCgpDUI
↪"
```

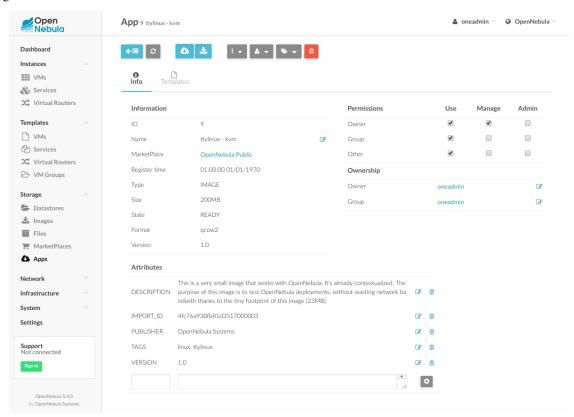Not that if we unpack that *VMTEMPLATE64* we obtain the following:

```
CONTEXT = [ NETWORK  ="YES",SSH_PUBLIC_KEY  ="$USER[SSH_PUBLIC_KEY]"]

CPU = "0.1"
GRAPHICS = [ LISTEN  ="0.0.0.0",TYPE  ="vnc"]

MEMORY = "128"
LOGO = "images/logos/linux.png"
```

Which demonstrates the capability of including a template into the appliance's data.

Using Sunstone:

### 7.6.5 Create a New MarketPlaceApp

In order to create a MarketPlaceApp you will need to prepare a new template file with the following attributes:

| Attribute | Description |
|---|---|
| NAME | Required |
| ORIGIN_ID | (**Required**) The ID of the source image. It must reference an available image and it must be in one of the supported datastores. |
| TYPE | (**Required**) Must be IMAGE. |
| MARKETPLACE_ID | (**Required**) The target marketplace ID. Alternatively you can specify the MARKETPLACE name. |
| MARKETPLACE | (**Required**) The target marketplace name. Alternatively you can specify the MARKETPLACE_ID name. |
| DESCRIPTION | (Optional) Text description of the MarketPlaceApp. |
| PUBLISHER | (Optional) If not provided, the username will be used. |
| VERSION | (Optional) A string indicating the MarketPlaceApp version. |
| VMTEMPLATE64 | (Optional) Creates this template (encoded in base64) pointing to the base image. |
| APPTEMPLATE64 | (Optional) This is the image template (encoded in base64) that will be added to the registered image. It is useful to include parameteres like DRIVER or DEV_PREFIX. |

Example:

```
$ cat marketapp.tpl
NAME=TTYlinux
ORIGIN_ID=0
TYPE=image

$ onemarketapp create marketapp.tpl -m "OpenNebula Public"
ID: 40
```
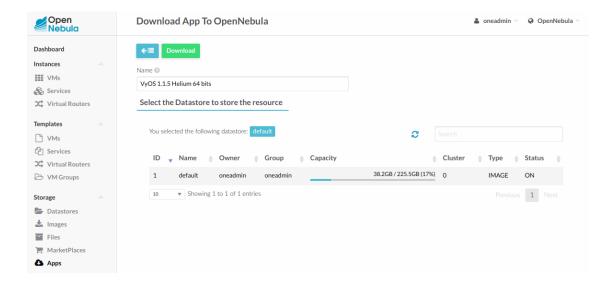
Using Sunstone:

### 7.6.6 Exporting a MarketPlaceApp

Using the CLI:

The command that exports the MarketPlaceApp is *onemarketapp export* which will return the ID of the new Image **and** the ID of the new associated template. If no template has been defined, it will return *-1*.

```
$ onemarketapp export 40 from_t1app -d 1
IMAGE
    ID: 1
VMTEMPLATE
    ID: -1
```

Using Sunstone:

## 7.6.7 Downloading a MarketPlaceApp

To download a MarketPlaceApp to a file:

```
$ onemarketapp download 40 /path/to/app
```

> **Warning:** This command requires that the *ONE_SUNSTONE* environment variable is set. Read here for more information.

> **Warning:** Make sure the Sunstone is properly deployed to handle this feature. Read here for more information.

## 7.6.8 Additional Commands

Like any other OpenNebula Resource, MarketPlaceApps respond to the base actions, namely:

- delete
- update
- chgrp
- chown
- chmod
- enable
- disable

Please take a look at the CLI reference to see how to use these actions. In Sunstone this options are also available.

### Tuning & Extending

System administrators and integrators are encouraged to modify these drivers in order to integrate them with their datacenter. Please refer to the Market Driver Development guide to learn about the driver details.

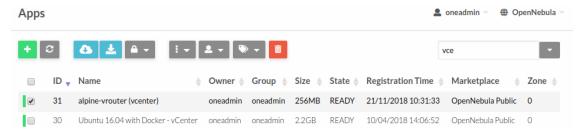## 7.7 Migrate images to/from KVM / vCenter DS

### 7.7.1 Overview

OpenNebula allows the management of hybrid environments, offering end-users a self-service portal to consume resources from both VMware-based infrastructures and KVM based ones to the user in a transparent way.

We are going to show you both transition below step-by-step trough Sunstone:
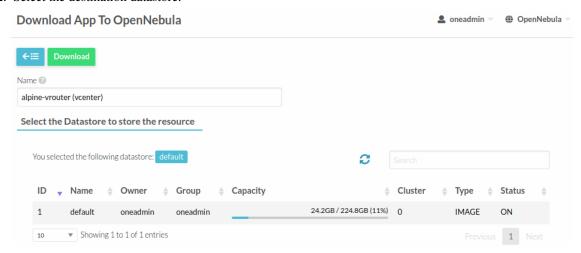
### 7.7.2 VMDK Image to QCOW2 Datastore

Here is how it works. We have a vmdk image within MarketPlace and we want to use it in KVM:

1. Go to MarketPlace and select the image, then click on **Download Button**.



> **Warning:** When the destination image datastore is qcow2/raw you **must** define the attribute `DRIVER=qcow2` or `DRIVER=raw`, respectively, in order to convert the image. If not, the image will be download without any change.

2. Select the destination datastore.



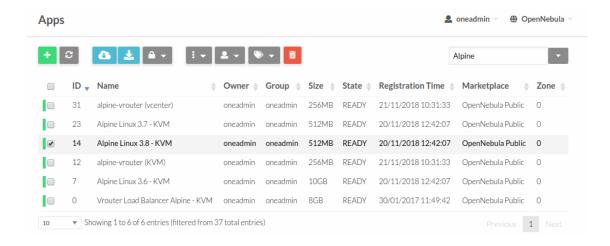3. Create a template in order to use the new image.

4. Instantiate the template and check that it works.
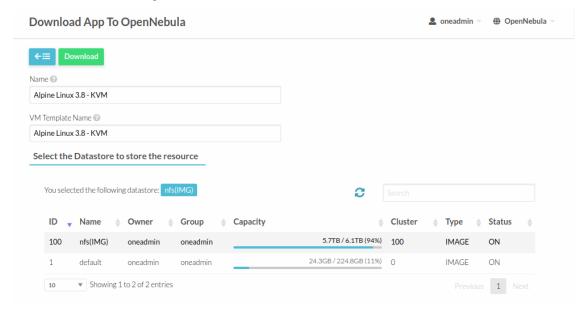


### 7.7.3 QCOW2 Image to VMDK Datastore

The process is very similar to the one described above.

1. Go to MarketPlace and select the image in qcow2 format to be used in a vCenter cluster, then click on **Download Button**.
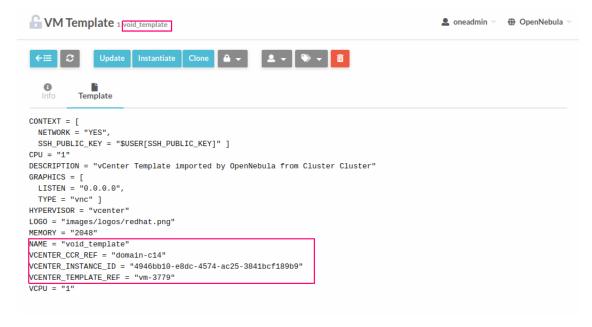
---

**Note:** In this case, when you import a vcenter datastore is automatically set `DRIVER=vcenter` so we dont need to define **DRIVER** attribute.
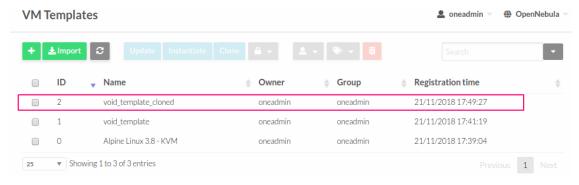
---

2. Select the destination image datastore.



3. When we download a vmdk image from the marketplace, a template is automatically created along with the image. However, we need a template with a valid vcenter ref for your cloud. You need to define an empty template in vcenter and import it in OpenNebula.
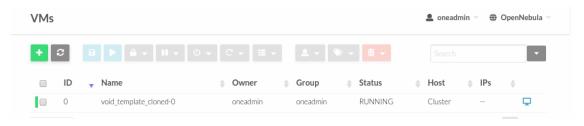
4. Now, clone the empty template to make use of the downloaded image.



5. Attach the image to the cloned template, so we can keep the original for other VMs.

6. Finally, instantiate the template.

In vCenter:



### 7.7.4 How it was implemented

Everytime the image that we selected from MarketPlace is downloaded to the frontend. Then, when the download process is finished, it is converted with the `qemu-img convert` tool as follows:

```
qemu-img convert -f <original_type> -O <destination_type> <original_file>
↪<destination_file>
```

Then, the file is sent to the destination datastore.

### 7.7.5 Limitations and restrictions

We have to take into account that when we convert an image from qcow2/raw to vmdk, the contextualization is lost so, we will have to install **VMWare tools** manually.

# APPLICATIONS CONTAINERIZATION

## 8.1 Overview

Application containerization is an OS-level virtualization method used to deploy and run distributed applications without launching an entire virtual machine (VM) for each application. OpenNebula supports the execution of multiple isolated applications within the VM instances, which means an instance could have multiple containers all sharing the same resources allocated to the running VM.

OpenNebula brings a built-in integration with Docker, the most common application containerization, to simplify the provision and management of Dockerized hosts. These virtualized Docker hosts are created using the OpenNebula Docker appliance (for KVM or vCenter) available on the OpenNebula Marketplace that should be previously downloaded and registered in the cloud datastore. OpenNebula provides cloud users with three approaches to use the Docker engine instances hosted by these virtualized Docker hosts.

The simpler approach is to directly instantiate and access the OpenNebula Docker image, and manage the dockerized hosts by using the OpenNebula GUI and CLI. OpenNebula also provides a driver for Docker Machine, which allows the remote provision and management of Docker hosts, and execution of Docker commands on the remote host from your Docker client.

### 8.1.1 How Should I Read This Chapter

We recommend you start reading the guide about how to *configure* and *use* the OpenNebula Docker appliance with the OpenNebula CLI and GUI, and continue with the *Docker Machine guide* only if you are interested in remote management of Docker hosts and clusters from your remote Docker client.

After reading this chapter you can continue configuring more *Advanced Components*.

### 8.1.2 Hypervisor Compatibility

This Chapter applies both to KVM and vCenter.

## 8.2 Docker Appliance Configuration

The Docker Appliance available on the OpenNebula marketplace brings a Docker Engine pre-installed and the contextualization packages configured to create Docker Hosts in a single click. Alternatively you can prepare your own image, using your favourite distribution, as long as it's supported by Docker Machine and it has the latest OpenNebula Contextualization packages.

The Docker Appliance is based on Ubuntu and have installed Docker, you can found more information about ths specific versions at the platform notes. In order to access the appliance once it's have been deployed it's necessary to update the template with the network and the password for the root user or the SSH key.

In order to prepare your cloud to serve Docker Engines please follow the next steps.

### 8.2.1 Step 1 - Create a Template

#### KVM

Download the appliance from the apps marketplace:

| | 34 | Ubuntu 16.04 with Docker - KVM | oneadmin | oneadmin | 2.2GB | READY | 10/04/2018 14:06:52 | OpenNebula Public | 0 |

When the appliance is downloaded a template with the same name it's created. It's recomended to update the template with a network for make the vm accessible from outside, set the disk size, the root password or the ssh key.

You can also create a template and include the appliance image in it.

Once the template is ready you can instantiate it.

#### vCenter

Download the appliance from the apps marketplace:

| | 33 | Ubuntu 16.04 with Docker - vCenter | oneadmin | oneadmin | 2.2GB | READY | 10/04/2018 14:06:52 | OpenNebula Public | 0 |

Create a vCenter template or update an existing one, the template must have a network for make the vm accessible from outside. Once the template is ready import both the template and the network into OpenNebula.

From OpenNebula update the template and attach the appliance disk to the vm. Also it's recommended to update the context of the template for set the root password or to include an SSH key.

### 8.2.2 Step 2 - Customize The Template

If you want to make any changes in the appliance and save them for latter use, you can set the image as persistent before launching the appliance. After the changes are performed, you need to shut the VM down and remove the persistent option from the image. This way you can create a golden image and new instances of the appliance won't overwrite it.



If you have already run the appliance as non persistent image you can take a look at Step 4 from the *Docker Application Usage* guide.

## 8.3 Docker Appliance Usage

In order to use the Docker Appliance it is needed to have the appliance configured at the OpenNebula installation. You can find more info at *Docker Appliance Configuration*.

### 8.3.1 Step 1 - Instantiate the template

First of all you need to identify the template. The default name is `Ubuntu 16.04 with Docker - KVM/ vCenter`.

Then you just need to instantiate it and wait for it to be in `running` state.

### 8.3.2 Step 2 - Running Hello World

Once the VM is running you can connect over SSH or VNC. You can run the docker "Hello world" and make sure every thing it's running well.

```
# ssh root@<vm_ip>
# root@ubuntu:~# docker run hello-world
  Unable to find image 'hello-world:latest' locally
  latest: Pulling from library/hello-world
  ca4f61b1923c: Pull complete
  Digest: sha256:97ce6fa4b6cdc0790cda65fe7290b74cfebd9fa0c9b8c38e979330d547d22ce1
  Status: Downloaded newer image for hello-world:latest


  Hello from Docker!
  This message shows that your installation appears to be working correctly.
```

Now you can see the "hello-world" image has been included in your images:

```
# root@ubuntu:~# docker images
  REPOSITORY          TAG                IMAGE ID            CREATED             SIZE
  hello-world         latest             f2a91732366c        4 months ago        1.
↪85kB
```

You can also see the container using the -a option (for show all the containers, including the ones that are not running):

```
# root@ubuntu:~# docker ps -a
  CONTAINER ID        IMAGE               COMMAND             CREATED             ␣
↪STATUS                    PORTS               NAMES
  3a0189f6b0af        hello-world         "/hello"            9 minutes ago       ␣
↪Exited (0) 9 minutes ago                       flamboyant_mirzakhani
```

### 8.3.3 Step 3 - Update the Docker Version

You can check your Docker version:

```
# root@ubuntu:~# docker version
  Client:
    Version:           18.03.0-ce
    API version:    1.37
    Go version:        go1.9.4
    Git commit:        0520e24
    Built:  Wed Mar 21 23:10:01 2018
    OS/Arch:           linux/amd64
    Experimental:    false
    Orchestrator:   swarm

  Server:
   Engine:
    Version:           18.03.0-ce
    API version:    1.37 (minimum version 1.12)
    Go version:        go1.9.4
    Git commit:        0520e24
    Built:  Wed Mar 21 23:08:31 2018
    OS/Arch:           linux/amd64
    Experimental:    false
```

And update it using the OS packages manager:

```
# root@ubuntu:~#apt-get update
# root@ubuntu:~#apt-get upgrade
```

### 8.3.4 Step 4 - Update a Docker Image

You can get an existing image and change it:

```
# root@ubuntu:~#docker run -i -t ubuntu /bin/bash
  Unable to find image 'ubuntu:latest' locally
  latest: Pulling from library/ubuntu
  a48c500ed24e: Pull complete
  1e1de00ff7e1: Pull complete
  0330ca45a200: Pull complete
  471db38bcfbf: Pull complete
  0b4aba487617: Pull complete
  Digest: sha256:c8c275751219dadad8fa56b3ac41ca6cb22219ff117ca98fe82b42f24e1ba64e
  Status: Downloaded newer image for ubuntu:latest
# root@0ac23d115db8:/# apt-get update
# root@0ac23d115db8:/# apt-get install ruby-full
# root@ubuntu:~#docker commit 0ac23d115db8  one/ubuntu-with-ruby
  sha256:eefdc54faeb5bafebd27012520a55b70c6818808997be2986d16b85b6c6f56e2
# root@ubuntu:~#docker image ls
  REPOSITORY             TAG               IMAGE ID           CREATED            ⌐
→SIZE
  one/ubuntu-with-ruby   latest            eefdc54faeb5       22 seconds ago     ⌐
→79.6MB
```

### 8.3.5 Step 5 - Save the Image

If you want to save changes like the ones performed in Step 3 and Step 4, the disk saveas functionality can be used to save this image as a new one. This option is available at the storage tab of the VM, this will automatically create a new image with the performed changes.

# 8.4 Docker Hosts Provision with Docker Machine

## 8.4.1 Introduction

This guide shows how to provision and manage remote Docker Hosts with Docker Machine on your OpenNebula cloud.

## 8.4.2 Prerequisites

To follow this guide, you will need the following (see the Platform Notes to see the version certified):

- Access to a fully working OpenNebula cloud running version 5.6 or later. You can check this by using any OpenNebula CLI command without parameters.
- A client computer with Docker CLI (the daemon is not required) and Docker Machine installed.
- Your OpenNebula Cloud must be accessible from your client computer.
- The OpenNebula Docker application available in the Marketplace should have been imported into the OpenNebula cloud, you can find more information at *Docker Appliance Configuration*.

This guide shows how to specify all the required command line attributes to create the Docker Engines. As an alternative you can specify a template registered in OpenNebula, in this case you can see all the available options at *Docker Machine Driver References* section.

## 8.4.3 Step 1 - Install Docker Machine OpenNebula Driver

In order to install the Docker Machine OpenNebula Driver you just need to install the package *docker-machine-opennebula* available here in your desktop.

In case you already have it installed in the OpenNebula frontend (or any other host) you instead can copy the *docker-machine-driver-opennebula* file from the /usr/share/docker_machine path of the frontend into any folder of your desktop that is included in your $PATH.

## 8.4.4 Step 2 - Configure Client Machine to Access the OpenNebula Cloud

It is assumed that you have a user with permissions to create / manage instances.

Set up env variables ONE_AUTH to contain user:password and ONE_XMLRPC to point to the OpenNebula cloud:

```
# export ONE_AUTH=~/.one/one_auth
# export ONE_XMLRPC=https://<ONE FRONTEND>:2633/RPC2
```

## 8.4.5 Step 3 - Use with vCenter

For vCenter hypervisor you will need to follow these steps to be able to use Docker Machine:

- Make sure you have a network defined in vCenter to connect Docker to.
- Create a template in vCenter with the desired capacity (CPU, Memory), a new hard disk (select the desired capacity) and new CD/DVD Drive (Datastore ISO File) with the ISO of the selected OS. Make sure you check Connect At Power On. Do not specify a network.

- In OpenNebula you will need to import the template and the desired networks. Make sure you make the network type ipv4.

### 8.4.6 Step 4 - Start your First Docker Host

To start your first Docker host you just need to use the *docker-machine create* command:

```
#docker-machine create --driver opennebula --opennebula-template-id $TEMPLATE_ID $VM_
↪NAME
```

This command creates a VM in OpenNebula using $TEMPLATE_ID as the template and $VM_NAME as the VM name.

Make sure the network attached to the template allows Docker Machine to connect to the VM.

If you want to create a VM without using a template (only for KVM) you can take a look at "Not Using a Template" section from *Docker Machine Driver References*.

### 8.4.7 Step 5 - Interact with your Docker Engine

You can list the VMs deployed by docker machine:

```
# docker-machine ls
  NAME            ACTIVE   DRIVER       STATE     URL                            SWARM   ␣
↪DOCKER       ERRORS
  ubuntu-docker   -        opennebula   Running   tcp://192.168.122.3:2376               ␣
↪v18.04.0-ce
```

Poweroff the remote host:

```
# docker-machine stop ubuntu-docker
  Stopping "ubuntu-docker"...
  Machine "ubuntu-docker" was stopped.
# docker-machine ls
  NAME            ACTIVE   DRIVER       STATE     URL     SWARM    DOCKER     ERRORS
  ubuntu-docker            opennebula   Timeout
```

Restart the remote host:

```
# docker-machine start ubuntu-docker
  Starting "ubuntu-docker"...
  (ubuntu-docker) Waiting for SSH..
  Machine "ubuntu-docker" was started.
  Waiting for SSH to be available...
  Detecting the provisioner...
# docker-machine ls
  NAME            ACTIVE   DRIVER       STATE     URL                            SWARM   ␣
↪DOCKER       ERRORS
  ubuntu-docker   -        opennebula   Running   tcp://192.168.122.3:2376               ␣
↪v18.04.0-ce
```

Remove the remote host (it will remove the VM from OpenNebula):

```
# docker-machine rm ubuntu-docker
  About to remove ubuntu-docker
  WARNING: This action will delete both local reference and remote instance.
```

(continues on next page)

```
Are you sure? (y/n): y
Successfully removed ubuntu-docker
```

Get more information about the host:

```
# docker-machine inspect ubuntu-docker
...
"EngineOptions": {
      "ArbitraryFlags": [],
      "Dns": null,
      "GraphDir": "",
      "Env": [],
      "Ipv6": false,
      "InsecureRegistry": [],
      "Labels": [],
      "LogLevel": "",
      "StorageDriver": "",
      "SelinuxEnabled": false,
      "TlsVerify": true,
      "RegistryMirror": [],
      "InstallURL": "https://get.docker.com"
   }
...
```

Get the IP address of the host:

```
# docker-machine ip ubuntu-docker
192.168.122.3
```

Connect to the host via SSH:

```
# docker-machine ssh ubuntu-docker
 $ docker ps -a
   CONTAINER ID        IMAGE               COMMAND             CREATED            ␣
→STATUS                       PORTS              NAMES
   787b15395f48        hello-world         "/hello"            16 seconds ago     ␣
→Exited (0) 15 seconds ago                      upbeat_bardeen
```

Activate the host, you can connect your Docker client to the remote host to run docker commands:

```
# eval $(docker-machine env ubuntu-docker)
# docker-machine ls
 NAME             ACTIVE    DRIVER        STATE     URL                              SWARM  ␣
→DOCKER        ERRORS
 ubuntu-docker    *         opennebula    Running   tcp://192.168.122.3:2376         ␣
→v18.04.0-ce
# docker ps -a
 CONTAINER ID        IMAGE               COMMAND             CREATED            ␣
→STATUS                       PORTS              NAMES
 787b15395f48        hello-world         "/hello"            6 minutes ago      ␣
→Exited (0) 6 minutes ago                       upbeat_bardeen
```

You can see how an "*" appears at the active field.

### 8.4.8 Containers Orchestration Platforms

### Swarm

Check the OpenNebula blog post to learn how to use Docker Swarm on an OpenNebula cloud.

### Swarmkit / Swarm mode

Check Docker documentation to use Swarmkit / Swarm mode. If you have discovery issues, please check your multicast support is OK.

As long as your VM template includes only one network, you should not even need to give –advertise-addr or –listen-addr

### Rancher

Check this OpenNebula blog post to learn how to use Rancher.

### Autoscaling via OneFlow

A service of Docker engines can be defined in *OneFlow*, and the autoscaling mechanisms of OneFlow used to automatically grow/decrease the number of Docker engines based on application metrics.

## 8.5 Docker Machine Driver Reference

### 8.5.1 Driver Options

- `--opennebula-user`: User identifier to authenticate with
- `--opennebula-password`: User password or token
- `--opennebula-xmlrpcurl`: XMLRPC endpoint
- `--opennebula-cpu`: CPU value for the VM
- `--opennebula-vcpu`: VCPUs for the VM
- `--opennebula-memory`: Size of memory for VM in MB
- `--opennebula-template-id`: Template ID to use
- `--opennebula-template-name`: Template to use
- `--opennebula-network-id`: Network ID to connect the machine to
- `--opennebula-network-name`: Network to connect the machine to
- `--opennebula-network-owner`: User ID of the Network to connect the machine to
- `--opennebula-image-id`: Image ID to use as the OS
- `--opennebula-image-name`: Image to use as the OS
- `--opennebula-image-owner`: Owner of the image to use as the OS
- `--opennebula-dev-prefix`: Dev prefix to use for the images: 'vd', 'sd', 'hd', etc. . .
- `--opennebula-disk-resize`: Size of disk for VM in MB
- `--opennebula-b2d-size`: Size of the Volatile disk in MB (only for b2d)

- `--opennebula-ssh-user`: Set the name of the SSH user

- `--opennebula-disable-vnc`: VNC is enabled by default. Disable it with this flag

- `--opennebula-start-retries`: number of retries to make for check if the VM is running, after each retry the driver sleeps for 2 seconds.

| CLI Option | Default Value | Environment Variable |
|---|---|---|
| *–opennebula-user* | | *ONE_USER* |
| *–opennebula-password* | | *ONE_PASSWORD* |
| *–opennebula-xmlrpcurl* | *http://localhost:2633/RPC2* | *ONE_XMLRPC* |
| *–opennebula-cpu* | *1* | *ONE_CPU* |
| *–opennebula-vcpu* | *1* | *ONE_VCPU* |
| *–opennebula-memory* | *1024* | *ONE_MEMORY* |
| *–opennebula-template-id* | | *ONE_TEMPLATE_ID* |
| *–opennebula-template-name* | | *ONE_TEMPLATE_NAME* |
| *–opennebula-network-id* | | *ONE_NETWORK_ID* |
| *–opennebula-network-name* | | *ONE_NETWORK_NAME* |
| *–opennebula-network-owner* | | *ONE_NETWORK_OWNER* |
| *–opennebula-image-id* | | *ONE_IMAGE_ID* |
| *–opennebula-image-name* | | *ONE_IMAGE_NAME* |
| *–opennebula-image-owner* | | *ONE_IMAGE_OWNER* |
| *–opennebula-dev-prefix* | | *ONE_IMAGE_DEV_PREFIX* |
| *–opennebula-disk-resize* | | *ONE_DISK_SIZE* |
| *–opennebula-b2d-size* | | *ONE_B2D_DATA_SIZE* |
| *–opennebula-ssh-user* | *docker* | *ONE_SSH_USER* |
| *–opennebula-disable-vnc* | Enabled | *ONE_DISABLE_VNC* |
| *–opennebula-start-retries* | 600 | *ONE_START_RETRIES* |

## 8.5.2 Using Templates

Using a VM template means specifying either *–opennebula-template-id* or *–opennebula-template-name*. If you specify either of these two options, the following table applies, indicating what incompatible and what overrideable parameters are available:

| Incompatible | Override |
|---|---|
| *–opennebula-image-id* | *–opennebula-cpu* |
| *–opennebula-image-name* | *–opennebula-vcpu* |
| *–opennebula-image-owner* | *–opennebula-memory* |
| *–opennebula-dev-prefix* | *–opennebula-network-id* |
| *–opennebula-disk-resize* | *–opennebula-network-name* |
| *–opennebula-b2d-size* | *–opennebula-network-owner* |
| *–opennebula-disable-vnc* | |

If you try to specify an attribute in the *incompatible* list, along with either *–opennebula-template-id* or *–opennebula-template-name*, then *docker-machine* will raise an error. If you specify an attribute in the *override* list, it will use that value instead of what is specified in the template.

The template must have a reference to an image, however, referencing a network is entirely optional. It the template has a network, the *–opennebula-network-\** options will override it, using the one in the template by default; if the template doesn't reference any networks, the *docker-machine* user **must** specify one.

```
# A template that references a network doesn't require any --opennebula-network-*
→attribute:
$ docker-machine create --driver opennebula --opennebula-template-id 10 mydockerengine

# However it can be overridden:
$ docker-machine create --driver opennebula --opennebula-template-id 10 --opennebula-
→network-id 2 mydockerengine
```

This is what the registered template in OpenNebula may look like:

```
NAME=b2d

CPU="1"
MEMORY="512"

# The OS Disl
DISK=[
IMAGE="b2d" ]

# The volatile disk (only for Boot2Docker)
DISK=[
FORMAT="raw",
SIZE="1024",
TYPE="fs" ]

# The network can be specified in the template or as a parameter
NIC=[
NETWORK="private" ]

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

Note that if there is a CONTEXT section in the template, it will be discarded and replaced with one by docker-machine.

### 8.5.3 Not Using Templates

if you don't specify neither `-opennebula-template-id` nor `--opennebula-template-name`, then you must specify the image: `--opennebula-image-*`, and the network: `--opennebula-network-*`, and optionally the other parameters.

# DISAGGREGATED DATA CENTERS

## 9.1 Overview

The aim of this advanced component is to provide the tools and methods needed to dynamically grow your private cloud infrastructure with physical resources running on remote bare-metal cloud providers.

Two of the use cases that will be supported by this new disaggregated cloud approach will be:

- **Edge Cloud Computing**. This approach will allow the transition from centralized clouds to distributed edge-like cloud environments. You will be able to grow your private cloud with resources at edge data center locations to meet latency and bandwidth needs of your workload.

- **Hybrid Cloud Computing**. This approach works as an alternative to the existing hybrid cloud drivers. So if there is a peak of demand and need for extra computing power you will be able to dynamically grow your underlying physical infrastructure. Compared with the use of hybrid drivers, this approach can be more efficient because it involves a single management layer. Also, it is a simpler approach because you can continue using the existing OpenNebula images and templates. Moreover, you always keep complete control over the infrastructure and avoid vendor lock-in.



There are several benefits of this approach over the traditional, more decoupled hybrid solution that involves using the provider cloud API. However, one of them stands tall among the rest and it is the ability to move offline workload between your local and rented resources.

### 9.1.1 How Should I Read This Chapter

You should be reading this Chapter as part of the *Advanced Components Guide*. You should be aware of the *Cloud Bursting* functionality, as the cluster provisioning shares some approaches.

In this Chapter, you can find a guide on how to provide additional resources from public bare-metal cloud providers into your OpenNebula. Cloud Administrators should follow the sections which cover *Installation* and *Basic Usage*, the Cloud Integrators and Developers might find useful the *Provision and Configuration Reference* and Provision Driver API specification.

After reading this chapter you can continue with other topics from *Advanced Components*.

## 9.2 OneProvision Installation

All functionality is distributed as an optional operating system package `opennebula-provision`, which must be installed on your **frontend alongside with the server packages**.

---

**Important:** The tool requires the Ansible to be installed on the same host(s) as well. There is **no automatic dependency** which would install Ansible automatically, you have to manage the installation of required Ansible version on your own.

Supported versions: Ansible 2.5.x or 2.6.x.

---

### 9.2.1 Step 1. Tools

Installation of tools is as easy as the install of any operating system package. Choose from the sections below based on your operating system. You also need to have installed the OpenNebula front-end packages.

#### CentOS/RHEL 7

```
$ sudo yum install opennebula-provision
```

#### Debian/Ubuntu

```
$ sudo apt-get install opennebula-provision
```

### 9.2.2 Step 2. Ansible

It's necessary to have Ansible installed. You can use a distribution package if there is a suitable version. Otherwise, you can install the required version via `pip` the following way:

#### CentOS/RHEL 7

```
$ sudo yum install python-pip
```

**Debian/Ubuntu**

```
$ sudo apt-get install python-pip
```

and, then install Ansible:

```
$ sudo pip install 'ansible>=2.5.0,<2.6.0'
```

Check the Ansible is installed properly:

```
$ ansible --version
ansible 2.5.3
  config file = /etc/ansible/ansible.cfg
  configured module search path = [u'/root/.ansible/plugins/modules', u'/usr/share/
↪ansible/plugins/modules']
  ansible python module location = /usr/lib/python2.7/site-packages/ansible
  executable location = /usr/bin/ansible
  python version = 2.7.5 (default, Apr 11 2018, 07:36:10) [GCC 4.8.5 20150623 (Red␣
↪Hat 4.8.5-28)]
```

# 9.3 OneProvision Basic Usage

Each new provision is described by the *provision template*, a YAML document specifying the OpenNebula resources to add (cluster, hosts, datastores, virtual networks), physical resources to provision from the remote infrastructure provider, the connection parameters for SSH and configuration steps (playbook) with tunables. The template is prepared by the experienced Cloud Administrator and passed to the command line tool `oneprovision`. At the end of the process, there is a new cluster available in the OpenNebula.



All operations with the provision and physical resources are performed only with the command line tool `oneprovision`. Create a new provision, manage (reboot, reset, power off, resume) the existing provisions, and delete the provision at the end.

In this chapter, we'll cover the basics of writing the provision templates and available commands to interact with the provision.

### 9.3.1 Provision Template

Provision template describes the resources to create in the OpenNebula (cluster, hosts, datastores, virtual networks), parameters for allocation of the new hosts on the remote bare-metal cloud provider, how to connect and configure them from the perspective of operating system and software. It's a YAML document, which needs to be prepared by the Cloud Administrator.

We'll explain the templating basics on a few simple examples. Continue to the *provision template reference* with comprehensive documentation.

#### Example 1: Empty cluster with datastores

```
---
name: example1

cluster:
  name: ex1-cluster

datastores:
  - name: ex1-default
    ds_mad: fs
    tm_mad: ssh
  - name: ex1-system
    type: system_ds
    tm_mad: ssh
    safe_dirs: '/var/tmp /tmp'
```

Execution of this provision will create new cluster `ex1-cluster` with datastores `ex1-default` and `ex1-system`. The cluster is always just a single one, datastores (hosts and virtual networks) are specified as a list (collection) of objects. Each object is described by a hash (associative array, map) of attributes, which would be otherwise specified in the OpenNebula INI-like template. I.e., it's an OpenNebula template represented as YAML hash.

---

**Note:** The system datastore `ex1-system` from the example matches the very same datastore which would be created over CLI and specified as OpenNebula INI-like template:

```
$ cat systemds.txt
NAME      = ex1-system
TYPE      = SYSTEM_DS
TM_MAD    = ssh
SAFE_DIRS = "/var/tmp /tmp"

$ onedatastore create systemds.txt
ID: 100
```

---

Check Datastores section in the Operation Guide for suitable attributes and values.

#### Example 2: Cluster with EC2 host

Following template describes a provision of a cluster with only single host deployed on Amazon EC2:

```
---
name: example2
```

```
cluster:
  name: ex2-cluster

hosts:
  - reserved_cpu: 100
    im_mad: kvm
    vm_mad: kvm
    provision:
      hostname: "ex2-host1"
      driver: ec2
      ec2_access: ********************
      ec2_secret: ****************************************
      region_name: "us-east-1"
      cloud_init: true
      ami: ami-66a7871c
      instancetype: "i3.metal"
      securitygroupsids: sg-*****************
      subnetid: subnet-*****************
```

As with the datastores in from *Example 1* above, the hosts are specified as a list as well. Each host is described by a hash with template attributes required by the OpenNebula. Parameters for the provisioning on the remote cloud providers must be set in own section `provision` of each host. The provision parameters are driver specific, you have to be aware of the available drivers and their parameters.

Check *Provision Drivers* reference for available drivers and parameters.

### Example 3: Host Configuration

The newly provisioned hosts are mostly a fresh installation without anything necessary for running the hypervisor. In this example, we add a few more parameters telling the OpenNebula how to connect and configure the new host:

```
---
name: example3
playbook: static_vxlan

cluster:
  name: ex3-cluster

hosts:
  - reserved_cpu: 100
    im_mad: kvm
    vm_mad: kvm
    provision:
      hostname: "ex3-host1"
      driver: ec2
      ec2_access: ********************
      ec2_secret: ****************************************
      region_name: "us-east-1"
      cloud_init: true
      ami: ami-66a7871c
      instancetype: "i3.metal"
      securitygroupsids: sg-*****************
      subnetid: subnet-*****************
    connection:
```

```
        remote_user: root
      configuration:
        opennebula_repository_version: 5.8.0
        opennebula_node_kvm_use_ev: true
        opennebula_node_kvm_param_nested: true
```

As part of provision creation, the new hosts are connected over SSH and required software is installed and configured. Custom SSH connection information can be set for each host in section `connection`. Installation is handled by the Ansible, which runs the template-global installation prescription called `playbook`. The playbook run can be slightly modified by optional `configuration` tunables.

Check the following subsections:

- *Playbooks* reference for available Ansible playbooks,

- *Roles* reference with a detailed description of individual roles and their configuration tunables.

## Example 4: Defaults

When deploying several hosts, repeating still same provision, configuration and connection parameters would be annoying and prone to errors.

In the following example, we explain how to use defaults:

```
---
name: example4
playbook: static_vxlan

defaults:
  provision:
    driver: ec2
    ec2_access: *******************
    ec2_secret: ****************************************
    region_name: "us-east-1"
    cloud_init: true
    ami: ami-66a7871c
    instancetype: "i3.metal"
    securitygroupsids: sg-*****************
    subnetid: subnet-*****************
  connection:
    remote_user: root
  configuration:
    opennebula_repository_version: 5.8.0
    opennebula_node_kvm_use_ev: true
    opennebula_node_kvm_param_nested: true

cluster:
  name: ex4-cluster

hosts:
  - reserved_cpu: 100
    im_mad: kvm
    vm_mad: kvm
    provision:
      hostname: "ex4-host1"
  - reserved_cpu: 100
    im_mad: kvm
```

```
    vm_mad: kvm
    provision:
      hostname: "ex4-host2"
      ami: ami-759bc50a
      cloud_init: false
    connection:
      remote_user: ubuntu
    configuration:
      opennebula_node_kvm_param_nested: false
```

Section `defaults` contains sub-sections for `provision`, `connection`, and `configuration` familiar from the previous examples. Defaults are applied to all objects, optionally you can override any of the parameters on the objects level. In the example, the first host `ex-host1` inherits all the **defaults** and extends them only with custom hostname. The second host `ex-host2` provides few more `provision`, `connection`, and `configuration` overrides (the rest defaults are taken untouched).

### Example 5: Full Cluster

Following example shows the provision of complete cluster with host, datastores, and networks.

```
---
name: example5
playbook: default

defaults:
  provision:
    driver: ec2
    ec2_access: ********************
    ec2_secret: ****************************************
    region_name: "us-east-1"
    cloud_init: true
    ami: ami-66a7871c
    instancetype: "i3.metal"
    securitygroupsids: sg-*****************
    subnetid: subnet-*****************
  connection:
    remote_user: root
  configuration:
    opennebula_node_kvm_manage_kvm: False
    opennebula_repository_version: 5.8.0
    opennebula_node_kvm_use_ev: true
    opennebula_node_kvm_param_nested: true

cluster:
  name: ex5-cluster

hosts:
  - reserved_cpu: 100
    im_mad: kvm
    vm_mad: kvm
    provision:
      hostname: "ex5-host1"

datastores:
  - name: ex5-default
```

```
      ds_mad: fs
      tm_mad: ssh
  - name: ex5-system
      type: system_ds
      tm_mad: ssh
      safe_dirs: '/var/tmp /tmp'

networks:
  - name: ex5-nat
      vn_mad: dummy
      bridge: br0
      dns: "8.8.8.8 8.8.4.4"
      gateway: "192.168.150.1"
      description: "Host-only networking with NAT"
      ar:
        - ip: "192.168.150.2"
          size: 253
          type: IP4
```

### Example 6: Template Inheritance

Similarly, as with **defaults** in *Example 4*, the reusable parts of the templates can be moved into their own templates. One provision template can include another provision template, extend or override the information from the included one. The template can directly extend only from one template, but several templates can be chained (for the recursive inheritance). Hosts, datastores, and networks sections are **merged** (appended) in the order they are defined and inherited, defaults are **deep merged** on the level of individual parameters.

In the following example, we separate datastores and networks definitions into own template `example-ds_vnets.yaml`:

```
---
datastores:
  - name: example-default
      ds_mad: fs
      tm_mad: ssh
  - name: example-system
      type: system_ds
      tm_mad: ssh
      safe_dirs: '/var/tmp /tmp'

networks:
  - name: example-nat
      vn_mad: dummy
      bridge: br0
      dns: "8.8.8.8 8.8.4.4"
      gateway: "192.168.150.1"
      description: "Host-only networking with NAT"
      ar:
        - ip: "192.168.150.2"
          size: 253
          type: IP4
```

Main template extends the datastores and network with one EC2 host:

```
---
name: example6
extends: example-ds_vnets.yaml

defaults:
  provision:
    driver: ec2
    ec2_access: ********************
    ec2_secret: ******************************************
    region_name: "us-east-1"
    cloud_init: true
    ami: ami-66a7871c
    instancetype: "i3.metal"
    securitygroupsids: sg-******************
    subnetid: subnet-******************
  connection:
    remote_user: root
  configuration:
    opennebula_node_kvm_manage_kvm: False
    opennebula_repository_version: 5.8.0
    opennebula_node_kvm_use_ev: true
    opennebula_node_kvm_param_nested: true

cluster:
  name: ex6-cluster

hosts:
  - reserved_cpu: 100
    im_mad: kvm
    vm_mad: kvm
    provision:
      hostname: "ex6-host1"
```

Check *Templates* reference for available base templates.

### 9.3.2 CLI Commands

This section covers available commands of the `oneprovision` tool.

> **Warning:** Commands should be run under `oneadmin` user on your frontend.

---

**Note:** Additional CLI arguments `--verbose/-d` and `--debug/-D` (applicable for all commands of `oneprovision` tool) provide additional levels of logging. Check *Logging Modes* for the detailed description.

---

#### Create

All deployment steps (create, provision, configuration) are covered by a single run of the command `oneprovision create`, it's necessary to provide *provision template* (with information about what to create, provision and how to configure the hosts). The OpenNebula provision ID is returned after successful provision.

Deployment of a new provision is a 4 steps process:

- **Add**. OpenNebula provision objects (cluster, hosts, datastores, networks) are created, but disabled for general use.

- *Provision*. Resources are allocated on the remote provider (e.g. use provider's API to get clean new hosts).

- *Configure*. Resources are reconfigured for the particular use (e.g. install virtualization tools on new hosts).

- **Enable**. Ready-to-use resources are enabled in the OpenNebula.

Parameters:

| Parameter | Description | Mandatory |
|---|---|---|
| `FILENAME` | File with *provision template* | **YES** |
| `--ping-retries number` | Number of SSH connection retries (default: 10) | NO |
| `--ping-timeout number` | Seconds between each SSH retry (default: 20) | NO |

Example:

```
$ oneprovision create myprovision.yaml -d
2018-11-27 11:32:03 INFO  : Creating provision objects
WARNING: This operation can take tens of minutes. Please be patient.
2018-11-27 11:32:05 INFO  : Deploying
2018-11-27 11:34:42 INFO  : Monitoring hosts
2018-11-27 11:34:46 INFO  : Checking working SSH connection
2018-11-27 11:34:49 INFO  : Configuring hosts
ID: 8fc831e6-9066-4c57-9ee4-4b11fea98f00
```

### Validate

The `validate` command checks the provided *provision template* is correct. Returns exit code 0 is template is valid.

Parameters:

| Parameter | Description | Mandatory |
|---|---|---|
| `FILENAME` | File with *provision template* | **YES** |
| `--dump` | Show complete provision template on standard output | NO |

Examples:

```
$ oneprovision validate simple.yaml
$ oneprovision validate simple.yaml --dump | head -4
---
name: myprovision
playbook: default
```

### List

The `list` command lists all provisions.

```
$ oneprovision list
                             ID NAME                         CLUSTERS HOSTS VNETS␣
↪DATASTORES STAT
8fc831e6-9066-4c57-9ee4-4b11fea98f00 myprovision                   1     1     1 ␣
↪       2 configured
```

### Show

The `show` command list all provisioned objects of the particular provision.

Parameters:

| Parameter | Description | Mandatory |
|---|---|---|
| `provision ID` | Valid provision ID | **YES** |
| `--csv` | Show output as CSV | NO |

Examples:

```
$ oneprovision show 8fc831e6-9066-4c57-9ee4-4b11fea98f00
PROVISION   INFORMATION
ID               : 8fc831e6-9066-4c57-9ee4-4b11fea98f00
NAME             : myprovision
STATUS           : configured

CLUSTERS
184

HOSTS
766

VNETS
135

DATASTORES
318
319
```

### Configure

> **Warning:** It's important to understand that the (re)configuration can happen only on physical hosts that aren't actively used by the users (e.g., no virtual machines running on the host) and with the operating systems/services configuration untouched since the last (re)configuration. It's not possible to (re)configure the host with manually modified OS/services configuration. It's not possible to fix a seriously broken host. Such situation needs to be handled manually by the experienced systems administrator.

The `configure` command offlines the OpenNebula hosts (making it unavailable for the users) and triggers again the deployment configuration phase. If the provision was already successfully configured before, the command line argument `--force` needs to be used. After successful configuration, the OpenNebula hosts are enabled back.

Parameters:

| Parameter | Description | Mandatory |
|---|---|---|
| `provision ID` | Valid provision ID | **YES** |
| `--force` | Force reconfiguration | NO |

Examples:

```
$ oneprovision configure 8fc831e6-9066-4c57-9ee4-4b11fea98f00 -d
ERROR: Hosts are already configured

$ oneprovision configure 8fc831e6-9066-4c57-9ee4-4b11fea98f00 -d --force
2018-11-27 12:43:31 INFO  : Checking working SSH connection
2018-11-27 12:43:34 INFO  : Configuring hosts
```

### Delete

The `delete` command releases the physical resources to the remote provider and deletes the provisioned OpenNebula objects.

```
$ oneprovision delete 8fc831e6-9066-4c57-9ee4-4b11fea98f00 -d
2018-11-27 12:45:21 INFO  : Deleting provision 8fc831e6-9066-4c57-9ee4-4b11fea98f00
2018-11-27 12:45:21 INFO  : Undeploying hosts
2018-11-27 12:45:23 INFO  : Deleting provision objects
```

Only provisions with no running VMs or images in the datastores can be easily deleted. You can force the `oneprovision` to terminate VMs running on provisioned hosts and delete all images in the datastores with `--cleanup` parameter.

Parameters:

| Parameter | Description | Mandatory |
|---|---|---|
| `provision ID` | Valid provision ID | **YES** |
| `--delete-all` | Delete all contained objects (VMs, images) | NO |

Examples:

```
$ oneprovision delete 8fc831e6-9066-4c57-9ee4-4b11fea98f00 -d
2018-11-27 13:44:40 INFO  : Deleting provision 8fc831e6-9066-4c57-9ee4-4b11fea98f00
ERROR: Provision with running VMs can't be deleted
```

```
$ oneprovision delete 8fc831e6-9066-4c57-9ee4-4b11fea98f00 -d --cleanup
2018-11-27 13:56:39 INFO  : Deleting provision 8fc831e6-9066-4c57-9ee4-4b11fea98f00
2018-11-27 13:56:44 INFO  : Undeploying hosts
2018-11-27 13:56:51 INFO  : Deleting provision objects
```

### Host Management

Individual hosts from the provision can be managed by the `oneprovision host` subcommands.

### List

The `host list` command lists all provisioned hosts, and `host top` command periodically refreshes the list until it's terminated.

```
$ oneprovision host list
  ID NAME            CLUSTER   RVM PROVIDER VM_MAD   STAT
 766 147.75.33.113   conf-prov   0 packet   kvm      on

$ oneprovision host top
```

**Host Power Off**

The `host poweroff` command offlines the host in the OpenNebula (making it unavailable for use by the users) and power off the physical resource.

```
$ oneprovision host poweroff 766 -d
2018-11-27 12:21:40 INFO  : Powering off host: 766
HOST 766: disabled
```

**Host Resume**

The `host resume` command power on the physical resource, and enables back the OpenNebula host (making it available again to the users).

```
$ oneprovision host resume 766 -d
2018-11-27 12:22:57 INFO  : Resuming host: 766
HOST 766: enabled
```

**Host Reboot**

The `host reboot` command offlines the OpenNebula host (making it unavailable for the users), cleanly reboots the physical resource and enables the OpenNebula host back (making it available again for the users after successful OpenNebula host monitoring).

```
$ oneprovision host reboot 766 -d
2018-11-27 12:25:10 INFO  : Rebooting host: 766
HOST 766: enabled
```

**Host Reset**

The `host reboot --hard` command offlines the OpenNebula host (making it unavailable for the users), resets the physical resource and enables the OpenNebula host back.

```
$ oneprovision host reboot --hard 766 -d
2018-11-27 12:27:55 INFO  : Resetting host: 766
HOST 766: enabled
```

**Host SSH**

The `host ssh` command opens the interactive SSH connection on the physical resource to the same (privileged) user used for the configuration.

```
$ oneprovision host ssh 766
Welcome to Ubuntu 18.04 LTS (GNU/Linux 4.15.0-20-generic x86_64)

 * Documentation:  https://help.ubuntu.com
 * Management:     https://landscape.canonical.com
 * Support:        https://ubuntu.com/advantage
```

```
Last login: Tue Nov 27 10:37:42 2018 from 213.175.39.66
root@myprovision-host1:~#
```

An additional argument may specify a command to run on the remote side.

```
$ oneprovision host ssh 766 hostname
ip-172-30-4-47.ec2.internal
```

### Host Configure

The physical host *configuration* is part of the initial deployment, but it's possible to trigger the reconfiguration on provisioned host anytime later (e.g. when a configured service stopped running, or the host needs to be reconfigured different way). Based on the initially provided connection and configuration parameters in the *provision template*, the configuration steps are applied again.

The `host configure` command offlines the OpenNebula host (making it unavailable for the users) and triggers again the deployment configuration phase. If provisioned the host was already successfully configured before, the command line argument `--force` needs to be used. After successful configuration, the OpenNebula host is enabled back.

```
$ oneprovision host configure 766 -d
ERROR: Hosts are already configured

$ oneprovision host configure 766 -d --force
2018-11-27 12:36:18 INFO  : Checking working SSH connection
2018-11-27 12:36:21 INFO  : Configuring hosts
HOST 766:
```

### Cluster Management

Individual clusters from the provision can be managed by the `oneprovision cluster` subcommands.

### Cluster List

The `oneprovision cluster list` command lists all provisioned clusters.

```
$ oneprovision cluster list
   ID NAME                   HOSTS VNETS DATASTORES
  184 myprovision                1     1          2
```

### Cluster Delete

The `oneprovision cluster delete` command deletes the cluster.

```
$ oneprovision cluster delete 184 -d
CLUSTER 184: deleted
```

The cluster needs to be without any datastores, virtual networks, or hosts. Please, check `oneprovision delete` command to remove all the related objects.

```
$ oneprovision cluster delete 184 -d
ERROR: [one.cluster.delete] Cannot delete cluster. Cluster 185 is not empty, it␣
→contains 1 datastores.
```

### Datastore Management

Individual datastores from the provision can be managed by the `oneprovision datastore` subcommands.

### Datastore List

The `oneprovision datastore list` command lists all provisioned datastores.

```
$ oneprovision datastore list
  ID NAME               SIZE AVAIL CLUSTERS     IMAGES TYPE DS      PROVIDER TM    ␣
→STA
 318 conf-provisio    271.1G 7%    184              0 img  fs      packet   ssh   ␣
→on
 319 conf-provisio       0M -      184              0 sys  -       packet   ssh   ␣
→on
```

### Datastore Delete

The `oneprovision datastore delete` command deletes the datastore.

```
$ oneprovision datastore delete 318 -d
2018-11-27 13:01:08 INFO  : Deleting datastore 318
DATASTORE 318: deleted
```

### Virtual Networks Management

Individual virtual networks from the provision can be managed by the `oneprovision vnet` subcommands.

### Vnet List

The `oneprovision vnet list` command lists all virtual networks.

```
$ oneprovision vnet list
  ID USER           GROUP       NAME                   CLUSTERS   BRIDGE   PROVIDER␣
→LEASES
 136 oneadmin       oneadmin    myprovision-hostonl 184           br0      packet    ␣
→  0
```

### Vnet Delete

The `oneprovision vnet delete` command deletes the virtual network.

```
$ oneprovision vnet delete 136 -d
2018-11-27 13:02:08 INFO  : Deleting vnet 136
VNET 136: deleted
```

### 9.3.3 Logging Modes

The `oneprovision` tool in the default mode returns only minimal requested output (e.g., provision IDs after create), or errors. The operations with the remote providers or the host configuration are complicated and time-consuming tasks. For the better insight and for debugging purposes there are 2 logging modes available providing more information on the standard error output.

- **verbose** (`--verbose`/`-d`). Only the main steps are logged.

Example:

```
$ oneprovision host reboot 766 -d
2018-11-27 12:58:32 INFO  : Rebooting host: 766
HOST 766: disabled
```

- **debug** (`--debug`/`-D`). All internal actions incl. generated configurations with **sensitive data** are logged.

Example:

```
$ oneprovision host reboot 766 -D
2018-11-27 12:59:02 DEBUG : Offlining OpenNebula host: 766
2018-11-27 12:59:02 INFO  : Rebooting host: 766
2018-11-27 12:59:02 DEBUG : Command run: /var/lib/one/remotes/pm/packet/reboot⌴
↪fa65c328-57c3-4890-831e-172c9d730b04 147.75.33.113 767 147.75.33.113
2018-11-27 12:59:09 DEBUG : Command succeeded
2018-11-27 12:59:09 DEBUG : Enabling OpenNebula host: 766
```

### 9.3.4 Running Modes

The `oneprovision` tool is ready to deal with common problems during the execution. It's able to retry some actions or clean up an uncomplete provision. Depending on where and how the tool is used, it offers 2 running modes:

- **interactive** (default). If the unexpected condition appears, the user is asked how to continue.

Example:

```
$ oneprovision host poweroff 0
ERROR: Driver action '/var/lib/one/remotes/pm/packet/shutdown' failed
Shutdown of Packet host 147.75.33.123 failed due to "{"errors"=>["Device must be⌴
↪powered on"]}"
1. quit
2. retry
3. skip
Choose failover method: 2
ERROR: Driver action '/var/lib/one/remotes/pm/packet/shutdown' failed
Shutdown of Packet host 147.75.33.123 failed due to "{"errors"=>["Device must be⌴
↪powered on"]}"
1. quit
2. retry
3. skip
Choose failover method: 1
$
```

- **batch** (`--batch`). It's expected to be run as part of the scripts. No question is raised to the user, but the tool tries to automatically deal with the problem according to the failover method specified as a command line parameter:

| Parameter | Description |
| --- | --- |
| `--fail-quit` | Set batch failover mode to quit (default) |
| `--fail-retry` number | Set batch failover mode to number of retries |
| `--fail-cleanup` | Set batch failover mode to clean up and quit |
| `--fail-skip` | Set batch failover mode to skip failing part |

Example of automatic retry:

```
$ oneprovision host poweroff 0 --batch --fail-retry 2
ERROR: Driver action '/var/lib/one/remotes/pm/packet/shutdown' failed
Shutdown of Packet host 147.75.33.123 failed due to "{"errors"=>["Device must be␣
↪powered on"]}"
ERROR: Driver action '/var/lib/one/remotes/pm/packet/shutdown' failed
Shutdown of Packet host 147.75.33.123 failed due to "{"errors"=>["Device must be␣
↪powered on"]}"
ERROR: Driver action '/var/lib/one/remotes/pm/packet/shutdown' failed
Shutdown of Packet host 147.75.33.123 failed due to "{"errors"=>["Device must be␣
↪powered on"]}"
```

Example of non-interactive provision with automatic clean up in case of failure:

```
$ oneprovision create simple.yaml -d --batch --fail-cleanup
2018-11-27 13:48:53 INFO  : Creating provision objects
WARNING: This operation can take tens of minutes. Please be patient.
2018-11-27 13:48:54 INFO  : Deploying
2018-11-27 13:51:32 INFO  : Monitoring hosts
2018-11-27 13:51:36 INFO  : Checking working SSH connection
2018-11-27 13:51:38 INFO  : Configuring hosts
2018-11-27 13:52:02 WARN  : Command FAILED (code=2): ANSIBLE_CONFIG=/tmp/d20181127-
↪11335-ktlqrb/ansible.cfg ansible-playbook --ssh-common-args='-o UserKnownHostsFile=/
↪dev/null' -i /tmp/d20181127-11335-ktlqrb/inventory -i /usr/share/one/oneprovision/
↪ansible/inventories/default/ /usr/share/one/oneprovision/ansible/default.yml
ERROR: Configuration failed
- 147.75.33.125   : TASK[opennebula-repository : Add OpenNebula repository (Ubuntu)] -
↪ MODULE FAILURE
2018-11-27 13:52:02 INFO  : Deleting provision 18e85ef4-b29f-4391-8d89-c72702ede54e
2018-11-27 13:52:02 INFO  : Undeploying hosts
2018-11-27 13:52:05 INFO  : Deleting provision objects
```

## 9.4 Provision and Configuration Reference

### 9.4.1 Provision Reference

#### Provision

The provision is a process of allocating new physical resources from the remote providers. *Provision drivers* are used for communication with the remote providers. Credentials for the communication and parameters of the required provision (hardware, operating system, IPs, etc.) need to be specified. All these information are stored in the *provision template* file and passed to the `oneprovision create` command.

In this chapter, we'll describe the format and content of the provision template.

### Template

**Provision template** is a YAML formatted file with parameters specifying the new physical resources to be provisioned. Contains:

- header (name, configuration playbook, parent template)
- global default parameters for
  - remote connection (SSH),
  - host provision driver,
  - host configuration tunables,
- list of provision objects (cluster, hosts, datastores, virtual networks) to deploy with overrides to the global defaults above.

Example:

```
---
name: myprovision
playbook: default

# Global defaults:
defaults:
  provision:
    driver: packet
    packet_token: ********************************
    packet_project: ************************************
    facility: ams1
    plan: baremetal_0
    os: centos7
  connection:
    public_key: '/var/lib/one/.ssh/id_rsa.pub'
    private_key: '/var/lib/one/.ssh/id_rsa'
  configuration:
    opennebula_node_kvm_param_nested: true

# List of provision objects to deploy with provision/connection/configuration
→overrides
cluster:
  name: mycluster

hosts:
  - reserved_cpu: 100
    im_mad: kvm
    vm_mad: kvm
    provision:
      hostname: "myhost1"
  - reserved_cpu: 100
    im_mad: kvm
    vm_mad: kvm
    provision:
      hostname: "myhost2"
      os: ubuntu18_04
    connection:
```

(continues on next page)

---

```
      remote_user: ubuntu

datastores:
  - name: "myprovision-images"
    ds_mad: fs
    tm_mad: ssh
  - name: "myprovision-system"
    type: system_ds
    tm_mad: ssh
    safe_dirs: "/var/tmp /tmp"

networks:
  - name: "myprovision-hostonly_nat"
    vn_mad: dummy
    bridge: br0
    dns: "8.8.8.8 8.8.4.4"
    gateway: "192.168.150.1"
    description: "Host-only networking with NAT"
    filter_ip_spoofing: "YES"
    filter_mac_spoofing: "YES"
    ar:
      - ip: 192.168.150.2
        size: 253
        type: IP4
```

### Header

| Parameter | Default | Description |
|---|---|---|
| name | none | Name of provision. |
| playbook | default | Ansible playbook used for hosts configuration. Provide the custom **absolute filename**, or one of predefined: <br>• *default* <br>• *default_lxd* <br>• *static_vxlan* |
| extends | none | Parent template to include and extend. Provide the custom **absolute filename**, or one of predefined: <br>• */usr/share/one/oneprovision/templates/default.yaml* <br>• */usr/share/one/oneprovision/templates/static_vxlan* |

### Shared sections

Following shared sections can be specified inside the template `defaults`, or directly inside each OpenNebula provision object (cluster, datastore, virtual network, and host). Parameters specified on the object side have higher priority and overrides the parameters from `defaults`.

### connection

This section contains parameters for the remote SSH connection on the privileged user or the user with escalation rights (via `sudo`) of the newly provisioned host(s).

| Parameter | Default | Description |
|---|---|---|
| `remote_user` | `root` | Remote user to connect via SSH. |
| `remote_port` | 22 | Remote SSH service port. |
| `public_key` | `/var/lib/one/.ssh/ddc/id_rsa.pub` | Path or content of the SSH public key. |
| `private_key` | `/var/lib/one/.ssh/ddc/id_rsa` | Path or content of the SSH private key. |

### provision

This section contains parameters for the provisioning driver. Most parameters are specific for each driver, the only valid common parameters are:

| Parameter | Default | Description |
|---|---|---|
| `driver` | none, needs to be specified | Host provision driver. Options:<br>• *packet*<br>• *ec2* |

### configuration

This section provides parameters for the host configuration process (e.g. KVM installation, host networking etc.). All parameters are passed to the external configuration tool (Ansible), all available parameters are covered by the *configuration* chapter.

### Provision objects

Sections `cluster`, `hosts`, `datastores`, `networks` contain list of provision objects to be deployed with all necessary parameters for the deployment and create in the OpenNebula. The object structure is a YAML representation of OpenNebula template with additional shared sections (`connection`, `provision`, `configuration`).

---

**Note:** It's possible to deploy only a single cluster, the section `cluster` is a dictionary. All other sections are lists.

---

Example of datastore defined from regular template:

```
$ cat ds.tpl
NAME="myprovision-images"
TM_MAD="ssh"
DS_MAD="fs"

$ onedatastore create ds.tpl
ID: 328
```

Example of the same datastore defined in provision template:

---

```
datastores:
  - name: "myprovision-images"
    ds_mad: fs
    tm_mad: ssh
```

### Templates

Several templates are shipped in the distribution package. Those are not the final templates, but only provide a partial definition of infrastructure and should be used (extended) as a base in own custom templates. Check the brief description of each template, and continue with reading the content of the template files in your installation.

### Template 'default'

---

**Note:** Installed into `/usr/share/one/oneprovision/templates/default.yaml`.

---

Template with private OpenNebula virtual network configured by *default* on physical hosts.

Following virtual network(s) are configured:

- nat

### Template 'static_vxlan'

---

**Note:** Installed into `/usr/share/one/oneprovision/templates/static_vxlan.yaml`.

---

Template with private OpenNebula virtual networks configured by *static_vxlan* on physical hosts.

Following virtual network(s) are configured:

- nat
- private

### Provision Drivers

### Packet Driver

Requirements:

- existing Packet project
- generated Packet API token

Supported Packet hosts:

- operating systems: `centos_7`, `ubuntu_16_04`, `ubuntu_18_04`
- plans: `baremetal_0`

### Provision parameters

The parameters specifying the Packet provision are set in the **provision section** of the *provision template*. The following table shows mandatory and optional provision parameters for the Packet driver. No other options/features are supported.

| Parameter | Mandatory | Description |
|---|---|---|
| `packet_project` | **YES** | ID of the existing project you want to deploy the new host |
| `packet_token` | **YES** | API token |
| `facility` | **YES** | Datacenter to deploy the host |
| `plan` | **YES** | Type of HW plan (Packet ID or slug) |
| `os` | **YES** | Operating system (Packet ID or slug) |
| `hostname` | NO | Hostname set in the booted operating system (and on Packet dashboard) |
| `billing_cycle` | NO | Billing period |
| `userdata` | NO | Custom EC2-like user data passed to the `cloud-init` running on host |
| `tags` | NO | Host tags in the Packet inventory |

### Example

Example of minimal Packet provision template:

```
---
name: myprovision

defaults:
  provision:
    driver: packet
    packet_token: *********************************
    packet_project: *************************************
    facility: ams1
    plan: baremetal_0
    os: ubuntu_18_04

hosts:
  - reserved_cpu: 100
    im_mad: kvm
    vm_mad: kvm
    provision:
      hostname: "host1"
```

### Amazon EC2 Driver

Requirements:

- AWS API credentials (access and secret key)

Supported EC2 hosts:

- operating systems: CentOS 7, Ubuntu 16.04 and 18.04

- AMIs (us-east1): `ami-66a7871c` (CentOS 7), `ami-759bc50a` (Ubuntu 16.04)

- instance types: `i3.metal`

### Provision parameters

The following table shows base provision parameters for the EC2 driver:

| Parameter | Mandatory | Description |
|---|---|---|
| `ec2_access` | **YES** | AWS access key |
| `ec2_secret` | **YES** | AWS secret key |
| `region_name` | **YES** | AWS deployment's region |
| `instancetype` | **YES** | Type of HW plan |
| `ami` | **YES** | AWS image ID (operating system) |
| `securitygroupids` | **YES** | AWS security group IDs |
| `subnetid` | **YES** | AWS subnet ID |
| `hostname` | NO | Hostname set in the booted operating system (and on EC2 dashboard) |
| `userdata` | NO | Custom user data |
| `keypair` | NO | EC2 keypair name |
| `tags` | NO | Host tags in the EC2 inventory |
| `cloud_init` | NO | Generate `cloud-init` contextualization data if no custom `userdata` specified (default: `no`). See *Cloud-init*. |

and, all the parameters supported by the *EC2 cloud-bursting* driver.

### Cloud-init

Cloud-init is a popular tool to contextualize the cloud resources. Although the OpenNebula mainly supports own contextualization method and tools (see packages for Linux and Windows), the EC2 driver can be forced to provide configuration for generic images relying on the cloud-init.

The cloud-init configuration is generated if

1. explicitly enabled by `cloud_init` parameter and

2. no custom `userdata` parameter is specified.

Parameter `cloud_init` allowed values:

- `yes` or `true`

- `no` or `false`

Supported functionality:

- provide authorized keys for default user based on the OpenNebula context keys

### Example

Example of minimal EC2 provision template:

```
---
name: myprovision

defaults:
  provision:
    driver: ec2
    ec2_access: ********************
```

```
    ec2_secret: *****************************************
    region_name: "us-east-1"
    cloud_init: true
    ami: ami-66a7871c
    instancetype: "i3.metal"
    securitygroupsids: sg-*****************
    subnetid: subnet-*****************

hosts:
  - reserved_cpu: 100
    im_mad: kvm
    vm_mad: kvm
    provision:
      hostname: "host1"
```

### 9.4.2 Configuration Reference

#### Configuration

Newly provisioned physical resources are mostly running only a base operating system without any additional services. Hosts need to pass the configuration phase to setup the additional software repositories, install packages, and configure and run necessary services to comply with the intended host purpose. This configuration process is fully handled by the `oneprovision` tool as part of the initial deployment (`oneprovision create`) or independent run anytime later (`oneprovision configure`).

---

**Note:** Tool `oneprovision` has a seamless integration with the Ansible (needs to be already installed on the system). It's not necessary to be familiar with the Ansible unless you need to make changes deep inside the configuration process.

---

As we use the Ansible for the host configuration, we'll also share their terminology in this section.

- **task** - single configuration step (e.g. package installation, service start)

- **role** - set of related tasks (e.g. role to deal with Linux bridges - utils install, bridge configure, and activation)

- **playbook** - set of roles/tasks to configure several components at once (e.g. fully prepare a host as KVM hypervisor)

The configuration phase can be parameterized to slightly change the configuration process (e.g. enable or disable particular OS feature, or force different repository location or version). These custom parameters are specified in the *configuration* section of the provision template. For most cases, the general defaults should match most requirements.

All code for the Ansible (tasks, roles, playbooks) is installed into `/usr/share/one/oneprovision/ansible/` with the following high-level structure:

```
/usr/share/one/oneprovision/ansible
|-- inventories
|   |-- default
|   `-- static_vxlan
|-- roles
|   |-- bridged-networking
|   |-- ddc
|   |-- iptables
|   |-- opennebula-node-kvm
```

```
|   |-- opennebula-node-lxd
|   |-- opennebula-p2p-vxlan
|   |-- opennebula-repository
|   |-- opennebula-ssh
|   |-- python
|   `-- tuntap
|-- default.yml
`-- static_packet.yml
```

Description:

- `*.yml` - playbooks

- `inventories/` - parameters for each playbook

- `roles/` - roles with tasks

The detailed description of all roles and their configuration parameters is in separate chapter *Roles*, which is intended for advanced users.

### Playbooks

### Playbooks

Playbooks are extensive descriptions of the configuration process (what and how is installed and configured on the physical host). Each configuration description prepares a physical host from the initial to the final ready-to-use state. Each description can configure the host in a completely different way (e.g. KVM host with private networking, KVM host with shared NFS filesystem, or KVM host supporting Packet elastic IPs, etc.). *Configuration parameters* are only a small tunables to the configuration process driven by the playbooks.

Before the deployment, a user must choose from the available playbooks to apply on the host(s).

### Playbook 'default'

---

**Note:  Description:** KVM host with host-only networking and NAT

---

This configuration prepares the host with

- KVM hypervisor

- network 1: bridge `br0` for the private host-only networking and NAT

- masquerade (NAT) to allow VMs from **network 1** access the public services

---

**Important:**  If more physical hosts are created, the private traffic of the virtual machines isn't routed between them. Virtual machines on different hosts are isolated although sharing the same private address space! This is the simplest configuration type.

---

### Networking 1 (host-only with NAT)

On the physical host, the IP configuration of prepared bridge `br0` (with TAP interface `tap0`) is same on all hosts:

| Parameter | Value |
|-----------|-------|
| Interface | `br0` |
| Slave | `tap0` |
| IP address | `192.168.150.1` |
| Netmask | `255.255.255.0` |

For **virtual machines**, the following IP configuration can be used:

| Parameter | Value |
|-----------|-------|
| IP address | any from range `192.168.150.2 - 192.168.150.254` |
| Netmask | `255.255.255.0` |
| Gateway (NAT) | `192.168.150.1` |

## Create OpenNebula Virtual Network

### From Provision Template

Put the full network definition into your provision template:

```
networks:
  - name: "nat"
    vn_mad: dummy
    bridge: br0
    dns: "8.8.8.8 8.8.4.4"
    gateway: "192.168.150.1"
    description: "Host-only networking with NAT"
    ar:
      - ip: "192.168.150.2"
        size: 253
        type: IP4
```

or, just easily extend the shipped template with above definition by setting the `extends` attribute in the provision template:

```
extends: /usr/share/one/oneprovision/templates/default.yaml
```

### Manually

In the OpenNebula, the virtual network for the virtual machines can be defined by the following template:

```
NAME         = "nat"
VN_MAD       = "dummy"
BRIDGE       = "br0"
DNS          = "8.8.8.8 8.8.4.4"
GATEWAY      = "192.168.150.1"
DESCRIPTION  = "Host-only networking with NAT"

AR=[
    TYPE = "IP4",
    IP   = "192.168.150.2",
    SIZE = "253"
]
```

Put the template above into a file and execute the following command to create a virtual network:

```
$ onevnet create net1.tpl
ID: 1
```

### Parameters

Main configuration parameters:

| Parameter | Value | Description |
|---|---|---|
| `bridged_networking_static_ip` | 192.168.150.1 | IP address of the bridge |
| `bridged_networking_static_netmask` | 255.255.255.0 | Netmask of the bridge |
| `opennebula_node_kvm_use_ev` | True or False | Whether to use the ev package for kvm |
| `opennebula_node_kvm_param_nested` | True or False | Enable nested KVM virtualization |
| `opennebula_repository_version` | 5.8 | OpenNebula repository version |
| `opennebula_repository_base` | http://downloads.opennebula.org/repo/{{ opennebula_repository_version }} | Repository of the OpenNebula packages |

All parameters are covered in the *Configuration Roles*

### Configuration Steps

The roles and tasks are applied during the configuration in the following order:

1. **python** - check and install Python required for Ansible

2. **ddc** - general asserts and cleanups

3. **opennebula-repository** - setup OpenNebula package repository

4. **opennebula-node-kvm** - install OpenNebula node KVM package

5. **opennebula-ssh** - deploy local SSH keys for the remote oneadmin

6. **tuntap** - create TAP `tap0` interface

7. **bridged-networking** - bridge Linux bridge `br0` with TAP interface

8. **iptables** - create basic iptables rules and enable NAT

with the following configuration overrides to the *roles defaults*:

| Parameter | Value |
|---|---|
| `opennebula_node_kvm_use_ev` | true |
| `bridged_networking_iface` | tap0 |
| `bridged_networking_iface_manage` | false |
| `bridged_networking_static_ip` | 192.168.150.1 |
| `iptables_masquerade_enabled` | true |
| `iptables_base_rules_strict` | false |

### Playbook 'default_lxd'

---

**Note:** **Description:** LXD host with host-only networking and NAT

---

This configuration prepares the host with

- LXD hypervisor

- network 1: bridge `br0` for the private host-only networking and NAT

- masquerade (NAT) to allow VMs from **network 1** access the public services

---

**Important:** If more physical hosts are created, the private traffic of the virtual machines isn't routed between them. Virtual machines on different hosts are isolated although sharing the same private address space! This is the simplest configuration type.

---

### Networking 1 (host-only with NAT)

On the physical host, the IP configuration of prepared bridge `br0` (with TAP interface `tap0`) is same on all hosts:

| Parameter | Value |
|-----------|-------|
| Interface | `br0` |
| Slave | `tap0` |
| IP address | `192.168.150.1` |
| Netmask | `255.255.255.0` |

For **virtual machines**, the following IP configuration can be used:

| Parameter | Value |
|-----------|-------|
| IP address | any from range `192.168.150.2 - 192.168.150.254` |
| Netmask | `255.255.255.0` |
| Gateway (NAT) | `192.168.150.1` |

### Create OpenNebula Virtual Network

### From Provision Template

Put the full network definition into your provision template:

```
networks:
  - name: "nat"
    vn_mad: dummy
    bridge: br0
    dns: "8.8.8.8 8.8.4.4"
    gateway: "192.168.150.1"
    description: "Host-only networking with NAT"
    ar:
      - ip: "192.168.150.2"
```

(continues on next page)

```
        size: 253
        type: IP4
```

or, just easily extend the shipped template with above definition by setting the `extends` attribute in the provision template:

```
extends: /usr/share/one/oneprovision/templates/default.yaml
```

### Manually

In the OpenNebula, the virtual network for the virtual machines can be defined by the following template:

```
NAME        = "nat"
VN_MAD      = "dummy"
BRIDGE      = "br0"
DNS         = "8.8.8.8 8.8.4.4"
GATEWAY     = "192.168.150.1"
DESCRIPTION = "Host-only networking with NAT"

AR=[
    TYPE = "IP4",
    IP   = "192.168.150.2",
    SIZE = "253"
]
```

Put the template above into a file and execute the following command to create a virtual network:

```
$ onevnet create net1.tpl
ID: 1
```

### Parameters

Main configuration parameters:

| Parameter | Value | Description |
|---|---|---|
| bridged_networking_static_ip | 192.168.150.1 | IP address of the bridge |
| bridged_networking_static_netmask | 255.255.255.0 | Netmask of the bridge |
| opennebula_repository_version | 5.8 | OpenNebula repository version |
| opennebula_repository_base | http://downloads.opennebula.org/repo/ {{ opennebula_repository_version }} | Repository of the OpenNebula packages |

All parameters are covered in the *Configuration Roles*

### Configuration Steps

The roles and tasks are applied during the configuration in the following order:

1. **python** - check and install Python required for Ansible

2. **ddc** - general asserts and cleanups

3. **opennebula-repository** - setup OpenNebula package repository

4. **opennebula-node-lxd** - install OpenNebula node LXD package

5. **opennebula-ssh** - deploy local SSH keys for the remote oneadmin

6. **tuntap** - create TAP `tap0` interface

7. **bridged-networking** - bridge Linux bridge `br0` with TAP interface

8. **iptables** - create basic iptables rules and enable NAT

with the following configuration overrides to the *roles defaults*:

| Parameter | Value |
| --- | --- |
| `bridged_networking_iface` | tap0 |
| `bridged_networking_iface_manage` | false |
| `bridged_networking_static_ip` | 192.168.150.1 |
| `iptables_masquerade_enabled` | true |
| `iptables_base_rules_strict` | false |

### Playbook 'static_vxlan'

**Note:** **Description:** KVM host with static private networking and NAT.

This configuration prepares the host with

- KVM hypervisor

- network 1: bridge `br0` for the private host-only networking and NAT

- network 2: bridge `vxbr100` with static VXLAN connections among all provisioned hosts

- masquerade (NAT) to allow VMs from **network 1** access the public services

### Networking 1 (host-only with NAT)

On the physical host, the IP configuration of prepared bridge `br0` (with TAP interface `tap0`) is same on all hosts:

| Parameter | Value |
| --- | --- |
| Interface | `br0` |
| Slave | `tap0` |
| IP address | `192.168.150.1` |
| Netmask | `255.255.255.0` |

For **virtual machines**, the following IP configuration can be used:

| Parameter | Value |
| --- | --- |
| IP address | any from range `192.168.150.2 - 192.168.150.254` |
| Netmask | `255.255.255.0` |
| Gateway (NAT) | `192.168.150.1` |

**Create OpenNebula Virtual Network**

**From Provision Template**

Put the first network definition into your provision template:

```
networks:
  - name: "nat"
    vn_mad: dummy
    bridge: br0
    dns: "8.8.8.8 8.8.4.4"
    gateway: "192.168.150.1"
    description: "Host-only networking with NAT"
    ar:
      - ip: "192.168.150.2"
        size: 253
        type: IP4
```

or, just easily extend the shipped template with both network definitions by setting the `extends` attribute in the provision template:

```
extends: /usr/share/one/oneprovision/templates/static_vxlan.yaml
```

**Manually**

In the OpenNebula, the virtual network for the virtual machines can be defined by the following template:

```
NAME        = "nat"
VN_MAD      = "dummy"
BRIDGE      = "br0"
DNS         = "8.8.8.8 8.8.4.4"
GATEWAY     = "192.168.150.1"
DESCRIPTION = "Host-only networking with NAT"

AR=[
    TYPE = "IP4",
    IP   = "192.168.150.2",
    SIZE = "253"
]
```

Put the template above into a file and execute the following command to create a virtual network:

```
$ onevnet create net1.tpl
ID: 1
```

**Networking 2 (private among hosts)**

On the physical host, another bridge `vxbr100` will be created without any IP configuration.

| Parameter | Value |
|---|---|
| Interface | `vxbr100` |
| Slave | `vxlan100` |
| Physical | `bond0:0` or `eth0` |
| IP address | none |
| Netmask | none |

For **virtual machines**, any IPs distinct to existing IP ranges configured on the host can be used. For example:

| Parameter | Value |
|---|---|
| IP address | any from range `192.168.160.2 - 192.168.160.254` |
| Netmask | `255.255.255.0` |
| Gateway (NAT) | none |

### Create OpenNebula Virtual Network

### From Provision Template

Put the second network definition into your provision template:

```
networks:
  - name:        "private"
    vn_mad:      "dummy"
    bridge:      "vxbr100"
    mtu:         "1450"
    description: "Private networking"
    ar:
      - ip:    "192.168.160.2"
        size: "253"
        type: "IP4"
```

or, just easily extend the shipped template with both network definitions by setting the `extends` attribute in the provision template:

```
extends: /usr/share/one/oneprovision/templates/static_vxlan.yaml
```

### Manually

In the OpenNebula, the virtual network for the virtual machines can be defined by the following template:

```
NAME        = "private"
VN_MAD      = "dummy"
BRIDGE      = "vxbr100"
MTU         = 1450
DESCRIPTION = "Private networking"

AR=[
    TYPE = "IP4",
    IP   = "192.168.160.2",
    SIZE = "253"
]
```

Put the template above into a file and execute the following command to create a virtual network:

```
$ onevnet create net2.tpl
ID: 2
```

### Parameters

Main configuration parameters:

| Parameter | Value | Description |
|---|---|---|
| `bridged_networking_static_ip` | 192.168.150.1 | IP address of the bridge |
| `bridged_networking_static_netmask` | 255.255.255.0 | Netmask of the bridge |
| `opennebula_node_kvm_use_ev` | **True** or False | Whether to use the ev package for kvm |
| `opennebula_node_kvm_param_nested` | True or **False** | Enable nested KVM virtualization |
| `opennebula_repository_version` | 5.8 | OpenNebula repository version |
| `opennebula_repository_base` | https://downloads.opennebula.org/repo/ {{ opennebula_repository_version }} | Repository of the OpenNebula packages |

All parameters are covered in the *Configuration Roles*

### Configuration Steps

The roles and tasks are applied during the configuration in the following order:

1. **python** - check and install Python required for Ansible
2. **ddc** - general asserts and cleanups
3. **opennebula-repository** - setup OpenNebula package repository
4. **opennebula-node-kvm** - install OpenNebula node KVM package
5. **opennebula-ssh** - deploy local SSH keys for the remote oneadmin
6. **tuntap** - create TAP `tap0` interface
7. **bridged-networking** - bridge Linux bridge `br0` with TAP interface
8. **opennebula-p2p-vxlan** - bridge `vxlan100` with static VXLAN connections among hosts
9. **iptables** - create basic iptables rules and enable NAT

with the following configuration overrides to the *roles defaults*:

| Parameter | Value |
|---|---|
| `opennebula_node_kvm_use_ev` | true |
| `bridged_networking_iface` | tap0 |
| `bridged_networking_iface_manage` | false |
| `bridged_networking_static_ip` | 192.168.150.1 |
| `iptables_masquerade_enabled` | true |
| `iptables_base_rules_strict` | false |
| `opennebula_p2p_vxlan_bridge` | vxbr100 |
| `opennebula_p2p_vxlan_phydev` | bond0:0 or eth0 |
| `opennebula_p2p_vxlan_vxlan_vni` | 100 |
| `opennebula_p2p_vxlan_vxlan_dev` | vxlan100 |
| `opennebula_p2p_vxlan_vxlan_local_ip` | autodetect IPv4 address on bond0:0 or eth0 |
| `opennebula_p2p_vxlan_remotes` | autodetect list of IPv4 on bond0:0 or eth0 from all hosts |

## Roles

> **Warning:** This chapter is only for the advanced users who need to modify the host configuration process significantly. Unless the configuration process doesn't meet your requirements, you don't need to be familiar with this part.

Following roles are shipped with the OpenNebula provision tool and installed into `/usr/share/one/oneprovision/ansible/roles/`.

## Role bridged-networking

Creates a new bridge (identified by `bridged_networking_bridge`) and connects the specified network interface (`bridged_networking_iface`) into.

| Parameter | Default | Description |
|---|---|---|
| `bridged_networking_bridge` | br0 | The bridge that will be created |
| `bridged_networking_bridge_manage` | True | Manage configuration for bridge |
| `bridged_networking_iface` | eth1 | The network device connected to the bridge |
| `bridged_networking_iface_manage` | True | Manage configuration of interface connected to bridge |
| `bridged_networking_static_ip` | NULL | IP address of the bridge |
| `bridged_networking_static_netmask` | 255.255.255.0 | Netmask of the bridge |
| `bridged_networking_static_gateway` | NULL | Gateway of the bridge |
| `bridged_networking_ip_iface` | eth1 | Name of interface to take the IP configuration for bridge, if `bridged_networking_static_ip` not defined |

## Role ddc

Set of internal clean up and check tasks. E.g. check if the target host operating system is supported, or network configuration cleanups.

No parameters.

**Role iptables**

Creates the basic set of IPv4 and IPv6 packet filter rules to ensure only the specified traffic is allowed. Masqurade (NAT) with IP port forwarding can be enabled.

| Parameter | Default | Description |
|---|---|---|
| `iptables_ip_forward_enable` | true | Enable IP forwarding |
| `iptables_manage_persistent` | true | Manage persistent configuration |
| `iptables_base_rules_enable` | true | Create set of base rules |
| `iptables_base_rules_interface` | NULL | Particular network interface to limit the base rules |
| `iptables_base_rules_strict` | true | Include the rules to drop any other traffic |
| `iptables_base_rules_services` | [{protocol:'tcp', port:22}] | List of whitelisted services |
| `iptables_masquerade_enable` | false | Enable NAT |
| `iptables_masquerade_interface` | ansible_default_ipv4.interface | NAT output interface |

**Role opennebula-node-kvm**

Installs the OpenNebula node-kvm package, optionally configures the KVM module for the nested virtualization, and ensures the libvirt is enabled and running.

| Parameter | Default | Description |
|---|---|---|
| `opennebula_node_kvm_use_ev` | False | Whether to use the ev package for kvm |
| `opennebula_node_kvm_param_nested` | False | Enable nested KVM virtualization |
| `opennebula_node_kvm_manage_kvm` | True | Enable KVM configuration |
| `opennebula_node_kvm_rhev_repo` | True | Name of Red Hat EV repository |
| `opennebula_node_selinux_booleans` | virt_use_nfs | SELinux booleans to configure |

**Role opennebula-node-lxd**

Installs the opennebula-node-lxd package.

No parameters.

**Role opennebula-p2p-vxlan**

Creates static VXLAN connections between several physical hosts. This allows to have the limited private networking in the infrastructures where VXLAN discovery over multicast isn't supported.

| Parameter | Default | Description |
|---|---|---|
| `opennebula_p2p_vxlan_bridge` | NULL | Name of VXLAN bridge |
| `opennebula_p2p_vxlan_phydev` | NULL | Name of VXLAN physical interface |
| `opennebula_p2p_vxlan_vxlan_vni` | NULL | VXLAN ID (VNI) |
| `opennebula_p2p_vxlan_vxlan_dev` | NULL | Name of VXLAN device |
| `opennebula_p2p_vxlan_vxlan_local_ip` | NULL | Source IP address to use by VXLAN device |
| `opennebula_p2p_vxlan_remotes` | [] | List of all remote VXLAN endpoints |

### Role opennebula-repository

Configures the OpenNebula package repository for the particular version.

| Parameter | Default | Description |
|---|---|---|
| `opennebula_repository_version` | `5.8` | OpenNebula repository version |
| `opennebula_repository_base` | `https://downloads.`<br>`opennebula.org/repo/` `{{`<br>`opennebula_repository_version }}` | Repository of the OpenNebula packages |
| `opennebula_repository_gpgcheck` | `yes` | Enable GPG check for the packages |
| `opennebula_repository_repo_gpgcheck` | `yes` | Enable GPG check for the repos (RHEL/CentOS only) |

### Role opennebula-ssh

Handles the SSH configuration and SSH keys distribution on the OpenNebula frontend/hosts.

| Parameter | Default | Description |
|---|---|---|
| `opennebula_ssh_manage_sshd` | True | Manage SSH server configuration |
| `opennebula_ssh_sshd_passwordauthentication` | no | SSH server option for Password Authentication |
| `opennebula_ssh_sshd_permitrootlogin` | without-password | SSH server option for PermitRootLogin |
| `opennebula_ssh_deploy_local` | True | Deploy local oneadmin's SSH key to remote host |

### Role python

Installs python2 for Debian and Ubuntu.

No parameters.

### Role tuntap

The role creates a TUN/TAP interface with persistent configuration.

| Parameter | Default | Description |
|---|---|---|
| `tuntap_name` | tap0 | Name of interface |
| `tuntap_mode` | tap | Interface mode |