



OpenNebula 5.4 Introduction and Release Notes

Release 5.4.0

OpenNebula Systems

Jul 20, 2017

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

1	Concepts and Terminology	1
1.1	Start Here: OpenNebula Overview	1
1.2	OpenNebula Key Features	5
1.3	Glossary	9
2	Release Notes 5.4	11
2.1	What's New in 5.4	11
2.2	Platform Notes	15
2.3	Compatibility Guide	19
2.4	Known Issues	23
2.5	Acknowledgements	25
3	Upgrading	26
3.1	Overview	26
3.2	Upgrading from OpenNebula 5.2.x	26
3.3	Upgrading from OpenNebula 5.0.x	31
3.4	Upgrading from OpenNebula 4.14.x	36
3.5	Upgrading from OpenNebula 4.12.x	42
3.6	Upgrading from OpenNebula 4.10.x	48
3.7	Upgrading from OpenNebula 4.8.x	54
3.8	Upgrading from OpenNebula 4.6.x	59
3.9	Upgrading from OpenNebula 4.4.x	65
3.10	Upgrading from OpenNebula 4.2	69
3.11	Upgrading from OpenNebula 4.0.x	74
3.12	Upgrading from OpenNebula 3.8.x	79
3.13	vCenter upgrade 5.2 to 5.4	87

CONCEPTS AND TERMINOLOGY

1.1 Start Here: OpenNebula Overview

Welcome to OpenNebula documentation!

OpenNebula is an open-source management platform to build IaaS private, public and hybrid clouds. Installing a cloud from scratch could be a complex process, in the sense that many components and concepts are involved. The degree of familiarity with these concepts (system administration, infrastructure planning, virtualization management...) will determine the difficulty of the installation process.

If you are new to OpenNebula you should go through this short introduction before proceeding to the deployment and administration guides.

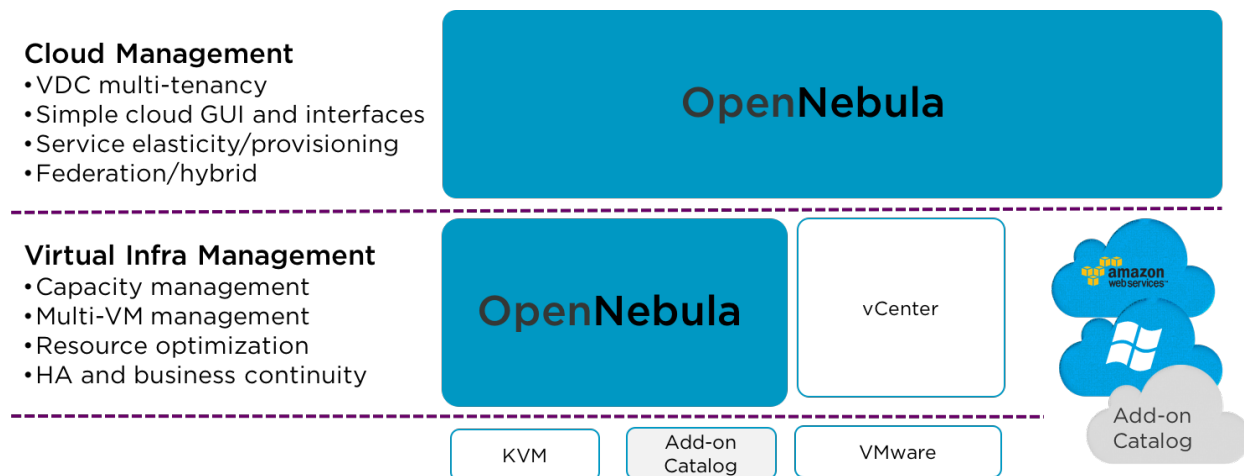
1.1.1 Step 1. Choose Your Hypervisor

The first step is to decide on the hypervisor that you will use in your cloud infrastructure. The main OpenNebula distribution provides full support for the two most widely used hypervisors, KVM and VMware (through vCenter), at different levels of functionality.

- **Virtualization and Cloud Management on KVM.** Many companies use OpenNebula to manage data center virtualization, consolidate servers, and integrate existing IT assets for computing, storage, and networking. In this deployment model, OpenNebula directly integrates with KVM and has complete control over virtual and physical resources, providing advanced features for capacity management, resource optimization, high availability and business continuity. Some of these deployments additionally use OpenNebula's **Cloud Management and Provisioning** features when they want to federate data centers, implement cloudbursting, or offer self-service portals for end users.
- **Cloud Management on VMware vCenter.** Other companies use OpenNebula to provide a multi-tenant, cloud-like provisioning layer on top of VMware vCenter. These deployments are looking for provisioning, elasticity and multi-tenancy cloud features like virtual data centers provisioning, datacenter federation or hybrid cloud computing to connect in-house infrastructures with public clouds, while the infrastructure is managed by already familiar tools for infrastructure management and operation, such as vSphere and vCenter Operations Manager.

After having installed the cloud with one hypervisor you may add another hypervisors. You can deploy heterogeneous multi-hypervisor environments managed by a single OpenNebula instance. An advantage of using OpenNebula on VMware is the strategic path to openness as companies move beyond virtualization toward a private cloud. OpenNebula can leverage existing VMware infrastructure, protecting IT investments, and at the same time gradually integrate other open-source hypervisors, therefore avoiding future vendor lock-in and strengthening the negotiating position of the company.

There are other virtualization technologies, like LXC or Xen, supported by the community. Please refer to the [OpenNebula Add-ons Catalog](#).



1.1.2 Step 2. Design and Install the Cloud

2.1. Design the Cloud Architecture

In order to get the most out of a OpenNebula Cloud, we recommend that you create a plan with the features, performance, scalability, and high availability characteristics you want in your deployment. We have prepared **Cloud Architecture Design guides** for KVM and vCenter to help you plan an OpenNebula installation, so you can easily architect your deployment and understand the technologies involved in the management of virtualized resources and their relationship. These guides have been created from the collective information and experiences from hundreds of users and cloud client engagements. Besides main logical components and interrelationships, this guides document software products, configurations, and requirements of infrastructure platforms recommended for a smooth OpenNebula installation.

2.2. Install the Front-end

Next step is the **installation of OpenNebula in the cloud front-end**. This installation process is the same for any underlying hypervisor.

Optionally you can setup a high available cluster for OpenNebula for OpenNebula to reduce downtime of core OpenNebula services, and configure a MySQL backend as an alternative to the default Sqlite backend if you are planning a large-scale infrastructure.

2.3. Install the Virtualization hosts

Now you are ready to **add the virtualization nodes**. The OpenNebula packages bring support for KVM and vCenter nodes. In the case of vCenter, a host represents a vCenter cluster with all its ESX hosts. You can add different hypervisors to the same OpenNebula instance, or any other virtualization technology, like LXC or Xen, supported by the community. Please refer to the [OpenNebula Add-ons Catalog](#).

1.1.3 Step 3. Set-up Infrastructure and Services

3.1. Integrate with Data Center Infrastructure

Now you should have an OpenNebula cloud up and running with at least one virtualization node. The next step is, if needed, to perform the integration of OpenNebula with your infrastructure platform and define the configuration of

its components. When using the vCenter driver, no additional integration is required because the interaction with the underlying networking, storage and compute infrastructure is performed through vCenter.

However when using KVM, in the open cloud architecture, OpenNebula directly manages the hypervisor, networking and storage platforms, and you may need additional configuration:

- **Networking setup** with 802.1Q VLANs, ebttables, Open vSwitch or VXLAN.
- **Storage setup** with filesystem datastore, LVM datastore, Ceph, Dev, or iSCSI datastore.
- **Host setup** with the configuration options for the KVM hosts, Monitoring subsystem, Virtual Machine HA or PCI Passthrough.

3.2. Configure Cloud Services

OpenNebula comes by default with an internal **user/password authentication system**. Optionally you can enable an external Authentication driver like ssh, x509, ldap or Active Directory.

Sunstone, the OpenNebula GUI, brings by default a pre-defined configuration of views. Optionally it can be customized and extended to meet your needs. You can customize the roles and views, improve security with x509 authentication and SSL or improve scalability for large deployments.

We also provide **references** with a detailed description of the different configuration files, and logging and debugging reports of the OpenNebula services.

1.1.4 Step 4. Operate your Cloud

4.1. Define a Provisioning Model

Before configuring multi-tenancy and defining the provisioning model of your cloud, we recommend you go through this introduction to the OpenNebula provisioning model. In a small installation with a few hosts, you can skip this guide and use OpenNebula without giving much thought to infrastructure partitioning and provisioning. But for medium and large deployments you will probably want to provide some level of isolation and structure.

- Regarding the **underlying infrastructure**, OpenNebula provides complete functionality for the management of the physical hosts and clusters in the cloud. A Cluster is a group of Hosts that can have associated Datastores and Virtual Networks.
- Regarding **user management**, OpenNebula features advanced multi-tenancy with powerful users and groups management, an Access Control List mechanism allowing different role management with fine grain permission granting over any resource, resource quota management to track and limit computing, storage and networking utilization, and a configurable accounting and showback systems to visualize and report resource usage data and to allow their integration with chargeback and billing platforms, or to guarantee fair share of resources among users.
- Last but not least, you can define VDCs (Virtual Data Center) as assignments of one or several user groups to a pool of physical resources. While clusters are used to group physical resources according to common characteristics such as networking topology or physical location, Virtual Data Centers (VDCs) allow to create “logical” pools of resources (which could belong to different clusters and cones) and allocate them to user groups.

4.2. Manage Virtual Resources

Now everything is ready for operation. OpenNebula provides full control to manage virtual resources.

- **Virtual machine image management** that allows to store disk images in catalogs (termed datastores), that can be then used to define VMs or shared with other users. The images can be OS installations, persistent data sets or empty data blocks that are created within the datastore.
- **Virtual network management** of Virtual networks that can be organized in network catalogs, and provide means to interconnect virtual machines. This kind of resources can be defined as IPv4, IPv6, or mixed networks, and can be used to achieve full isolation between virtual networks. Networks can be easily interconnected by using virtual routers and KVM users can also dynamically configure security groups
- **Virtual machine template management** with template catalog system that allows to register virtual machine definitions in the system, to be instantiated later as virtual machine instances.
- **Virtual machine instance management** with a number of operations that can be performed to control lifecycle of the virtual machine instances, such as migration (live and cold), stop, resume, cancel, power-off, etc.

Several reference guides are provided for more information about definition files, templates and CLI.

4.3. Create Virtual Machines

One of the most important aspects of the cloud is the **preparation of the images** for our users. OpenNebula uses a method called contextualization to send information to the VM at boot time. Its most basic usage is to share networking configuration and login credentials with the VM so it can be configured. More advanced cases can be starting a custom script on VM boot or preparing configuration to use OpenNebula Gate.

1.1.5 Step 5. Install Advanced Components

This step is optional and only for advanced users. We recommend you familiarize with OpenNebula before installing these components.

OpenNebula brings the following advanced components:


- Implementation of the EC2 Query and EBS **public cloud** interfaces.
- OneFlow allows **multi-VM application and auto-scaling** to define, execute and manage multi-tiered elastic applications, or services composed of interconnected Virtual Machines with deployment dependencies between them and auto-scaling rules.
- The datacenter federation functionality allows for the **centralized management of multiple instances of OpenNebula for scalability, isolation and multiple-site support**.
- **Application insight** with OneGate 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.
- Marketplaces for sharing, provisioning and consuming cloud images. They can be seen as external datastores, where images can be easily imported, exported and shared by a federation of OpenNebula instances.
- **Cloud bursting** gives support to build a hybrid cloud, an extension of a private cloud to combine local resources with resources from remote cloud providers. A whole public cloud provider can be encapsulated as a local resource to be able to use extra computational capacity to satisfy peak demands. Out of the box connectors are shipped to support Amazon EC2 and Microsoft Azure cloudbursting.

1.1.6 Step 6. Integrate with other Components

This step is optional and only for integrators and builders.

Because no two clouds are the same, OpenNebula provides many different interfaces that can be used to interact with the functionality offered to manage physical and virtual resources.

- **Modular and extensible architecture** with customizable plug-ins for integration with any third-party data center infrastructure platform for storage, monitoring, networking, authentication, virtualization, cloud bursting and market.
- **API for integration** with higher level tools such as billing, self-service portals... that offers all the rich functionality of the OpenNebula core, with bindings for ruby and java and XMLRPC API,
- **OneFlow API** to create, control and monitor multi-tier applications or services composed of interconnected Virtual Machines.
- **Sunstone custom routes and tabs** to extend the sunstone server.
- **Hook Manager** to trigger administration scripts upon VM state change.



1.2 OpenNebula Key Features

OpenNebula offers a **simple but feature-rich and flexible solution** to build and manage data center virtualization and enterprise clouds. This guide summarizes its key features(*). You can also refer to the *Platform Notes* included in the documentation of each version to know about the infrastructure platforms and services supported by OpenNebula.

INTERFACES FOR CLOUD CONSUMERS

- De-facto standard cloud APIs with compatibility with cloud ecosystem tools
- Simple, clean, intuitive GUI for cloud consumers to allow non-IT end users to easily create, deploy and manage compute, storage and network resources

VIRTUAL MACHINE MANAGEMENT

- Virtual infrastructure management adjusted to enterprise data centers with full control, monitoring and accounting of virtual resources

- Virtual machine image management through catalogs of disk images (termed datastores) with OS installations, persistent data sets or empty data blocks that are created within the datastore
- Virtual machine template management through catalogs of templates that allow to register virtual machine definitions in the system to be instantiated later as virtual machine instances
- Virtual machine instance management with full control of virtual machine lifecycle
- Programmable VM operations allowing users to schedule actions
- Volume hotplugging and disk snapshot capabilities and disk resizing for KVM virtual machines

VIRTUAL NETWORK MANAGEMENT

- Advanced network virtualization capabilities with traffic isolation, address reservation, flexible definition of address ranges to accommodate any address distribution, definition of generic attributes to define multi-tier services...
- IPv6 support with definition site and global unicast addresses
- Virtual routers
- Security Groups to define firewall rules and apply them to KVM Virtual Machines

APPLICATION CONFIGURATION AND INSIGHT

- Automatic installation and configuration of application environments
- VM attributes can be provided by the user when the template is instantiated
- Wide range of guest operating system including Microsoft Windows and Linux
- Gain insight cloud applications so their status and metrics can be easily queried through OpenNebula interfaces and used in auto-scaling rules

MULTI-VM APPLICATION MANAGEMENT

- Automatic execution of multi-tiered (multi-VM) applications and their provision from a catalog and self-service portal
- Automatic scaling of multi-tiered applications according to performance metrics and time schedule

INTERFACES FOR ADMINISTRATORS AND ADVANCED USERS

- Powerful Command Line Interface that resembles typical UNIX commands applications
- Easy-to-use Sunstone Graphical Interface providing usage graphics and statistics with cloudwatch-like functionality, remote access through VNC or SPICE, different system views for different roles, catalog access, multiple-zone management...
- Sunstone is easily customizable to define multiple cloud views for different user groups

APPLIANCE MARKETPLACE

- Access to the public [OpenNebula Systems Marketplace](#) with a catalog of OpenNebula-ready cloud images
- Create your private centralized catalog (external satastore) of cloud applications (images and templates)
- Move VM images and templates across different types of datastores within the same OpenNebula instance
- Share VM images in Federation environments across several OpenNebula instances

ACCOUNTING AND SHOWBACK

- Configurable accounting system to report resource usage data and guarantee fair share of resources among users
- Easy integration with chargeback and billing platforms
- Showback capabilities to define cost associated to CPU/hours and MEMORY/hours per VM Template

MULTI-TENANCY AND SECURITY

- Fine-grained ACLs for resource allocation
- Powerful user and role management
- Administrators can group users into organizations that can represent different projects, division. . .
- Integration with external identity management services
- Special authentication mechanisms for SunStone (OpenNebula GUI) and the Cloud Services (EC2)
- Login token functionality to password based logins
- Fine-grained auditing
- Support for isolation at different levels
- Advanced access control policies for VMs to redefine the access level (ADMIN, MANAGE and USE) required for each VM action
- Traceability on VM actions, VM history records logs the data associated to the action performed on a VM

ON-DEMAND PROVISION OF VIRTUAL DATA CENTERS

- A VDC (Virtual Data Center) is a fully-isolated virtual infrastructure environment where a Group of users, optionally under the control of the group admin, can create and manage compute and storage capacity
- There is a pre-configured Sunstone view for group admins

CAPACITY AND PERFORMANCE MANAGEMENT

- Host management with complete functionality for the management of the virtualization nodes in the cloud
- Dynamic creation of Clusters as pools of hosts that share datastores and virtual networks for load balancing, high availability, and high performance computing
- Customizable and highly scalable monitoring system and also can be integrated with external data center monitoring tools.
- Powerful and flexible scheduler for the definition of workload and resource-aware allocation policies such as packing, striping, load-aware, affinity-aware. . .
- Definition of groups of related VMs and set VM affinity rules across them.
- Resource quota management to track and limit computing, storage and networking resource utilization
- Support for multiple data stores to balance I/O operations between storage servers, or to define different SLA policies (e.g. backup) and performance features for different KVM VM types or users
- PCI passthrough available for KVM VMs that need consumption of raw GPU devices

FEDERATED CLOUD ENVIRONMENTS

- Federation of multiple OpenNebula Zones for scalability, isolation or multiple-site support
- Users can seamlessly provision virtual machines from multiple zones with an integrated interface both in Sunstone and CLI

HIGH AVAILABILITY AND BUSINESS CONTINUITY

- High availability architecture in active-passive configuration
- Persistent database backend with support for high availability configurations
- Configurable behavior in the event of host or KVM VM failure to provide easy to use and cost-effective failover solutions

CLOUD BURSTING

- Build a hybrid cloud to combine your local resources with resources from remote cloud provider and use extra computational capacity to satisfy peak demands

PLATFORM

- Fully platform independent
- Hypervisor agnostic with broad hypervisor support (KVM and VMware vCenter) and centralized management of environments with multiple hypervisors
- *Broad support for commodity and enterprise-grade hypervisor; monitoring, storage, networking and user management services*
- Packages for major Linux distributions

CUSTOMIZATION AND INTEGRATION

- Modular and extensible architecture to fit into any existing datacenter
- Customizable drivers for the main subsystems to easily leverage existing IT infrastructure and system management products: storage, monitoring, networking, authentication, virtualization, cloud bursting and market
- API for integration with higher level tools such as billing, self-service portals...
- Hook manager to trigger administration scripts upon VM state change
- Sunstone custom routes and tabs to extend the sunstone server
- OneFlow API to create, control and monitor multi-tier applications or services composed of interconnected Virtual Machines.
- [OpenNebula Add-on Catalog](#) with components enhancing the functionality provided by OpenNebula
- Configuration and tuning parameters to adjust behavior of the cloud management instance to the requirements of the environment and use cases

LICENSING

- Fully open-source software released under Apache license

INSTALLATION AND UPGRADE PROCESS

- *Configurable to deploy public, private and hybrid clouds*
- All key functionalities for enterprise cloud computing, storage and networking in a single install
- Long term stability and performance through a *single integrated patching and upgrade process*
- Automatic import of existing VMs running in local hypervisors and public clouds for hybrid cloud computing
- Optional building from source code
- System features a small footprint, less than 10Mb

QUALITY ASSURANCE

- Internal quality assurance process for functionality, scalability, performance, robustness and stability
- Technology matured through an active and engaged large community
- Scalability, reliability and performance tested on many massive scalable production deployments consisting of hundreds of thousands of cores and VMs

PRODUCT SUPPORT

- Best-effort community support
- SLA-based commercial support directly from the developers

- Integrated tab in Sunstone to access OpenNebula Systems professional support

(*) *Because OpenNebula leverages the functionality exposed by the underlying platform services, its functionality and performance may be affected by the limitations imposed by those services.*

- *The list of features may change on the different platform configurations*
- *Not all platform configurations exhibit a similar performance and stability*
- *The features may change to offer users more features and integration with other virtualization and cloud components*
- *The features may change due to changes in the functionality provided by underlying virtualization services*

1.3 Glossary

1.3.1 OpenNebula Components

- **Front-end:** Machine running the OpenNebula services.
- **Host:** Physical machine running a supported hypervisor. See the Host subsystem.
- **Cluster:** Pool of hosts that share datastores and virtual networks. Clusters are used for load balancing, high availability, and high performance computing.
- **Datastore:** Storage medium used as disk images repository or to hold images for running VMs.
- **Sunstone:** OpenNebula web interface. Learn more about Sunstone
- **Self-Service** OpenNebula web interfaced towards the end user. It is implemented by configuring a user view of the Sunstone Portal.
- **EC2 Service:** Server that enables the management of OpenNebula with EC2 interface. Learn more about EC2 Service.
- **OCA:** OpenNebula Cloud API. It is a set of libraries that ease the communication with the XML-RPC management interface. Learn more about ruby and java APIs.

1.3.2 OpenNebula Resources

- **Template:** Virtual Machine definition. These definitions are managed with the onetemplate command.
- **Image:** Virtual Machine disk image, created and managed with the oneimage command.
- **Virtual Machine:** Instantiated Template. A Virtual Machine represents one life-cycle, and several Virtual Machines can be created from a single Template. Check out the VM management guide.
- **Virtual Network:** A group of IP leases that VMs can use to automatically obtain IP addresses. See the Networking subsystem.
- **Virtual Data Center (VDC):** Defines an assignment of one or several Groups to a pool of Physical Resources. Typically this pool of Physical Resources consists of resources from one or several Clusters that could belong to different Zones or public external clouds for hybrid cloud computing.
- **Zone:** A group of interconnected physical hosts with hypervisors controlled by the same OpenNebula.

1.3.3 OpenNebula Management

- **ACL:** Access Control List. Check the managing ACL rules guide.
- **oneadmin:** Special administrative account. See the Users and Groups guide.
- **User:** An OpenNebula user account.
- **Group:** A group of Users.
- **Federation:** Several OpenNebula instances can be configured as zones.

RELEASE NOTES 5.4

2.1 What's New in 5.4

OpenNebula 5.4 (Medusa) is the third release of the OpenNebula 5 series. A significant effort has been applied in this release to stabilize features introduced in 5.2 Excession, while keeping an eye in implementing those features more demanded by the community.

As usual almost every component of OpenNebula has been reviewed to target usability and functional improvements, trying to keep API changes to a minimum to avoid disrupting ecosystem components. An important focus has been on the vCenter integration, with an enhanced network and storage management. Also, new components have been added to improve the OpenNebula experience.

A major overhaul has been applied to the vCenter integration. The team decided to go all the way and level the vCenter integration with the KVM support. This means:

- Full storage management. Non-Persistent images are now supported as well as volatile disks. OpenNebula is now aware of all VM disks and storage quotas can be enforced. Support for linked clones and Marketplace.
- Full network management. It is now possible to create vCenter standard and distributed port groups and even vSwitches directly from within OpenNebula. You can assign a VLAN ID to a port group created by OpenNebula.
- Improved monitoring. Up to two orders of magnitude of speedup.
- An enhanced import process where naming limitations in imported resources has been removed and the ability to enable VNC automatically for Wild VMs.
- Disk resizing, VM and Templates folder selection when a VM is deployed... and many more changes!

The screenshot shows the 'Create Virtual Network' wizard in the OpenNebula interface. The 'Conf' tab is active, showing configuration options for a virtual network. The 'Bridge' is set to 'sample_standard_port_group'. The 'Network mode' is set to 'vCenter'. The 'VLAN ID' is set to 'Manual VLAN ID' with a value of '113'. The 'Physical device' is 'vmmic2vmmic3' and the 'MTU of the interface' is '1350'. The 'Switch name' is 'sample_standard_switch', the 'Number of ports' is '100', and the 'Port group type' is 'Port group'. The 'OpenNebula's Host ID' is '66:devel_lvcentercenter5-devel-Datacenter_lfa8ddcad1f2'.

A new resource, the VM Groups, has been added to implement VM-VM and VM-Host affinity/anti-affinity rules. A VM group is a set of related virtual machines that may impose placement constraints based on affinity and anti-affinity

rules. A VM group is defined as a set of Roles. A Role defines a VM type or class, and expressions to the VM Group can be added to define affinity between VM roles, or between VM and hosts. This ensures a dynamic approach to affinity/anti-affinity since new VMs can be enrolled to a particular Role at boot time, after the VM Group has been defined and other VMs added to it.

```

ROLE = [
  NAME = backup
  POLICY = ANTI_AFFINED,
  HOST_ANTI_AFFINED="5,6"
]

ROLE = [
  NAME = monitor
]


ROLE = [
  NAME = db
  POLICY = AFFINED,
  HOST_AFFINED="2,3,4"
]

ROLE = [
  NAME = frontend
]

ANTI_AFFINED = "db, backup"
AFFINED = "monitor, front-end"

```

PLACEMENT EXAMPLE:



To top it all, OpenNebula 5.4 brings to the table a native implementation of a consensus algorithm, which enables the High Availability deployment of the OpenNebula front-end without relying to third party components. This distributed consensus protocol provides fault-tolerance and state consistency across OpenNebula services. A consensus algorithm is built around two concepts, System State -the data stored in the database tables- and Log -a sequence of SQL statements that are consistently applied to the OpenNebula DB in all servers-. 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 (follower*) to keep its leadership. If a leader fails to send the heartbeat, followers promote to candidates and start a new election. This feature, with support from floating IPs and a proper Sunstone configuration, gives robustness to OpenNebula clouds. This new functionality of distributed system state is also used to implement OpenNebula federation. In both cases (Federation and HA) no support is needed from MySQL to create a clustered DB, so admins can forget about MySQL replication.

There are many other improvements in 5.4, like improved VM lifecycle, flexible resource permissions, life disk resizing, improved Ceph support, enhanced disk I/O feedback, showback cost estimate in Sunstone, flexible IPv6 definition, http proxy support for marketplace, purge tools for the OpenNebula database, resource group isolation, multiple Sunstone improvements (VNC, password dialogs, confirmation dialogs, better vCenter support, persistent labels, usability enhancements), networking improvements, user inputs in OneFlow and many many more features to enrich your cloud experience. As with previous releases, and in order to achieve a reliable cloud management platform, the team has gone great lengths to fix reported bugs and improve general usability.

This OpenNebula release is named after the [Medula Nebula](#), a large planetary nebula in the constellation of Gemini on the Canis Minor border. It also known as Abell 21 and Sharpless 2-274. It was originally discovered in 1955 by UCLA astronomer George O. Abell, who classified it as an old planetary nebula. The braided serpentine filaments of glowing gas suggests the serpent hair of Medusa found in ancient Greek mythology.

OpenNebula 5.4 Medusa is considered to be a stable release and as such, and update is available in production environments.

In the following list you can check the highlights of OpenNebula 5.4 ([a detailed list of changes can be found here](#)):

2.1.1 OpenNebula Core

- **New HA model**, providing native HA (based on RAFT consensus algorithm) in OpenNebula components, including Sunstone without third party dependencies.
- **Simplified Federation setups**, it is no longer required to setup and maintain a MySQL replicated DB using the new distributed system state feature implemented in OpenNebula.
- **Advanced access control policies** for VMs, now it is possible to redefine the access level (ADMIN, MANAGE and USE) required for each VM action.
- **Improved traceability on VM actions**, VM history records logs the data associated to the action performed on a VM.
- **VM Groups** to define groups of related VMs and set affinity rules across them.
- **Database maintenance tools** to purge history records, update corrupted data and more through the onedb command.

2.1.2 Storage

- **Improved VM disk management**, including options to resize disks for running VMs, enhanced I/O feedback and monitoring.
- **Improved Ceph integration**, by default Ceph disk snapshots are in a flat hierarchy (this can be also be selected for other storage backends). Also it is now easier to setup multiple clusters with different authentication attributes and finally a new option has been added for trim/discard option.

2.1.3 Networking

- **Better IPv6 support**, including support for Security Groups and the definition of Non-SLAAC IPv6 Address Range.
- **Improved network settings** that may override multiple default options used when creating links in the hypervisors including MAC spoofing, arp cache poisoning, interface MTU or STP among others.

2.1.4 Hybrid Clouds: Amazon EC2

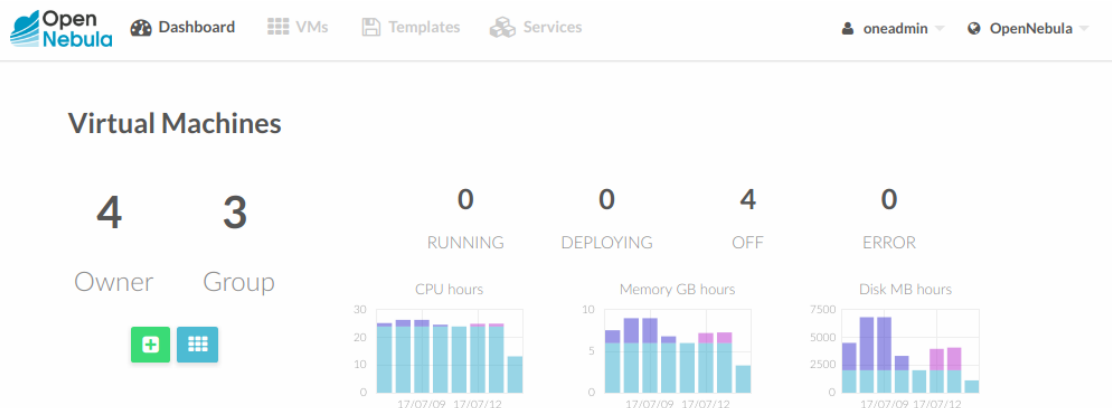
- **Enhanced EC2 monitoring and VM lifecycle**, with better handling of CloudWatch datapoints to avoid errors after long-term network problems.
- **Better EC2 resource characterization** the information to access and EC2 zone is now stored in the corresponding OpenNebula Host including EC2 credentials, capacity limits and EC2 zone name.

2.1.5 Scheduler

- **Affinity/Anti-affinity** for VM-VM, VM-Role and VM-Host using the new VM Group resource.
- **VM prioritization**, a static VM priority can be assigned to pending/reschedule VMs to alter the default FIFO ordering when dispatching VMs.

2.1.6 Sunstone

- **Improved customization** with more flags to restrict action usage and enhanced logo customization.
- **Persistent resource labels** that do not expire if no resource is tagged with a label.
- **Enhanced image upload control** with progress feedback and resume capabilities.
- **Better groups isolation** allowing to change the primary and secondary groups directly from the groups panel. Also group switch only shows current group resources to work by project easily.
- **Extended user inputs**, with new types like booleans and the possibility to define the order.
- **Fixed multilanguage keyboard support** in VNC feature.
- **Improved showback support**, with better dialogs to define and estimate the VM Template showback section.
- **A significant number of usability enhancements:**
 - More secure password change dialog.
 - ESC support for VNC dialog.
 - Improved overcommitment dialogs.
 - More presence of the VM logo in the VM Template and instance dialogs and tabs.
 - Confirmation dialog for destructive actions like reverting disks or erasing VMs.
 - Cloud view improved. Diff between own VMs and group VMs.



2.1.7 vCenter

The significant milestone is that vCenter is no longer treated as a public cloud by OpenNebula, but rather as a fully fledged hypervisor. The monitoring and import process have been optimized with a two orders of magnitude improvement in time efficiency.

- **Improved VM and VM Template management**, attach CDROM to a VM without a drive, add VNC capabilities to imported wild VMs, save VM as an OpenNebula template, linked clones capabilities, images and networks representing disks and nics are created for imported vCenter template and folder placement features, among others.
- **Network creation support**, a new vCenter network mode is available in virtual network definition, standard and different port groups and vSwitches can be created from within OpenNebula. VLAN IDs, MTUs and number of ports can be specified when a port group is created.

- **Improved Storage (datastore, Image and disk management)**, non-persistent images and volatiles disks are now supported, clustered datastores are clearly differentiated at import time, disk resize capabilities, save disk functionality and disks statistics monitoring.
- **vCenter default values**, some default values for vCenter attributes e.g NIC model, can be specified in a new configuration file.
- **Removed naming limitations**, like for instance vCenter cluster and datastore names with spaces are now supported.

2.2 Platform Notes

This page will show you the specific considerations at the time of using an OpenNebula cloud, according to the different supported platforms.

This is the list of the individual platform components that have been through the complete [OpenNebula Quality Assurance and Certification Process](#).

2.2.1 Certified Components Version

Front-End Components

Component	Version	More information
RedHat Enterprise Linux	7	Front-End Installation
Ubuntu Server	14.04 (LTS), 16.04 (LTS), 17.04	Front-End Installation
CentOS	7	Front-End Installation
Debian	8, 9	Front-End Installation
MariaDB or MySQL	Version included in the Linux distribution	MySQL Setup
SQLite	Version included in the Linux distribution	Default DB, no configuration needed
Ruby Gems	Versions installed by packages and install_gems utility	front-end installation
Corosync+Pacemaker	Version included in the Linux distribution	Front-end HA Setup

vCenter Nodes

Component	Version	More information
vCenter	5.5/6.0/6.5, managing ESX 5.5/6.0/6.5	vCenter Node Installation

KVM Nodes

Component	Version	More information
RedHat Enterprise Linux	7	KVM Driver
Ubuntu Server	14.04 (LTS) , 16.04 (LTS), 17.04	KVM Driver
CentOS/RHEL	7	KVM Driver
Debian	8, 9	KVM Driver
KVM/Libvirt	Support for version included in the Linux distribution. For CentOS/RedHat the packages from <code>qemu-ev</code> are used.	KVM Node Installation

Open Cloud Networking Infrastructure

Component	Version	More information
etables	Version included in the Linux distribution	Etables
8021q kernel module	Version included in the Linux distribution	802.1Q VLAN
Open vSwitch	Version included in the Linux distribution	Open vSwitch
iproute2	Version included in the Linux distribution	VXLAN

Open Cloud Storage Infrastructure

Component	Version	More information
iSCSI	Version included in the Linux distribution	LVM Drivers
LVM2	Version included in the Linux distribution	LVM Drivers
Ceph	Hammer (LTS) v0.94, Jewel v10.2.0	The Ceph Datastore

Authentication

Component	Version	More information
net-ldap ruby library	0.12.1	LDAP Authentication
openssl	Version included in the Linux distribution	x509 Authentication

Cloud Bursting

Component	Version	More information
aws-sdk	2.5.10	Amazon EC2 Driver
azure	0.7.9	Azure Driver

Note: Generally for all Linux platforms, it is worth noting that gems should be installed with the `install_gems`, avoiding the platform's package version.

2.2.2 Frontend Platform Notes

The following applies to all Front-Ends:

- XML-RPC tuning parameters (MAX_CONN, MAX_CONN_BACKLOG, KEEPALIVE_TIMEOUT, KEEPALIVE_MAX_CONN and TIMEOUT) are only available with packages distributed by us as they are compiled with a newer xmlrpc-c library.
- For **cloud bursting**, a newer nokogiri gem than the one packed by current distros is required. If you are planning to use cloud bursting, you need to install nokogiri \geq 1.4.4 prior to run `install_gems`: `# sudo gem install nokogiri -v 1.4.4`.
- Only **ruby versions \geq 1.9.3 are supported**.

Ubuntu 14.04 Platform Notes

Package ruby-ox shouldn't be installed as it contains a version of the gem incompatible with the CLI

CentOS 7.0 Platform Notes

This distribution lacks some packaged ruby libraries. This makes some components unusable until they are installed. In the front-end, just after package installation these commands should be executed as root to install extra dependencies:

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

When using Apache to serve Sunstone, it is required that you disable or comment the `PrivateTmp=yes` directive in `/usr/lib/systemd/system/httpd.service`.

There is an automatic job that removes all data from `/var/tmp/`, in order to disable this, please edit the `/usr/lib/tmpfiles.d/tmp.conf` and remove the line that removes `/var/tmp`.

There is a bug in libvirt that prevents the use of the save/restore mechanism if `cpu_model` is set to `'host-passthrough'` via RAW. The work around if needed is described in [this issue](#).

Debian 8

Make sure that the packages `ruby-treetop` and `treetop` are not installed before running `install_gems`, as the version of `treetop` that comes packaged in Debian 8 is incompatible with OpenNebula. **OneFlow** requires a version \geq 1.6.3 for `treetop`, packages distributed with Debian 8 includes an older version (1.4.5).

2.2.3 Nodes Platform Notes

The following items apply to all distributions:

- Since OpenNebula 4.14 there is a new monitoring probe that gets information about PCI devices. By default it retrieves all the PCI devices in a host. To limit the PCI devices that it gets info and appear in `onehost show` refer to `kvm_pci_passthrough`.
- When using `qcow2` storage drivers you can make sure that the data is written to disk when doing snapshots setting its `cache` parameter to `writethrough`. This change will make writes slower than other cache modes but safer. To do this edit the file `/etc/one/vmm_exec/vmm_exec_kvm.conf` and change the line for DISK:

```
DISK = [ driver = "qcow2", cache = "writethrough" ]
```

CentOS/RedHat 7.0 Platform Notes

Ruby Dependencies

In order to install ruby dependencies, the Server Optional channel needs to be enabled. Please refer to [RedHat documentation](#) to enable the channel.

Alternatively, use CentOS 7 repositories to install ruby dependencies.

Libvirt Version

The libvirt/qemu packages used in the testing infrastructure are the ones in the `qemu-ev` repository. To add this repository you can install the following packages:

```
# yum install centos-release-qemu-ev
# yum install qemu-kvm-ev
```

Disable PolicyKit for Libvirt

It is recommended that you disable PolicyKit for Libvirt:

```
$ cat /etc/libvirt/libvirtd.conf
...
auth_unix_ro = "none"
auth_unix_rw = "none"
unix_sock_group = "oneadmin"
unix_sock_ro_perms = "0770"
unix_sock_rw_perms = "0770"
...
```

2.2.4 Unsupported Platforms Notes

Warning: Use the following distributions at your own risk. They are not officially supported by OpenNebula.

CentOS 6.5 Usage Platform Notes

- As a front-end, because home directory of oneadmin is located in `/var`, it violates SELinux default policy. So in ssh passwordless configuration you should disable SELinux by setting `SELINUX=disabled` in `/etc/selinux/config`.
- As a node, to accomplish disk hotplugging:
 - to accomplish disk hotplugging, disks need to be attached through SCSI, so their images should have a `DEV_PREFIX="sd"`
 - to accomplish disk hotplugging, VM template that will permit SCSI disk attaches afterwards needs to have an explicitly defined SCSI controller:

```
RAW=[TYPE = "kvm",
      DATA = "<devices><controller type='scsi' index='0' model='virtio-scsi'></>
<-controller></devices>"]
```

```
* due to libvirt version <= 0.10.2, there is a bug in libvirt/qemu attach/detach nic_  
↳ functionality that prevents the reuse of net IDs. This means that after a_  
↳ successful attach/detach NIC, a new attach will fail.
```

Installing on ArchLinux

OpenNebula is available at the Arch User Repository (AUR), [please check the opennebula package page](#).

Installing on Gentoo

There is an ebuild contributed by Thomas Stein in the following repository:

<https://github.com/himbeere/opennebula>

Still, if you want to compile it manually you need to install the `xmlrpc-c` package with threads support, as:

```
USE="threads" emerge xmlrpc-c
```

Installing on Devuan

Packages for Devuan Jessie 1.0 Beta are provided by Alberto Zuin. You can download them at:

<http://downloads.opennebula.org/extra/packages/devuan/>

2.3 Compatibility Guide

This guide is aimed at OpenNebula 5.2.x users and administrators who want to upgrade to the latest version. The following sections summarize the new features and usage changes that should be taken into account, or prone to cause confusion. You can check the upgrade process in the following [section](#)

Visit the [Features list](#) and the [Release Notes](#) for a comprehensive list of what's new in OpenNebula 5.4.

2.3.1 OpenNebula Administrators

OpenNebula Daemon

OpenNebula 5.4 features an implementation of the Raft consensus algorithm. Although current HA deployments may still work it is advised to update the installation to the new system.

Federated zones has been also updated to reduce its requirements and to better integrate it with HA zones. Current Federated zones are no longer compatible and have to be re-imported.

Restricted attributes for Virtual Machines does not support vector attributes (e.g. `USER_INPUTS/CPU`). This also a known issue for prior OpenNebula versions.

EC2 hybrid drivers

Configuration attributes of the EC2 driver has been moved from the `ec2_driver.conf` file to the host attributes. The migration process automatically moves the information to each host template. Any further update of these attributes should be performed directly on each the OpenNebula host.

2.3.2 Developers and Integrators

XML-RPC API

There are no compatibility changes on API calls, there are new RPC methods to expose the new functionality, visit the complete reference for more information.

2.3.3 vCenter

vCenter has experienced major changes from previous releases and this section tries to summarize the most important changes:

vCenter is no longer considered a PUBLIC CLOUD provider

In previous releases, vCenter was considered a PUBLIC CLOUD provider as Amazon EC2 or Microsoft's Azure. This had adverse effects for vCenter e.g the CONTEXT information was not regenerated when a VM was powered off and resume.

As vCenter was considered a PUBLIC CLOUD provider, several attributes were found inside VM templates under the PUBLIC_CLOUD section, that is no longer the case.

vCenter Objects reference

vCenter objects are no longer referenced from OpenNebula by its name. Previously, the names used by OpenNebula should match the name of the resource in vCenter and hence it could not be modified.

In OpenNebula 5.4, vCenter objects (clusters, datastores, port groups, templates and VMs) are referenced using vCenter's Managed Object Reference and vCenter's Instance UUID. So you will be able to change the name of these resources in vCenter and/or OpenNebula and the reference should remain the same.

You have more information about the managed object references in the docs.

This change on how vCenter objects are referenced implies that a migration must be performed in two phases as it will explained later.

Disks and networks that exist in vCenter templates are now visible

OpenNebula 5.4 can represent disks and port groups used by templates when they are imported. In previous versions those disks and NICs were invisible and therefore you could not manage those resources e.g you could not detach a disk or network interface card.

This change implies that a migration tool will identify existing disks and nics and will create OpenNebula images and virtual networks for templates that are running in OpenNebula. A new attribute OPENNEBULA_MANAGED=NO will identify DISKs and NICs that are cloned by vCenter when a VM is deployed, so these elements are not created by OpenNebula. You have more information about this attribute and its limitations here.

Also there's an important change. If you want to import vCenter templates you MUST first import the vCenter datastores where the VMDK files associated to template's virtual hard disks are located. In previous releases you could import a template after a vCenter cluster was importe because existing disks were not visible to OpenNebula.

Names generated by OpenNebula when resources are imported

When vCenter resources are imported, OpenNebula creates objects using a name that cannot be used by another object. To prevent name collisions OpenNebula has changed the way it names the resources:

- Hosts. The vcenter instance name, the datacenter name and a 12 character hash have been added to the vCenter's cluster name. Thanks to this, if you have the same cluster name in different folders or datacenters you may now import them.
- Templates. A 12 character hash has been added to the template name - cluster name.
- Datastores. The vcenter instance name and the datacenter name have been added to the datastore name. This way two datastores with the same name but placed in different datacenters can be imported.
- Networks. The vcenter instance name, the datacenter name and a 12 character hash have been added to the vCenter's port group name.

OpenNebula hosts, datastores and networks are assigned to OpenNebula clusters

An OpenNebula Cluster is a group of Hosts. Clusters can also have associated Datastores and Virtual Networks, and this is how the administrator sets which Hosts have the underlying requirements for each Datastore and Virtual Network configured.

If you recall, a vCenter cluster is represented as a Host. In OpenNebula 5.4 this Host will be added to an OpenNebula cluster and also vCenter datastores and networks will be added to that OpenNebula cluster when they imported.

Thanks to OpenNebula clusters, the `SCHED_REQUIREMENTS="NAME=XXXXX"` attribute that was used to tell OpenNebula's scheduler which host should be used when a VM is deployed, it's no longer needed. You can of course still use the `SCHED_REQUIREMENTS` and `SCHED_DS_REQUIREMENTS` attributes in a template to force how the scheduler behaves but they won't be mandatory. Also it will be easier for the scheduler to select the vCenter cluster based on the DISKS (datastores) and NICs (virtual networks) which are defined in the VM Template.

Please review the import resources section to know more.

KEEP DISKS ON DONE has been deprecated

The Keep Disks on Done option that you could use to prevent OpenNebula from erasing the VM disks upon reaching the done state (either via shutdown or cancel) has been deprecated in an attempt that KVM and vCenter storage management can converge.

If you want to create a copy of one disk you can use the save as option when the VM is in poweroff state. A new image will be created pointing to a new VMDK file.

VCENTER_DATASTORE is no longer used, Scheduler chooses the datastore

In previous OpenNebula releases you could add the `VCENTER_DATASTORE` attribute and force what datastore was going to be used when a VM was cloned before deploying the VM.

That attribute is no longer valid. OpenNebula's scheduler will be the responsible of choosing the datastore where the VM template is going to be cloned in. The scheduler behavior is ruled by the `/etc/one/sched.conf` file and you can override its policy using the `SCHED_REQUIREMENTS` and `SCHED_DS_REQUIREMENTS` attributes.

This also means that the end user cannot choose the target DS using a `USER_INPUT`, this will be addressed in future revisions.

Instantiate as Persistent

Instantiate as Persistent is still available when a VM is instantiated but now the template will detect disks and nics that in previous versions were invisible. Due to these new visible disks, note that you must not detach disks from the VM or resize any disk of the VM once you've deployed it with Instantiate as Persistent, as when the VM is terminated the OpenNebula template that was created before the VM was deployed will differ from the template created in vCenter. Differences between the templates may affect operations on VMs based on unsynced templates.

Datstores now have vCenter credentials inside its templates

In previous releases, datastore templates had an attribute called VCENTER_CLUSTER. That attribute helped OpenNebula to get vCenter credentials from a vcenter cluster (represented as an OpenNebula host) when datastore actions were executed.

OpenNebula 5.4 stores the VCENTER_HOST, VCENTER_USER and VCENTER_PASSWORD attributes inside datstores templates so datstores and hosts (vCenter clusters) are no longer coupled. Datastore can still be monitored even if no vCenter cluster is associated to it in OpenNebula.

Poweroff VMs are destroyed when they are deleted from OpenNebula

In previous releases, the hook `delete_poweroff_vms.rb` was required to clean up VMs that were deleted from OpenNebula when those VMs were in the POWEROFF state. That hook is no longer needed as the VMs will be destroyed in vCenter when a VM is deleted from OpenNebula no matter the state.

VLAN_TAGGED_ID no longer reported when a distributed port group is imported

In previous releases, when a distributed port group was imported some information about the VLAN ID that was assigned to the port group was reported and added to the VLAN_TAGGED_ID attribute.

That information was not accurate and trying to provide the same information for standard port groups would require a significant amount of time and CPU so VLAN ID is no longer retrieved when a distributed port group is imported.

Attributes that have changed its name

In an attempt to ease the task of identifying vCenter related attributes many attributes have changed its name. Here is a table with the old name and the new name.

Old Name	New name
VMWARETOOLS_RUNNING_STATUS	VCENTER_VMWARETOOLS_RUNNING_STATUS
VMWARETOOLS_VERSION	VCENTER_VMWARETOOLS_VERSION
VMWARETOOLS_VERSION	VCENTER_VMWARETOOLS_VERSION
CUSTOMIZATION_SPEC	VCENTER_CUSTOMIZATION_SPEC
GUEST_STATE	VCENTER_GUEST_STATE
ADAPTER_TYPE	VCENTER_ADAPTER_TYPE
ESX_HOST	VCENTER_ESX_HOST
RESOURCE_POOL	VCENTER_RESOURCE_POOL

In general, vCenter attributes will be preceded by the suffix **VCENTER_**

2.3.4 Sunstone

Labels

You will now have persistent labels that will be associated with a view or a user. You can read more about this topic at [Persistent labels](#)

Custom Tags

Added the possibility of adding a new object of type vector inside Custom tags section. In this section Sunstone will enter all the data that has not been collected by the Wizard and that exists in the Template For more information, you can see [Custom tags](#)

2.4 Known Issues

2.4.1 CLI

- [#3037](#) Different ruby versions need different time formats
- [#3337](#) Wrong headers when output is piped for oneacct and oneshowback

2.4.2 Core & System

- [#3020](#) OpenNebula should check the available space in the frontend before doing an undeploy
- [#2880](#) Unicode chars in VM name are truncated
- [#2502](#) deleting image in locked state leaves current operation in progress and files not cleaned
- [#3888](#) An image can be used twice by a VM, but this breaks the used/ready logic
- [#4563](#) onetemplate instantiate –persistent uses always the same name
- [#4154](#) after cold migration custom vnc password is lost
- [#4015](#) Showback records authorization relies on current VM ownership, instead of Showback record owner
- [#2998](#) Deleting an image deletes the record from the database even if the delete fails
- [#5164](#) Hybrid VMs in poweroff -> terminate are not removed from public cloud provider
- [#5146](#) Host Monitoring does not fail if probe does not come back
- [#5029](#) PUBLIC_CLOUD variables do not accept ,
- [#3098](#) onevnet release should return error if the IP to free is not part of the vnet
- [#2943](#) If the driver is not defined in oned.conf new history records are created each time an action is performed
- [#1342](#) Usernames are case sensitive/insensitive depending on the DB backend
- [#2507](#) XML-RPC session string with an space gives a timeout instead of an authentication error
- [#2937](#) If an user removes USE or MANAGE rights from VM it cannot access it anymore

2.4.3 Drivers - Network

- #3093 Review the Open vSwitch flows
- #4005 NIC defaults not honoured in attach NIC

2.4.4 Drivers - Storage

- #1573 If an image download fails, the file is never deleted from the datastore
- #3929 CEPH_HOST not IPv6 friendly
- #3727 Downloads can fail, but still not return in error
- #1763 User is not notified with error_message in rm action / recoverable actions

2.4.5 Drivers - VM

- #4335 vCenter password cannot be longer than 22 characters
- #4514 Spaces not allowed in SOURCE image attribute
- #4169 The state of VMs terminated through the AWS console should not be POWEROFF
- #3696 unsupported actions in the vmm drivers should end with exit 1

2.4.6 OneFlow

- #3134 Service Templates with dynamic networks cannot be instantiated from the CLI, unless the a template file with the required attributes is merged
- #3797 oneflow and oneflow-template ignore the no_proxy environment variable
- #2155 Scheduled policy start_time cannot be defined as a POSIX time number
- #4694 Throw an error an error if a template update is invalid

2.4.7 Marketplace

- #4975 DISPOSE=YES in market_mad/remotes/http/import is not honored

2.4.8 Scheduler

- #1811 If more than one scheduled actions fit in a scheduler cycle, the behavior is unexpected

2.4.9 Sunstone

- #1877 if syslog enabled disable the logs tab in the VM detailed view
- #3796 sunstone ignores the no_proxy environment variable
- #3902 LIMIT_MB is not used to calculate the available DS storage
- #3692 Sunstone image upload - not enough space
- #2867 Sunstone template update does not select images without User Name

- #2801 Template update: placement does not select the hosts/clusters
- #4652 noVNC mouse doesn't work when a touchscreen is present
- #4266 Uploading big image via sunstone got error message
- #5159 Quotas for user should take into account more restrictive group quotas

2.4.10 vCenter

- #4990 importing network on vcenter with vlan doesn't report the vlan id
- #5246 Wrong import of vCenter VM Templates with NICs in Distributed vSwitches or Distributed Ports

2.5 Acknowledgements

The OpenNebula project would like to thank the [community members](#) and [users](#) who have contributed to this software release by being active with the discussions, answering user questions, or providing patches for bugfixes, features and documentation.

The VM Groups functionality, the configurable semantics of the VM operations permissions (ADMIN, MANAGE and USE) and the improved VM history functionality were funded by [BlackBerry](#) in the context of the Fund a Feature Program. The configurable image persistency setting and the new Non-SLAAC IPv6 Address Range were funded by [University of Louvain](#).

UPGRADING

3.1 Overview

Keeping your OpenNebula up-to-date is very important, as you will receive the latest functionality and more importantly, the latest security patches. It is possible to upgrade to the latest OpenNebula release from earlier versions.

3.1.1 Hypervisor Compatibility

The upgrade procedure can be followed regardless of the hypervisor.

3.1.2 How Should I Read This Chapter

You only need to read this chapter if you are upgrading OpenNebula to a newer release. Make sure you have read the *Release Notes* and particularly the *Compatibility* section first.

Upgrading is a sequential procedure. The system will upgrade from the currently installed release to the latest release going through each release (if any). Therefore it's important to read each section.

After the upgrade procedure you can continue using your upgraded OpenNebula Cloud.

3.2 Upgrading from OpenNebula 5.2.x

This section describes the installation procedure for systems that are already running a 5.2.x OpenNebula. The upgrade to OpenNebula 5.4 can be done directly following this section, you don't need to perform intermediate version upgrades. The upgrade will preserve all current users, hosts, resources and configurations; for both Sqlite and MySQL backends.

Read the *Compatibility Guide* and *Release Notes* to know what is new in OpenNebula 5.4.

3.2.1 Upgrading a Federation

If you have two or more 5.2.x OpenNebulas working as a Federation, you need to upgrade all of them. The upgrade for this version has to occur simultaneously on all zones including the master.

The steps to follow are:

1. Stop the MySQL replication in all the slaves and master zone. The MySQL replication is no longer needed.
2. Upgrade the **master** zone to 5.4 version

3. Upgrade the **slaves** zones to 5.4 version
4. Replicate the state of the shared tables from the master zone into each slave zone.

During steps 1 and 2 the slave OpenNebula's can be running, and users can keep accessing them if each zone has a local Sunstone instance. However all the shared database tables (users, groups, ACL...) will not be updated in the slaves zones till step 3 is completed.

To perform the first step, you must stop and reset each slave (and master), remove any configuration attribute for replication in `my.cnf` file and finally restart `mysqld`. Please refer to mysql documentation for more details on how to perform this step.

Then follow this section for the **master zone**. After the master has been updated to 5.4, upgrade each **slave zone** following this same section.

3.2.2 Upgrading from a High Availability deployment

You need to restore the HA deployment according to the new implementation. Upgrade the active OpenNebula instance as described in this section and then regenerate the HA instances as described in the in the HA guide.

3.2.3 Preparation

Before proceeding, make sure you don't have any VMs in a transient state (prolog, migr, epil, save). Wait until these VMs get to a final state (runn, suspended, stopped, done). Check the Managing Virtual Machines guide for more information on the VM life-cycle.

vCenter

Important: Read this section carefully if you are using vCenter!

If you are using vCenter you will need to follow some extra steps while **still running OpenNebula 5.2**.

Follow the *vCenter upgrade 5.2 to 5.4 Pre-migration phase*.

Stop OpenNebula

Stop OpenNebula and any other related services you may have running: OneFlow, EC2, and Sunstone. Use preferably the system tools, like `systemctl` or `service` as `root` in order to stop the services.

3.2.4 Backup

Backup the configuration files located in **/etc/one**. You don't need to do a manual backup of your database, the `onedb` command will perform one automatically.

```
# cp -r /etc/one /etc/one.$(date +%Y-%m-%d')
```

3.2.5 Installation of New Version

Follow the *Platform Notes* and the Installation guide, taking into account that you will already have configured the passwordless ssh access for oneadmin.

Make sure to run the `install_gems` tool, as the new OpenNebula version may have different gem requirements.

Note: If executing `install_gems` you get a message asking to overwrite files for aws executables you can safely answer “yes”.

It is highly recommended **not to keep** your current `oned.conf`, and update the `oned.conf` file shipped with OpenNebula 5.4 to your setup. If for any reason you plan to preserve your current `oned.conf` file, read the *Compatibility Guide* and the complete `oned.conf` reference for 5.4 and 5.4 versions.

3.2.6 Configuration Files Upgrade

If you haven’t modified any configuration files, the package managers will replace the configuration files with their newer versions and no manual intervention is required.

If you have customized **any** configuration files under `/etc/one` we recommend you to follow these steps regardless of the platform/linux distribution.

1. Backup `/etc/one` (already performed)
2. Install the new packages (already performed)
3. Compare the old and new configuration files: `diff -ur /etc/one.YYYY-MM-DD /etc/one`. Or you can use graphical diff-tools like `meld` to compare both directories, which are very useful in this step.
4. Edit the **new** files and port all the customizations from the previous version.
5. You should **never** overwrite the configuration files with older versions.

EC2 Configuration File

The credentials and capacity from `ec2` zones have been moved from its configuration file to the template of the host. You don’t need to update the file `/etc/one/ec2_driver.conf` with the data from the old file. To make this data available to the migrator copy the old configuration file to `/etc/one/ec2_driver.conf.old`:

```
# cp /etc/one.$(date +%Y-%m-%d)/ec2_driver.conf /etc/one/ec2_driver.conf.old
```

After migration you can delete the old file:

```
# rm /etc/one/ec2_driver.conf.old
```

3.2.7 Database Upgrade

vCenter Migration Tool

Important: Read this section carefully if you are using vCenter!

If you are using vCenter you will need to run the vCenter migration tool before running the *onedb upgrade* command from the next section.

Follow the *vCenter upgrade 5.2 to 5.4 Migration phase*.

Perform the Database Upgrade

The database schema and contents are incompatible between versions. The OpenNebula daemon checks the existing DB version, and will fail to start if the version found is not the one expected, with the message 'Database version mismatch'.

You can upgrade the existing DB with the 'onedb' command. You can specify any Sqlite or MySQL database. Check the onedb reference for more information.

Note: Make sure at this point that OpenNebula is not running. If you installed from packages, the service may have been started automatically.

Note: For environments in a Federation: Before upgrading the **master**, make sure that all the slaves have the MySQL replication paused.

After you install the latest OpenNebula, and fix any possible conflicts in oned.conf, you can issue the 'onedb upgrade -v' command. The connection parameters have to be supplied with the command line options, see the onedb manpage for more information. Some examples:

```
$ onedb upgrade -v --sqlite /var/lib/one/one.db
```

```
$ onedb upgrade -v -S localhost -u oneadmin -p oneadmin -d opennebula
```

If everything goes well, you should get an output similar to this one:

```
$ onedb upgrade -v -u oneadmin -d opennebula
MySQL Password:
Version read:
Shared tables 4.11.80 : OpenNebula 4.12.1 daemon bootstrap
Local tables 4.11.80 : OpenNebula 4.12.1 daemon bootstrap

MySQL dump stored in /var/lib/one/mysql_localhost_opennebula.sql
Use 'onedb restore' or restore the DB using the mysql command:
mysql -u user -h server -P port db_name < backup_file

>>> Running migrators for shared tables
...

>>> Running migrators for local tables
...
> Done in 41.93s

Database migrated from 4.11.80 to 4.13.80 (OpenNebula 4.13.80) by onedb command.

Total time: 41.93s
```


Note: Make sure you keep the backup file. If you face any issues, the `onedb` command can restore this backup, but it won't downgrade databases to previous versions.

3.2.8 Check DB Consistency

After the upgrade is completed, you should run the command `onedb fsck`.

First, move the 5.2.x backup file created by the upgrade command to a safe place.

```
$ mv /var/lib/one/mysql_localhost_opennebula.sql /path/for/one-backups/
```

Then execute the following command:

```
$ onedb fsck -S localhost -u oneadmin -p oneadmin -d opennebula
MySQL dump stored in /var/lib/one/mysql_localhost_opennebula.sql
Use 'onedb restore' or restore the DB using the mysql command:
mysql -u user -h server -P port db_name < backup_file

Total errors found: 0
```

3.2.9 Recreate the Federation salves

This section applies only to environments working in a Federation.

For the **master zone**: Snapshot the shared tables using the `onedb` tool. Please refer to the federation guide for more details.

For a **slave zone**: Each slave should be already configured, i.e. `oned.conf` should include the `ZONE_ID` for the slave, auth files present and OpenNebula updated to last version. You only need to restore the shared tables saved in the previous step and start the slave zone.

3.2.10 Reload Start Scripts

Follow this section if you are using a `systemd` base distribution, like CentOS 7+, Ubuntu 16.04+, etc.

In order for the system to re-read the configuration files you should issue the following command after the installation of the new packages:

```
# systemctl daemon-reload
```

3.2.11 Update the Drivers

You should be able now to start OpenNebula as usual, running `service opennebula start` as `root`. At this point, as `oneadmin` user, execute `onehost sync` to update the new drivers in the hosts.

Warning: Doing `onehost sync` is important. If the monitorization drivers are not updated, the hosts will behave erratically.

3.2.12 Testing

OpenNebula will continue the monitoring and management of your previous Hosts and VMs.

As a measure of caution, look for any error messages in `oned.log`, and check that all drivers are loaded successfully. After that, keep an eye on `oned.log` while you issue the `onevm`, `onevnet`, `oneimage`, `oneuser`, `onehost` **list** commands. Try also using the **show** subcommand for some resources.

3.2.13 Restoring the Previous Version

If for any reason you need to restore your previous OpenNebula, follow these steps:

- With OpenNebula 5.4 still installed, restore the DB backup using `'onedb restore -f'`
- Uninstall OpenNebula 5.4, and install again your previous version.
- Copy back the backup of `/etc/one` you did to restore your configuration.

3.2.14 Known Issues

If the MySQL database password contains special characters, such as `@` or `#`, the `onedb` command will fail to connect to it.

The workaround is to temporarily change the `oneadmin`'s password to an ASCII string. The `set password` statement can be used for this:

```
$ mysql -u oneadmin -p
mysql> SET PASSWORD = PASSWORD('newpass');
```

3.3 Upgrading from OpenNebula 5.0.x

This section describes the installation procedure for systems that are already running a 5.0.x OpenNebula. The upgrade to OpenNebula 5.4 can be done directly following this section, you don't need to perform intermediate version upgrades. The upgrade will preserve all current users, hosts, resources and configurations; for both SQLite and MySQL backends.

Read the *Compatibility Guide* and *Release Notes* to know what is new in OpenNebula 5.4.

3.3.1 Upgrading a Federation

If you have two or more 5.0.x OpenNebulas working as a Federation, you need to upgrade all of them. The upgrade does not have to be simultaneous, the slaves can be kept running while the master is upgraded.

The steps to follow are:

1. Stop the MySQL replication in all the slaves
2. Upgrade the **master** OpenNebula
3. Upgrade each **slave**
4. Resume the replication

During the time between steps 1 and 4 the slave OpenNebulas can be running, and users can keep accessing them if each zone has a local Sunstone instance. There is however an important limitation to note: all the shared database tables will not be updated in the slaves zones. This means that new user accounts, password changes, new ACL rules, etc. will not have any effect in the slaves. Read the federation architecture documentation for more details.

It is recommended to upgrade all the slave zones as soon as possible.

To perform the first step, [pause the replication](#) in each **slave MySQL**:

```
mysql> STOP SLAVE;

mysql> SHOW SLAVE STATUS\G

Slave_IO_Running: No
Slave_SQL_Running: No
```

Then follow this section for the **master zone**. After the master has been updated to 5.4, upgrade each **slave zone** following this same section.

3.3.2 Upgrading from a High Availability deployment

The recommended procedure to upgrade two OpenNebulas configured in HA is to follow the upgrade procedure in a specific order. Some steps need to be executed in both servers, and others in just the active node. For the purpose of this section, we will still refer to the *active node* as such even after stopping the cluster, so we run the single node steps always in the same node:

- *Preparation* in the active node.
- *Backup* in the active node.
- Stop the cluster in the active node: `pcs cluster stop`.
- *Installation* in both nodes. Before running `install_gems`, run `gem list > previous_gems.txt` so we can go back to those specific `sinatra` and `rack` gems if the `pcsd` refuses to start.
- *Configuration Files Upgrade* in the active node.
- *Database Upgrade* in the active node.
- *Check DB Consistency* in the active node.
- *Reload Start Scripts in CentOS 7* in both nodes.
- Start the cluster in the active node.

3.3.3 Preparation

Before proceeding, make sure you don't have any VMs in a transient state (prolog, migr, epil, save). Wait until these VMs get to a final state (runn, suspended, stopped, done). Check the [Managing Virtual Machines](#) guide for more information on the VM life-cycle.

Stop OpenNebula and any other related services you may have running: OneFlow, EC2, and Sunstone. Use preferably the system tools, like `systemctl` or `service` as `root` in order to stop the services.

3.3.4 Backup

Backup the configuration files located in `/etc/one`. You don't need to do a manual backup of your database, the `onedb` command will perform one automatically.

```
# cp -r /etc/one /etc/one.$(date +%Y-%m-%d')
```

3.3.5 Installation

Follow the *Platform Notes* and the Installation guide, taking into account that you will already have configured the passwordless ssh access for oneadmin.

Make sure to run the `install_gems` tool, as the new OpenNebula version may have different gem requirements.

Note: If executing `install_gems` you get a message asking to overwrite files for aws executables you can safely answer “yes”.

It is highly recommended **not to keep** your current `oned.conf`, and update the `oned.conf` file shipped with OpenNebula 5.4 to your setup. If for any reason you plan to preserve your current `oned.conf` file, read the *Compatibility Guide* and the complete `oned.conf` reference for 5.0 and 5.4 versions.

3.3.6 Configuration Files Upgrade

If you haven’t modified any configuration files, the package managers will replace the configuration files with their newer versions and no manual intervention is required.

If you have customized **any** configuration files under `/etc/one` we recommend you to follow these steps regardless of the platform/linux distribution.

1. Backup `/etc/one` (already performed)
2. Install the new packages (already performed)
3. Compare the old and new configuration files: `diff -ur /etc/one.YYYY-MM-DD /etc/one`. Or you can use graphical diff-tools like `meld` to compare both directories, which are very useful in this step.
4. Edit the **new** files and port all the customizations from the previous version.
5. You should **never** overwrite the configuration files with older versions.

3.3.7 Database Upgrade

The database schema and contents are incompatible between versions. The OpenNebula daemon checks the existing DB version, and will fail to start if the version found is not the one expected, with the message ‘Database version mismatch’.

You can upgrade the existing DB with the ‘`onedb`’ command. You can specify any Sqlite or MySQL database. Check the `onedb` reference for more information.

Warning: Make sure at this point that OpenNebula is not running. If you installed from packages, the service may have been started automatically.

Warning: For environments in a Federation: Before upgrading the **master**, make sure that all the slaves have the MySQL replication paused.

After you install the latest OpenNebula, and fix any possible conflicts in `oned.conf`, you can issue the `'onedb upgrade -v'` command. The connection parameters have to be supplied with the command line options, see the `onedb` manpage for more information. Some examples:

```
$ onedb upgrade -v --sqlite /var/lib/one/one.db
```

```
$ onedb upgrade -v -S localhost -u oneadmin -p oneadmin -d opennebula
```

If everything goes well, you should get an output similar to this one:

```
$ onedb upgrade -v -u oneadmin -d opennebula
MySQL Password:
Version read:
Shared tables 4.11.80 : OpenNebula 4.12.1 daemon bootstrap
Local tables  4.11.80 : OpenNebula 4.12.1 daemon bootstrap

MySQL dump stored in /var/lib/one/mysql_localhost_opennebula.sql
Use 'onedb restore' or restore the DB using the mysql command:
mysql -u user -h server -P port db_name < backup_file

>>> Running migrators for shared tables
...

>>> Running migrators for local tables
...
> Done in 41.93s

Database migrated from 4.11.80 to 4.13.80 (OpenNebula 4.13.80) by onedb command.

Total time: 41.93s
```

Note: Make sure you keep the backup file. If you face any issues, the `onedb` command can restore this backup, but it won't downgrade databases to previous versions.

3.3.8 Check DB Consistency

After the upgrade is completed, you should run the command `onedb fsck`.

First, move the 5.0.x backup file created by the upgrade command to a safe place.

```
$ mv /var/lib/one/mysql_localhost_opennebula.sql /path/for/one-backups/
```

Then execute the following command:

```
$ onedb fsck -S localhost -u oneadmin -p oneadmin -d opennebula
MySQL dump stored in /var/lib/one/mysql_localhost_opennebula.sql
Use 'onedb restore' or restore the DB using the mysql command:
mysql -u user -h server -P port db_name < backup_file

Total errors found: 0
```

3.3.9 Resume the Federation

This section applies only to environments working in a Federation.

For the **master zone**: This step is not necessary.

For a **slave zone**: The MySQL replication must be resumed now.

Warning: Do not copy the server-id from this example, each slave should already have a unique ID.

- Start the **slave MySQL** process and check its status. It may take a while to copy and apply all the pending commands.

```
mysql> START SLAVE;
mysql> SHOW SLAVE STATUS\G
```

The `SHOW SLAVE STATUS` output will provide detailed information, but to confirm that the slave is connected to the master MySQL, take a look at these columns:

```
Slave_IO_State: Waiting for master to send event
Slave_IO_Running: Yes
Slave_SQL_Running: Yes
```

3.3.10 Reload Start Scripts in CentOS 7

In order for the system to re-read the configuration files you should issue the following command after the installation of the new packages:

```
# systemctl daemon-reload
```

3.3.11 Update the Drivers

You should be able now to start OpenNebula as usual, running `service opennebula start` as root. At this point, as `oneadmin` user, execute `onehost sync` to update the new drivers in the hosts.

Warning: Doing `onehost sync` is important. If the monitorization drivers are not updated, the hosts will behave erratically.

3.3.12 Testing

OpenNebula will continue the monitoring and management of your previous Hosts and VMs.

As a measure of caution, look for any error messages in `oned.log`, and check that all drivers are loaded successfully. After that, keep an eye on `oned.log` while you issue the `onevm`, `onevnet`, `oneimage`, `oneuser`, `onehost list` commands. Try also using the **show** subcommand for some resources.

3.3.13 Restoring the Previous Version

If for any reason you need to restore your previous OpenNebula, follow these steps:

- With OpenNebula 5.4 still installed, restore the DB backup using ‘onedb restore -f’
- Uninstall OpenNebula 5.4, and install again your previous version.
- Copy back the backup of /etc/one you did to restore your configuration.

3.3.14 Known Issues

If the MySQL database password contains special characters, such as @ or #, the onedb command will fail to connect to it.

The workaround is to temporarily change the oneadmin’s password to an ASCII string. The `set password` statement can be used for this:

```
$ mysql -u oneadmin -p
mysql> SET PASSWORD = PASSWORD('newpass');
```

3.4 Upgrading from OpenNebula 4.14.x

This section describes the installation procedure for systems that are already running a 4.14.x OpenNebula. The upgrade to OpenNebula 5.4 can be done directly following this section, you don’t need to perform intermediate version upgrades. The upgrade will preserve all current users, hosts, resources and configurations; for both Sqlite and MySQL backends.

Read the Compatibility Guide for [5.0](#) and [5.4](#), and the [Release Notes](#) to know what is new in OpenNebula 5.4.

3.4.1 Upgrading a Federation

If you have two or more 4.14.x OpenNebulas working as a Federation, you need to upgrade all of them. The upgrade does not have to be simultaneous, the slaves can be kept running while the master is upgraded.

The steps to follow are:

1. Stop the MySQL replication in all the slaves
2. Upgrade the **master** OpenNebula
3. Upgrade each **slave**
4. Resume the replication

During the time between steps 1 and 4 the slave OpenNebulas can be running, and users can keep accessing them if each zone has a local Sunstone instance. There is however an important limitation to note: all the shared database tables will not be updated in the slaves zones. This means that new user accounts, password changes, new ACL rules, etc. will not have any effect in the slaves. Read the federation architecture documentation for more details.

It is recommended to upgrade all the slave zones as soon as possible.

To perform the first step, [pause the replication](#) in each **slave MySQL**:

```
mysql> STOP SLAVE;
mysql> SHOW SLAVE STATUS\G

Slave_IO_Running: No
Slave_SQL_Running: No
```

Then follow this section for the **master zone**. After the master has been updated to 5.4, upgrade each **slave zone** following this same section.

3.4.2 Upgrading from a High Availability deployment

The recommended procedure to upgrade two OpenNebulas configured in HA is to follow the upgrade procedure in a specific order. Some steps need to be executed in both servers, and others in just the active node. For the purpose of this section, we will still refer to the *active node* as such even after stopping the cluster, so we run the single node steps always in the same node:

- *Preparation* in the active node.
- *Backup* in the active node.
- Stop the cluster in the active node: `pcs cluster stop`.
- *Installation* in both nodes. Before running `install_gems`, run `gem list > previous_gems.txt` so we can go back to those specific `sinatra` and `rack` gems if the `pcsd` refuses to start.
- *Configuration Files Upgrade* in the active node.
- *Database Upgrade* in the active node.
- *Check DB Consistency* in the active node.
- *Reload Start Scripts in CentOS 7* in both nodes.
- Start the cluster in the active node.

3.4.3 Preparation

Before proceeding, make sure you don't have any VMs in a transient state (`prolog`, `migr`, `epil`, `save`). Wait until these VMs get to a final state (`runn`, `suspended`, `stopped`, `done`). Check the Managing Virtual Machines guide for more information on the VM life-cycle.

The network drivers since OpenNebula 5.0 are located in the Virtual Network, rather than in the host. The upgrade process may ask you questions about your existing VMs, Virtual Networks and hosts, and as such it is wise to have the following information saved beforehand, since in the upgrade process OpenNebula will be stopped.

```
$ onevnet list -x > networks.txt
$ onehost list -x > hosts.txt
$ onevm list -x > vms.txt
```

The list of valid network drivers since 5.0 Wizard are:

- `802.1Q`
- `dummy`
- `eatables`
- `fw`
- `ovswitch`
- `vxlan`

Stop OpenNebula and any other related services you may have running: OneFlow, EC2, and Sunstone. Use preferably the system tools, like `systemctl` or `service` as `root` in order to stop the services.

3.4.4 Backup

Backup the configuration files located in `/etc/one`. You don't need to do a manual backup of your database, the `onedb` command will perform one automatically.

```
# cp -r /etc/one /etc/one.$(date +%Y-%m-%d')
```

3.4.5 Installation

Follow the *Platform Notes* and the Installation guide, taking into account that you will already have configured the passwordless ssh access for `oneadmin`.

Make sure to run the `install_gems` tool, as the new OpenNebula version may have different gem requirements.

Note: If executing `install_gems` you get a message asking to overwrite files for aws executables you can safely answer “yes”.

It is highly recommended **not to keep** your current `oned.conf`, and update the `oned.conf` file shipped with OpenNebula 5.4 to your setup. If for any reason you plan to preserve your current `oned.conf` file, read the *Compatibility Guide* and the complete `oned.conf` reference for 4.14 and 5.4 versions.

3.4.6 Configuration Files Upgrade

If you haven't modified any configuration files, the package managers will replace the configuration files with their newer versions and no manual intervention is required.

If you have customized **any** configuration files under `/etc/one` we recommend you to follow these steps regardless of the platform/linux distribution.

1. Backup `/etc/one` (already performed)
2. Install the new packages (already performed)
3. Compare the old and new configuration files: `diff -ur /etc/one.YYYY-MM-DD /etc/one`. Or you can use graphical diff-tools like `meld` to compare both directories, which are very useful in this step.
4. Edit the **new** files and port all the customizations from the previous version.
5. You should **never** overwrite the configuration files with older versions.

3.4.7 Database Upgrade

The database schema and contents are incompatible between versions. The OpenNebula daemon checks the existing DB version, and will fail to start if the version found is not the one expected, with the message ‘Database version mismatch’.

You can upgrade the existing DB with the ‘`onedb`’ command. You can specify any SQLite or MySQL database. Check the `onedb` reference for more information.

Warning: Make sure at this point that OpenNebula is not running. If you installed from packages, the service may have been started automatically.

Warning: For environments in a Federation: Before upgrading the **master**, make sure that all the slaves have the MySQL replication paused.

After you install the latest OpenNebula, and fix any possible conflicts in `oned.conf`, you can issue the `'onedb upgrade -v'` command. The connection parameters have to be supplied with the command line options, see the `onedb manpage` for more information. Some examples:

```
$ onedb upgrade -v --sqlite /var/lib/one/one.db
```

```
$ onedb upgrade -v -S localhost -u oneadmin -p oneadmin -d opennebula
```

If everything goes well, you should get an output similar to this one:

```
$ onedb upgrade -v -u oneadmin -d opennebula
MySQL Password:
Version read:
Shared tables 4.11.80 : OpenNebula 4.12.1 daemon bootstrap
Local tables 4.11.80 : OpenNebula 4.12.1 daemon bootstrap

MySQL dump stored in /var/lib/one/mysql_localhost_opennebula.sql
Use 'onedb restore' or restore the DB using the mysql command:
mysql -u user -h server -P port db_name < backup_file

>>> Running migrators for shared tables
...

>>> Running migrators for local tables
...
> Done in 41.93s

Database migrated from 4.11.80 to 4.13.80 (OpenNebula 4.13.80) by onedb command.

Total time: 41.93s
```

Note: Make sure you keep the backup file. If you face any issues, the `onedb` command can restore this backup, but it won't downgrade databases to previous versions.

3.4.8 Check DB Consistency

After the upgrade is completed, you should run the command `onedb fsck`.

First, move the 4.14 backup file created by the upgrade command to a safe place.

```
$ mv /var/lib/one/mysql_localhost_opennebula.sql /path/for/one-backups/
```

Then execute the following command:

```
$ onedb fsck -S localhost -u oneadmin -p oneadmin -d opennebula
MySQL dump stored in /var/lib/one/mysql_localhost_opennebula.sql
Use 'onedb restore' or restore the DB using the mysql command:
mysql -u user -h server -P port db_name < backup_file
```

```
Total errors found: 0
```

3.4.9 Resume the Federation

This section applies only to environments working in a Federation.

For the **master zone**: This step is not necessary.

For a **slave zone**: The MySQL replication must be resumed now.

- First, add two new tables, `marketplace_pool` and `marketplaceapp_pool`, to the replication configuration.

Warning: Do not copy the server-id from this example, each slave should already have a unique ID.

```
# vi /etc/my.cnf
[mysqld]
server-id          = 100
replicate-do-table = opennebula.user_pool
replicate-do-table = opennebula.group_pool
replicate-do-table = opennebula.vdc_pool
replicate-do-table = opennebula.zone_pool
replicate-do-table = opennebula.db_versioning
replicate-do-table = opennebula.acl
replicate-do-table = opennebula.marketplace_pool
replicate-do-table = opennebula.marketplaceapp_pool

# service mysqld restart
```

- Start the **slave MySQL** process and check its status. It may take a while to copy and apply all the pending commands.

```
mysql> START SLAVE;
mysql> SHOW SLAVE STATUS\G
```

The `SHOW SLAVE STATUS` output will provide detailed information, but to confirm that the slave is connected to the master MySQL, take a look at these columns:

```
Slave_IO_State: Waiting for master to send event
Slave_IO_Running: Yes
Slave_SQL_Running: Yes
```

3.4.10 Reload Start Scripts in CentOS 7

In order for the system to re-read the configuration files you should issue the following command after the installation of the new packages:

```
# systemctl daemon-reload
```

3.4.11 Update the Drivers

You should be able now to start OpenNebula as usual, running `service opennebula start` as `root`. At this point, as `oneadmin` user, execute `onehost sync` to update the new drivers in the hosts.

Warning: Doing `onehost sync` is important. If the monitorization drivers are not updated, the hosts will behave erratically.

3.4.12 Create the Virtual Router ACL Rule

There is a new kind of resource introduced in 5.0: Virtual Routers. If you want your existing users to be able to create their own Virtual Routers, create the following ACL Rule:

```
$ oneacl create "* VROUTER/* CREATE *"
```

Note: For environments in a Federation: This command needs to be executed only once in the master zone, after it is upgraded to 5.4.

3.4.13 Testing

OpenNebula will continue the monitoring and management of your previous Hosts and VMs.

As a measure of caution, look for any error messages in `oned.log`, and check that all drivers are loaded successfully. After that, keep an eye on `oned.log` while you issue the `onevm`, `onevnet`, `oneimage`, `oneuser`, `onehost list` commands. Try also using the **show** subcommand for some resources.

3.4.14 Restoring the Previous Version

If for any reason you need to restore your previous OpenNebula, follow these steps:

- With OpenNebula 5.4 still installed, restore the DB backup using `'onedb restore -f'`
- Uninstall OpenNebula 5.4, and install again your previous version.
- Copy back the backup of `/etc/one` you did to restore your configuration.

3.4.15 Known Issues

If the MySQL database password contains special characters, such as `@` or `#`, the `onedb` command will fail to connect to it.

The workaround is to temporarily change the `oneadmin`'s password to an ASCII string. The `set password` statement can be used for this:

```
$ mysql -u oneadmin -p
mysql> SET PASSWORD = PASSWORD('newpass');
```

3.5 Upgrading from OpenNebula 4.12.x

This section describes the installation procedure for systems that are already running a 4.12.x OpenNebula. The upgrade to OpenNebula 5.4 can be done directly following this section, you don't need to perform intermediate version upgrades. The upgrade will preserve all current users, hosts, resources and configurations; for both Sqlite and MySQL backends.

Read the Compatibility Guide for 4.14, 5.0 and 5.4, and the [Release Notes](#) to know what is new in OpenNebula 5.4.

3.5.1 Upgrading a Federation

If you have two or more 4.12.x OpenNebulas working as a Federation, you need to upgrade all of them. The upgrade does not have to be simultaneous, the slaves can be kept running while the master is upgraded.

The steps to follow are:

1. Stop the MySQL replication in all the slaves
2. Upgrade the **master** OpenNebula
3. Upgrade each **slave**
4. Resume the replication

During the time between steps 1 and 4 the slave OpenNebulas can be running, and users can keep accessing them if each zone has a local Sunstone instance. There is however an important limitation to note: all the shared database tables will not be updated in the slaves zones. This means that new user accounts, password changes, new ACL rules, etc. will not have any effect in the slaves. Read the federation architecture documentation for more details.

It is recommended to upgrade all the slave zones as soon as possible.

To perform the first step, [pause the replication](#) in each **slave MySQL**:

```
mysql> STOP SLAVE;

mysql> SHOW SLAVE STATUS\G

Slave_IO_Running: No
Slave_SQL_Running: No
```

Then follow this section for the **master zone**. After the master has been updated to 5.4, upgrade each **slave zone** following this same section.

3.5.2 Upgrading from a High Availability deployment

The recommended procedure to upgrade two OpenNebulas configured in HA is to follow the upgrade procedure in a specific order. Some steps need to be executed in both servers, and others in just the active node. For the purpose of this section, we will still refer to the *active node* as such even after stopping the cluster, so we run the single node steps always in the same node:

- *Preparation* in the active node.
- *Backup* in the active node.
- Stop the cluster in the active node: `pcs cluster stop`.
- *Installation* in both nodes. Before running `install_gems`, run `gem list > previous_gems.txt` so we can go back to those specific `sinatra` and `rack` gems if the `pcsd` refuses to start.

- *Configuration Files Upgrade* in the active node.
- *Database Upgrade* in the active node.
- *Check DB Consistency* in the active node.
- *Reload Start Scripts in CentOS 7* in both nodes.
- Start the cluster in the active node.

3.5.3 Preparation

Before proceeding, make sure you don't have any VMs in a transient state (prolog, migr, epil, save). Wait until these VMs get to a final state (runn, suspended, stopped, done). Check the Managing Virtual Machines guide for more information on the VM life-cycle.

Warning: In 4.14 the FAILED state disappears. You need to delete all the VMs in this state **before** the new version is installed.

The network drivers since OpenNebula 5.0 are located in the Virtual Network, rather than in the host. The upgrade process may ask you questions about your existing VMs, Virtual Networks and hosts, and as such it is wise to have the following information saved beforehand, since in the upgrade process OpenNebula will be stopped.

```
$ onevnet list -x > networks.txt
$ onehost list -x > hosts.txt
$ onevm list -x > vms.txt
```

The list of valid network drivers since 5.0 Wizard are:

- 802.1Q
- dummy
- ebttables
- fw
- ovswitch
- vxlan

Stop OpenNebula and any other related services you may have running: OneFlow, EC2, and Sunstone. Use preferably the system tools, like *systemctl* or *service* as *root* in order to stop the services.

3.5.4 Backup

Backup the configuration files located in **/etc/one**. You don't need to do a manual backup of your database, the `onedb` command will perform one automatically.

```
# cp -r /etc/one /etc/one.$(date +%Y-%m-%d')
```

3.5.5 Installation

Follow the *Platform Notes* and the Installation guide, taking into account that you will already have configured the passwordless ssh access for `oneadmin`.

Make sure to run the `install_gems` tool, as the new OpenNebula version may have different gem requirements.

It is highly recommended **not to keep** your current `oned.conf`, and update the `oned.conf` file shipped with OpenNebula 5.4 to your setup. If for any reason you plan to preserve your current `oned.conf` file, read the [Compatibility Guide](#) and the complete `oned.conf` reference for 4.12 and 5.0 versions.

3.5.6 Configuration Files Upgrade

If you haven't modified any configuration files, the package managers will replace the configuration files with their newer versions and no manual intervention is required.

If you have customized **any** configuration files under `/etc/one` we recommend you to follow these steps regardless of the platform/linux distribution.

1. Backup `/etc/one` (already performed)
2. Install the new packages (already performed)
3. Compare the old and new configuration files: `diff -ur /etc/one.YYYY-MM-DD /etc/one`. Or you can use graphical diff-tools like `meld` to compare both directories, which are very useful in this step.
4. Edit the **new** files and port all the customizations from the previous version.
5. You should **never** overwrite the configuration files with older versions.

3.5.7 Database Upgrade

The database schema and contents are incompatible between versions. The OpenNebula daemon checks the existing DB version, and will fail to start if the version found is not the one expected, with the message 'Database version mismatch'.

You can upgrade the existing DB with the 'onedb' command. You can specify any SQLite or MySQL database. Check the `onedb` reference for more information.

Warning: Make sure at this point that OpenNebula is not running. If you installed from packages, the service may have been started automatically.

Warning: For environments in a Federation: Before upgrading the **master**, make sure that all the slaves have the MySQL replication paused.

After you install the latest OpenNebula, and fix any possible conflicts in `oned.conf`, you can issue the 'onedb upgrade -v' command. The connection parameters have to be supplied with the command line options, see the `onedb` manpage for more information. Some examples:

```
$ onedb upgrade -v --sqlite /var/lib/one/one.db
```

```
$ onedb upgrade -v -S localhost -u oneadmin -p oneadmin -d opennebula
```

If everything goes well, you should get an output similar to this one:

```
$ onedb upgrade -v -u oneadmin -d opennebula
MySQL Password:
Version read:
```

```

Shared tables 4.11.80 : OpenNebula 4.12.1 daemon bootstrap
Local tables 4.11.80 : OpenNebula 4.12.1 daemon bootstrap

MySQL dump stored in /var/lib/one/mysql_localhost_opennebula.sql
Use 'onedb restore' or restore the DB using the mysql command:
mysql -u user -h server -P port db_name < backup_file

>>> Running migrators for shared tables
Database already uses version 4.11.80

>>> Running migrators for local tables
> Running migrator /usr/lib/one/ruby/onedb/local/4.11.80_to_4.13.80.rb
*****
* WARNING WARNING WARNING WARNING WARNING WARNING WARNING *
*****

OpenNebula 4.13.80 improves the management of FAILED VMs
Please remove (onevm delete) any FAILED VM before continuing.

*****
* WARNING WARNING WARNING WARNING WARNING WARNING WARNING *
*****

The scheduler (and oned) has been update to enforce access
rights on system datastores. This new version also checks that
the user can access the System DS.
This *may require* to update system DS rights of your cloud

Do you want to proceed ? [y/N]y
> Done in 41.93s

Database migrated from 4.11.80 to 4.13.80 (OpenNebula 4.13.80) by onedb command.

Total time: 41.93s

```

Note: Make sure you keep the backup file. If you face any issues, the onedb command can restore this backup, but it won't downgrade databases to previous versions.

3.5.8 Check DB Consistency

After the upgrade is completed, you should run the command `onedb fsck`.

First, move the 4.12 backup file created by the upgrade command to a safe place.

```
$ mv /var/lib/one/mysql_localhost_opennebula.sql /path/for/one-backups/
```

Then execute the following command:

```

$ onedb fsck -S localhost -u oneadmin -p oneadmin -d opennebula
MySQL dump stored in /var/lib/one/mysql_localhost_opennebula.sql
Use 'onedb restore' or restore the DB using the mysql command:
mysql -u user -h server -P port db_name < backup_file

```



```
Total errors found: 0
```

3.5.9 Resume the Federation

This section applies only to environments working in a Federation.

For the **master zone**: This step is not necessary.

For a **slave zone**: The MySQL replication must be resumed now.

- First, add two new tables, `marketplace_pool` and `marketplaceapp_pool`, to the replication configuration.

Warning: Do not copy the server-id from this example, each slave should already have a unique ID.

```
# vi /etc/my.cnf
[mysqld]
server-id           = 100
replicate-do-table = opennebula.user_pool
replicate-do-table = opennebula.group_pool
replicate-do-table = opennebula.vdc_pool
replicate-do-table = opennebula.zone_pool
replicate-do-table = opennebula.db_versioning
replicate-do-table = opennebula.acl
replicate-do-table = opennebula.marketplace_pool
replicate-do-table = opennebula.marketplaceapp_pool

# service mysqld restart
```

- Start the **slave MySQL** process and check its status. It may take a while to copy and apply all the pending commands.

```
mysql> START SLAVE;
mysql> SHOW SLAVE STATUS\G
```

The `SHOW SLAVE STATUS` output will provide detailed information, but to confirm that the slave is connected to the master MySQL, take a look at these columns:

```
Slave_IO_State: Waiting for master to send event
Slave_IO_Running: Yes
Slave_SQL_Running: Yes
```

3.5.10 Reload Start Scripts in CentOS 7

In order for the system to re-read the configuration files you should issue the following command after the installation of the new packages:

```
# systemctl daemon-reload
```

3.5.11 Update the Drivers

You should be able now to start OpenNebula as usual, running ‘one start’ as oneadmin. At this point, execute `onehost sync` to update the new drivers in the hosts.

Warning: Doing `onehost sync` is important. If the monitorization drivers are not updated, the hosts will behave erratically.

3.5.12 Default Auth

If you are using LDAP as default auth driver, you will need to update `/etc/one/oned.conf` and set the new `DEFAULT_AUTH` variable:

```
DEFAULT_AUTH = "ldap"
```

3.5.13 Create the Virtual Router ACL Rule

There is a new kind of resource introduced in 5.0: Virtual Routers. If you want your existing users to be able to create their own Virtual Routers, create the following ACL Rule:

```
$ oneacl create "* VROUTER/* CREATE *"
```

Note: For environments in a Federation: This command needs to be executed only once in the master zone, after it is upgraded to 5.4.

3.5.14 Testing

OpenNebula will continue the monitoring and management of your previous Hosts and VMs.

As a measure of caution, look for any error messages in `oned.log`, and check that all drivers are loaded successfully. After that, keep an eye on `oned.log` while you issue the `onevm`, `onevnet`, `oneimage`, `oneuser`, `onehost list` commands. Try also using the **show** subcommand for some resources.

3.5.15 Restoring the Previous Version

If for any reason you need to restore your previous OpenNebula, follow these steps:

- With OpenNebula 5.4 still installed, restore the DB backup using ‘`onedb restore -f`’
- Uninstall OpenNebula 5.4, and install again your previous version.
- Copy back the backup of `/etc/one` you did to restore your configuration.

3.5.16 Known Issues

If the MySQL database password contains special characters, such as `@` or `#`, the `onedb` command will fail to connect to it.

The workaround is to temporarily change the oneadmin's password to an ASCII string. The `set password` statement can be used for this:

```
$ mysql -u oneadmin -p
mysql> SET PASSWORD = PASSWORD('newpass');
```

3.6 Upgrading from OpenNebula 4.10.x

This section describes the installation procedure for systems that are already running a 4.10.x OpenNebula. The upgrade to OpenNebula 5.4 can be done directly following this section, you don't need to perform intermediate version upgrades. The upgrade will preserve all current users, hosts, resources and configurations; for both Sqlite and MySQL backends.

Read the Compatibility Guide for 4.12, 4.14, 5.0 and 5.4, and the [Release Notes](#) to know what is new in OpenNebula 5.4.

3.6.1 Upgrading a Federation

If you have two or more 4.10.x OpenNebulas working as a Federation, you need to upgrade all of them. The upgrade does not have to be simultaneous, the slaves can be kept running while the master is upgraded.

The steps to follow are:

1. Stop the MySQL replication in all the slaves
2. Upgrade the **master** OpenNebula
3. Upgrade each **slave**
4. Resume the replication

During the time between steps 1 and 4 the slave OpenNebulas can be running, and users can keep accessing them if each zone has a local Sunstone instance. There is however an important limitation to note: all the shared database tables will not be updated in the slaves zones. This means that new user accounts, password changes, new ACL rules, etc. will not have any effect in the slaves. Read the federation architecture documentation for more details.

It is recommended to upgrade all the slave zones as soon as possible.

To perform the first step, [pause the replication](#) in each **slave MySQL**:

```
mysql> STOP SLAVE;

mysql> SHOW SLAVE STATUS\G

Slave_IO_Running: No
Slave_SQL_Running: No
```

Then follow this section for the **master zone**. After the master has been updated to 5.4, upgrade each **slave zone** following this same section.

3.6.2 Preparation

Before proceeding, make sure you don't have any VMs in a transient state (prolog, migr, epil, save). Wait until these VMs get to a final state (runn, suspended, stopped, done). Check the [Managing Virtual Machines](#) guide for more information on the VM life-cycle.

Warning: In 4.14 the FAILED state disappears. You need to delete all the VMs in this state **before** the new version is installed.

The network drivers since OpenNebula 5.0 are located in the Virtual Network, rather than in the host. The upgrade process may ask you questions about your existing VMs, Virtual Networks and hosts, and as such it is wise to have the following information saved beforehand, since in the upgrade process OpenNebula will be stopped.

```
$ onevnet list -x > networks.txt
$ onehost list -x > hosts.txt
$ onevm list -x > vms.txt
```

The list of valid network drivers since 5.0 Wizard are:

- 802.1Q
- dummy
- ebttables
- fw
- ovswitch
- vxlan

Stop OpenNebula and any other related services you may have running: OneFlow, EC2, and Sunstone. Use preferably the system tools, like *systemctl* or *service as root* in order to stop the services.

3.6.3 Backup

Backup the configuration files located in **/etc/one**. You don't need to do a manual backup of your database, the `onedb` command will perform one automatically.

```
# cp -r /etc/one /etc/one.$(date +%Y-%m-%d')
```

3.6.4 Installation

Follow the *Platform Notes* and the Installation guide, taking into account that you will already have configured the passwordless ssh access for `oneadmin`.

Make sure to run the `install_gems` tool, as the new OpenNebula version may have different gem requirements.

It is highly recommended **not to keep** your current `oned.conf`, and update the `oned.conf` file shipped with OpenNebula 5.4 to your setup. If for any reason you plan to preserve your current `oned.conf` file, read the *Compatibility Guide* and the complete `oned.conf` reference for 4.10 and 5.0 versions.

3.6.5 Configuration Files Upgrade

If you haven't modified any configuration files, the package managers will replace the configuration files with their newer versions and no manual intervention is required.

If you have customized **any** configuration files under `/etc/one` we recommend you to follow these steps regardless of the platform/linux distribution.

1. Backup `/etc/one` (already performed)

2. Install the new packages (already performed)
3. Compare the old and new configuration files: `diff -ur /etc/one.YYYY-MM-DD /etc/one`. Or you can use graphical diff-tools like `meld` to compare both directories, which are very useful in this step.
4. Edit the **new** files and port all the customizations from the previous version.
5. You should **never** overwrite the configuration files with older versions.

3.6.6 Database Upgrade

The database schema and contents are incompatible between versions. The OpenNebula daemon checks the existing DB version, and will fail to start if the version found is not the one expected, with the message 'Database version mismatch'.

You can upgrade the existing DB with the 'onedb' command. You can specify any Sqlite or MySQL database. Check the onedb reference for more information.

Warning: Make sure at this point that OpenNebula is not running. If you installed from packages, the service may have been started automatically.

Warning: For environments in a Federation: Before upgrading the **master**, make sure that all the slaves have the MySQL replication paused.

After you install the latest OpenNebula, and fix any possible conflicts in `oned.conf`, you can issue the 'onedb upgrade -v' command. The connection parameters have to be supplied with the command line options, see the onedb manpage for more information. Some examples:

```
$ onedb upgrade -v --sqlite /var/lib/one/one.db
```

```
$ onedb upgrade -v -S localhost -u oneadmin -p oneadmin -d opennebula
```

If everything goes well, you should get an output similar to this one:

```
$ onedb upgrade -v -u oneadmin -d opennebula
MySQL Password:
Version read:
Shared tables 4.4.0 : OpenNebula 4.4.0 daemon bootstrap
Local tables 4.4.0 : OpenNebula 4.4.0 daemon bootstrap

>>> Running migrators for shared tables
> Running migrator /usr/lib/one/ruby/onedb/shared/4.4.0_to_4.4.1.rb
> Done in 0.00s

> Running migrator /usr/lib/one/ruby/onedb/shared/4.4.1_to_4.5.80.rb
> Done in 0.75s

Database migrated from 4.4.0 to 4.5.80 (OpenNebula 4.5.80) by onedb command.

>>> Running migrators for local tables
Database already uses version 4.5.80
Total time: 0.77s
```

Note: Make sure you keep the backup file. If you face any issues, the `onedb` command can restore this backup, but it won't downgrade databases to previous versions.

3.6.7 Check DB Consistency

After the upgrade is completed, you should run the command `onedb fsck`.

First, move the 4.10 backup file created by the upgrade command to a safe place.

```
$ mv /var/lib/one/mysql_localhost_opennebula.sql /path/for/one-backups/
```

Then execute the following command:

```
$ onedb fsck -S localhost -u oneadmin -p oneadmin -d opennebula
MySQL dump stored in /var/lib/one/mysql_localhost_opennebula.sql
Use 'onedb restore' or restore the DB using the mysql command:
mysql -u user -h server -P port db_name < backup_file

Total errors found: 0
```

3.6.8 Resume the Federation

This section applies only to environments working in a Federation.

For the **master zone**: This step is not necessary.

For a **slave zone**: The MySQL replication must be resumed now.

- First, add 3 new tables, `vdc_pool`, `marketplace_pool` and `marketplaceapp_pool` to the replication configuration.

Warning: Do not copy the server-id from this example, each slave should already have a unique ID.

```
# vi /etc/my.cnf
[mysqld]
server-id          = 100
replicate-do-table = opennebula.user_pool
replicate-do-table = opennebula.group_pool
replicate-do-table = opennebula.vdc_pool
replicate-do-table = opennebula.zone_pool
replicate-do-table = opennebula.db_versioning
replicate-do-table = opennebula.acl
replicate-do-table = opennebula.marketplace_pool
replicate-do-table = opennebula.marketplaceapp_pool

# service mysqld restart
```

- Start the **slave MySQL** process and check its status. It may take a while to copy and apply all the pending commands.

```
mysql> START SLAVE;
mysql> SHOW SLAVE STATUS\G
```

The `SHOW SLAVE STATUS` output will provide detailed information, but to confirm that the slave is connected to the master MySQL, take a look at these columns:

```
Slave_IO_State: Waiting for master to send event
Slave_IO_Running: Yes
Slave_SQL_Running: Yes
```

3.6.9 Reload Start Scripts in CentOS 7

In order for the system to re-read the configuration files you should issue the following command after the installation of the new packages:

```
# systemctl daemon-reload
```

3.6.10 Enable Start Scripts in CentOS 7

CentOS 7 packages now come with systemd scripts instead of the old systemV ones. You will need to enable the services again so they are started on system boot. The names of the services are the same as the previous one. For example, to enable `opennebula`, `opennebula-sunstone`, `opennebula-flow` and `opennebula-gate` you can issue these commands:

```
# systemctl enable opennebula
# systemctl enable opennebula-sunstone
# systemctl enable opennebula-flow
# systemctl enable opennebula-gate
```

3.6.11 Update the Drivers

You should be able now to start OpenNebula as usual, running 'one start' as `oneadmin`. At this point, execute `onehost sync` to update the new drivers in the hosts.

Warning: Doing `onehost sync` is important. If the monitorization drivers are not updated, the hosts will behave erratically.

3.6.12 Default Auth

If you are using LDAP as default auth driver, you will need to update `/etc/one/oned.conf` and set the new `DEFAULT_AUTH` variable:

```
DEFAULT_AUTH = "ldap"
```

3.6.13 vCenter Password

Note: This step only applies if you are upgrading from OpenNebula **4.10.0**. If you are already using 4.10.1 or 4.10.2 you can skip this step.

If you already have a host with vCenter drivers you need to update the password as version >4.10.0 expects it to be encrypted. To do so, proceed to Sunstone -> Infrastructure -> Hosts, click on the vCenter host(s) and change the value in `VCENTER_PASSWORD` field. It will be automatically encrypted.

3.6.14 Create the Security Group ACL Rule

There is a new kind of resource introduced in 4.12: Security Groups. If you want your existing users to be able to create their own Security Groups, create the following ACL Rule:

```
$ oneacl create "*" SECGROUP/* CREATE *
```

Note: For environments in a Federation: This command needs to be executed only once in the master zone, after it is upgraded to 5.4.

3.6.15 Create the Virtual Router ACL Rule

There is a new kind of resource introduced in 5.0: Virtual Routers. If you want your existing users to be able to create their own Virtual Routers, create the following ACL Rule:

```
$ oneacl create "*" VROUTER/* CREATE *
```

Note: For environments in a Federation: This command needs to be executed only once in the master zone, after it is upgraded to 5.4.

3.6.16 Testing

OpenNebula will continue the monitoring and management of your previous Hosts and VMs.

As a measure of caution, look for any error messages in `oned.log`, and check that all drivers are loaded successfully. After that, keep an eye on `oned.log` while you issue the `onevm`, `onevnet`, `oneimage`, `oneuser`, `onehost` **list** commands. Try also using the **show** subcommand for some resources.

3.6.17 Restoring the Previous Version

If for any reason you need to restore your previous OpenNebula, follow these steps:

- With OpenNebula 5.4 still installed, restore the DB backup using `onedb restore -f`
- Uninstall OpenNebula 5.4, and install again your previous version.
- Copy back the backup of `/etc/one` you did to restore your configuration.

3.6.18 Known Issues

If the MySQL database password contains special characters, such as `@` or `#`, the `onedb` command will fail to connect to it.

The workaround is to temporarily change the `oneadmin`'s password to an ASCII string. The `set password` statement can be used for this:


```
$ mysql -u oneadmin -p
mysql> SET PASSWORD = PASSWORD('newpass');
```

3.7 Upgrading from OpenNebula 4.8.x

This section describes the installation procedure for systems that are already running a 4.8.x OpenNebula. The upgrade to OpenNebula 5.4 can be done directly following this section, you don't need to perform intermediate version upgrades. The upgrade will preserve all current users, hosts, resources and configurations; for both Sqlite and MySQL backends.

Read the Compatibility Guide for [4.10](#), [4.12](#), [4.14](#), [5.0](#) and [5.4](#), and the [Release Notes](#) to know what is new in OpenNebula 5.4.

3.7.1 Upgrading a Federation

If you have two or more 4.8 OpenNebulas working as a Federation, you need to upgrade all of them. The upgrade does not have to be simultaneous, the slaves can be kept running while the master is upgraded.

The steps to follow are:

1. Stop the MySQL replication in all the slaves
2. Upgrade the **master** OpenNebula
3. Upgrade each **slave**
4. Resume the replication

During the time between steps 1 and 4 the slave OpenNebulas can be running, and users can keep accessing them if each zone has a local Sunstone instance. There is however an important limitation to note: all the shared database tables will not be updated in the slaves zones. This means that new user accounts, password changes, new ACL rules, etc. will not have any effect in the slaves. Read the federation architecture documentation for more details.

It is recommended to upgrade all the slave zones as soon as possible.

To perform the first step, [pause the replication](#) in each **slave MySQL**:

```
mysql> STOP SLAVE;

mysql> SHOW SLAVE STATUS\G

Slave_IO_Running: No
Slave_SQL_Running: No
```

Then follow this section for the **master zone**. After the master has been updated to 5.4, upgrade each **slave zone** following this same section.

3.7.2 Preparation

Before proceeding, make sure you don't have any VMs in a transient state (prolog, migr, epil, save). Wait until these VMs get to a final state (runn, suspended, stopped, done). Check the Managing Virtual Machines guide for more information on the VM life-cycle.

Warning: In 4.14 the FAILED state disappears. You need to delete all the VMs in this state **before** the new version is installed.

The network drivers since OpenNebula 5.0 are located in the Virtual Network, rather than in the host. The upgrade process may ask you questions about your existing VMs, Virtual Networks and hosts, and as such it is wise to have the following information saved beforehand, since in the upgrade process OpenNebula will be stopped.

```
$ onevnet list -x > networks.txt
$ onehost list -x > hosts.txt
$ onevm list -x > vms.txt
```

The list of valid network drivers since 5.0 Wizard are:

- 802.1Q
- dummy
- ebttables
- fw
- ovswitch
- vxlan

Stop OpenNebula and any other related services you may have running: EC2, OCCI, and Sunstone. As `oneadmin`, in the front-end:

```
$ sunstone-server stop
$ oneflow-server stop
$ econe-server stop
$ one stop
```

3.7.3 Backup

Backup the configuration files located in `/etc/one`. You don't need to do a manual backup of your database, the `onedb` command will perform one automatically.

```
# cp -r /etc/one /etc/one.YYYY-MM-DD
```

Note: Substitute YYYY-MM-DD with the date.

3.7.4 Installation

Follow the *Platform Notes* and the Installation guide, taking into account that you will already have configured the passwordless ssh access for `oneadmin`.

Make sure to run the `install_gems` tool, as the new OpenNebula version may have different gem requirements.

It is highly recommended **not to keep** your current `oned.conf`, and update the `oned.conf` file shipped with OpenNebula 5.4 to your setup. If for any reason you plan to preserve your current `oned.conf` file, read the *Compatibility Guide* and the complete `oned.conf` reference for 4.8 and 5.0 versions.

3.7.5 Configuration Files Upgrade

If you haven't modified any configuration files, the package managers will replace the configuration files with their newer versions and no manual intervention is required.

If you have customized **any** configuration files under `/etc/one` we recommend you to follow these steps regardless of the platform/linux distribution.

1. Backup `/etc/one` (already performed)
2. Install the new packages (already performed)
3. Compare the old and new configuration files: `diff -ur /etc/one.YYYY-MM-DD /etc/one`. Or you can use graphical diff-tools like `meld` to compare both directories, which are very useful in this step.
4. Edit the **new** files and port all the customizations from the previous version.
5. You should **never** overwrite the configuration files with older versions.

3.7.6 Database Upgrade

The database schema and contents are incompatible between versions. The OpenNebula daemon checks the existing DB version, and will fail to start if the version found is not the one expected, with the message 'Database version mismatch'.

You can upgrade the existing DB with the 'onedb' command. You can specify any Sqlite or MySQL database. Check the onedb reference for more information.

Warning: Make sure at this point that OpenNebula is not running. If you installed from packages, the service may have been started automatically.

Warning: For environments in a Federation: Before upgrading the **master**, make sure that all the slaves have the MySQL replication paused.

After you install the latest OpenNebula, and fix any possible conflicts in `oned.conf`, you can issue the 'onedb upgrade -v' command. The connection parameters have to be supplied with the command line options, see the onedb manpage for more information. Some examples:

```
$ onedb upgrade -v --sqlite /var/lib/one/one.db
```

```
$ onedb upgrade -v -S localhost -u oneadmin -p oneadmin -d opennebula
```

If everything goes well, you should get an output similar to this one:

```
$ onedb upgrade -v -u oneadmin -d opennebula
MySQL Password:
Version read:
Shared tables 4.4.0 : OpenNebula 4.4.0 daemon bootstrap
Local tables 4.4.0 : OpenNebula 4.4.0 daemon bootstrap

>>> Running migrators for shared tables
> Running migrator /usr/lib/one/ruby/onedb/shared/4.4.0_to_4.4.1.rb
> Done in 0.00s
```

```
> Running migrator /usr/lib/one/ruby/onedb/shared/4.4.1_to_4.5.80.rb
> Done in 0.75s
```

Database migrated from 4.4.0 to 4.5.80 (OpenNebula 4.5.80) by onedb command.

```
>>> Running migrators for local tables
Database already uses version 4.5.80
Total time: 0.77s
```

Note: Make sure you keep the backup file. If you face any issues, the onedb command can restore this backup, but it won't downgrade databases to previous versions.

3.7.7 Check DB Consistency

After the upgrade is completed, you should run the command `onedb fsck`.

First, move the 4.8 backup file created by the upgrade command to a safe place.

```
$ mv /var/lib/one/mysql_localhost_opennebula.sql /path/for/one-backups/
```

Then execute the following command:

```
$ onedb fsck -S localhost -u oneadmin -p oneadmin -d opennebula
MySQL dump stored in /var/lib/one/mysql_localhost_opennebula.sql
Use 'onedb restore' or restore the DB using the mysql command:
mysql -u user -h server -P port db_name < backup_file
```

```
Total errors found: 0
```

3.7.8 Resume the Federation

This section applies only to environments working in a Federation.

For the **master zone**: This step is not necessary.

For a **slave zone**: The MySQL replication must be resumed now.

- First, add 3 new tables, `vdc_pool`, `marketplace_pool` and `marketplaceapp_pool` to the replication configuration.

Warning: Do not copy the server-id from this example, each slave should already have a unique ID.

```
# vi /etc/my.cnf
[mysqld]
server-id                = 100
replicate-do-table      = opennebula.user_pool
replicate-do-table      = opennebula.group_pool
replicate-do-table      = opennebula.vdc_pool
replicate-do-table      = opennebula.zone_pool
replicate-do-table      = opennebula.db_versioning
replicate-do-table      = opennebula.acl
replicate-do-table      = opennebula.marketplace_pool
```

```
replicate-do-table = opennebula.marketplaceapp_pool  
  
# service mysqld restart
```

- Start the **slave MySQL** process and check its status. It may take a while to copy and apply all the pending commands.

```
mysql> START SLAVE;  
mysql> SHOW SLAVE STATUS\G
```

The `SHOW SLAVE STATUS` output will provide detailed information, but to confirm that the slave is connected to the master MySQL, take a look at these columns:

```
Slave_IO_State: Waiting for master to send event  
Slave_IO_Running: Yes  
Slave_SQL_Running: Yes
```

3.7.9 Update the Drivers

You should be able now to start OpenNebula as usual, running 'one start' as oneadmin. At this point, execute `onehost sync` to update the new drivers in the hosts.

Warning: Doing `onehost sync` is important. If the monitorization drivers are not updated, the hosts will behave erratically.

3.7.10 Default Auth

If you are using LDAP as default auth driver, you will need to update `/etc/one/oned.conf` and set the new `DEFAULT_AUTH` variable:

```
DEFAULT_AUTH = "ldap"
```

3.7.11 Create the Security Group ACL Rule

There is a new kind of resource introduced in 4.12: Security Groups. If you want your existing users to be able to create their own Security Groups, create the following ACL Rule:

```
$ oneacl create "* SECGROUP/* CREATE *"
```

Note: For environments in a Federation: This command needs to be executed only once in the master zone, after it is upgraded to 5.4.

3.7.12 Create the Virtual Router ACL Rule

There is a new kind of resource introduced in 5.0: Virtual Routers. If you want your existing users to be able to create their own Virtual Routers, create the following ACL Rule:

```
$ oneacl create "* VROUTER/* CREATE *"
```

Note: For environments in a Federation: This command needs to be executed only once in the master zone, after it is upgraded to 5.4.

3.7.13 Testing

OpenNebula will continue the monitoring and management of your previous Hosts and VMs.

As a measure of caution, look for any error messages in `oned.log`, and check that all drivers are loaded successfully. After that, keep an eye on `oned.log` while you issue the `onevm`, `onevnet`, `oneimage`, `oneuser`, `onehost` **list** commands. Try also using the **show** subcommand for some resources.

3.7.14 Restoring the Previous Version

If for any reason you need to restore your previous OpenNebula, follow these steps:

- With OpenNebula 5.4 still installed, restore the DB backup using ‘`onedb restore -f`’
- Uninstall OpenNebula 5.4, and install again your previous version.
- Copy back the backup of `/etc/one` you did to restore your configuration.

3.7.15 Known Issues

If the MySQL database password contains special characters, such as `@` or `#`, the `onedb` command will fail to connect to it.

The workaround is to temporarily change the `oneadmin`’s password to an ASCII string. The `set password` statement can be used for this:

```
$ mysql -u oneadmin -p
mysql> SET PASSWORD = PASSWORD('newpass');
```

3.8 Upgrading from OpenNebula 4.6.x

This section describes the installation procedure for systems that are already running a 4.6.x OpenNebula. The upgrade to OpenNebula 5.4 can be done directly following this section, you don’t need to perform intermediate version upgrades. The upgrade will preserve all current users, hosts, resources and configurations; for both Sqlite and MySQL backends.

Read the [Compatibility Guide for 4.8, 4.10, 4.12, 4.14, 5.0 and 5.4](#), and the [Release Notes](#) to know what is new in OpenNebula 5.4.

3.8.1 Upgrading a Federation

If you have two or more 4.6 OpenNebulas working as a Federation, you need to upgrade all of them. The upgrade does not have to be simultaneous, the slaves can be kept running while the master is upgraded.

The steps to follow are:

1. Stop the MySQL replication in all the slaves
2. Upgrade the **master** OpenNebula
3. Upgrade each **slave**
4. Resume the replication

During the time between steps 1 and 4 the slave OpenNebulas can be running, and users can keep accessing them if each zone has a local Sunstone instance. There is however an important limitation to note: all the shared database tables will not be updated in the slaves zones. This means that new user accounts, password changes, new ACL rules, etc. will not have any effect in the slaves. Read the federation architecture documentation for more details.

It is recommended to upgrade all the slave zones as soon as possible.

To perform the first step, [pause the replication](#) in each **slave MySQL**:

```
mysql> STOP SLAVE;

mysql> SHOW SLAVE STATUS\G

Slave_IO_Running: No
Slave_SQL_Running: No
```

Then follow this section for the **master zone**. After the master has been updated to 5.4, upgrade each **slave zone** following this same section.

3.8.2 Preparation

Before proceeding, make sure you don't have any VMs in a transient state (prolog, migr, epil, save). Wait until these VMs get to a final state (runn, suspended, stopped, done). Check the Managing Virtual Machines guide for more information on the VM life-cycle.

Warning: In 4.14 the FAILED state disappears. You need to delete all the VMs in this state **before** the new version is installed.

The network drivers since OpenNebula 5.0 are located in the Virtual Network, rather than in the host. The upgrade process may ask you questions about your existing VMs, Virtual Networks and hosts, and as such it is wise to have the following information saved beforehand, since in the upgrade process OpenNebula will be stopped.

```
$ onevnet list -x > networks.txt
$ onehost list -x > hosts.txt
$ onevm list -x > vms.txt
```

The list of valid network drivers since 5.0 Wizard are:

- 802.1Q
- dummy
- ebttables
- fw
- ovswitch
- vxlan

Stop OpenNebula and any other related services you may have running: EC2, OCCI, and Sunstone. As `oneadmin`, in the front-end:

```
$ sunstone-server stop
$ oneflow-server stop
$ econe-server stop
$ occi-server stop
$ one stop
```

3.8.3 Backup

Backup the configuration files located in `/etc/one`. You don't need to do a manual backup of your database, the `onedb` command will perform one automatically.

```
# cp -r /etc/one /etc/one.YYYY-MM-DD
```

Note: Substitute `YYYY-MM-DD` with the date.

3.8.4 Installation

Follow the *Platform Notes* and the Installation guide, taking into account that you will already have configured the passwordless ssh access for `oneadmin`.

Make sure to run the `install_gems` tool, as the new OpenNebula version may have different gem requirements.

It is highly recommended **not to keep** your current `oned.conf`, and update the `oned.conf` file shipped with OpenNebula 5.4 to your setup. If for any reason you plan to preserve your current `oned.conf` file, read the *Compatibility Guide* and the complete `oned.conf` reference for 4.6 and 5.0 versions.

3.8.5 Configuration Files Upgrade

If you haven't modified any configuration files, the package managers will replace the configuration files with their newer versions and no manual intervention is required.

If you have customized **any** configuration files under `/etc/one` we recommend you to follow these steps regardless of the platform/linux distribution.

1. Backup `/etc/one` (already performed)
2. Install the new packages (already performed)
3. Compare the old and new configuration files: `diff -ur /etc/one.YYYY-MM-DD /etc/one`. Or you can use graphical diff-tools like `meld` to compare both directories, which are very useful in this step.
4. Edit the **new** files and port all the customizations from the previous version.
5. You should **never** overwrite the configuration files with older versions.

3.8.6 Database Upgrade

The database schema and contents are incompatible between versions. The OpenNebula daemon checks the existing DB version, and will fail to start if the version found is not the one expected, with the message 'Database version mismatch'.

You can upgrade the existing DB with the ‘onedb’ command. You can specify any Sqlite or MySQL database. Check the onedb reference for more information.

Warning: Make sure at this point that OpenNebula is not running. If you installed from packages, the service may have been started automatically.

Warning: For environments in a Federation: Before upgrading the **master**, make sure that all the slaves have the MySQL replication paused.

Note: If you have a MAC_PREFIX in oned.conf different than the default 02:00, open /usr/lib/one/ruby/onedb/local/4.5.80_to_4.7.80.rb and change the value of the ONEDCONF_MAC_PREFIX constant.

After you install the latest OpenNebula, and fix any possible conflicts in oned.conf, you can issue the ‘onedb upgrade -v’ command. The connection parameters have to be supplied with the command line options, see the onedb manpage for more information. Some examples:

```
$ onedb upgrade -v --sqlite /var/lib/one/one.db
```

```
$ onedb upgrade -v -S localhost -u oneadmin -p oneadmin -d opennebula
```

If everything goes well, you should get an output similar to this one:

```
$ onedb upgrade -v -u oneadmin -d opennebula
MySQL Password:
Version read:
Shared tables 4.4.0 : OpenNebula 4.4.0 daemon bootstrap
Local tables 4.4.0 : OpenNebula 4.4.0 daemon bootstrap

>>> Running migrators for shared tables
  > Running migrator /usr/lib/one/ruby/onedb/shared/4.4.0_to_4.4.1.rb
  > Done in 0.00s

  > Running migrator /usr/lib/one/ruby/onedb/shared/4.4.1_to_4.5.80.rb
  > Done in 0.75s

Database migrated from 4.4.0 to 4.5.80 (OpenNebula 4.5.80) by onedb command.

>>> Running migrators for local tables
Database already uses version 4.5.80
Total time: 0.77s
```

Note: Make sure you keep the backup file. If you face any issues, the onedb command can restore this backup, but it won't downgrade databases to previous versions.

3.8.7 Check DB Consistency

After the upgrade is completed, you should run the command `onedb fsck`.

First, move the 4.6 backup file created by the upgrade command to a safe place.

```
$ mv /var/lib/one/mysql_localhost_opennebula.sql /path/for/one-backups/
```

Then execute the following command:

```
$ onedb fsck -S localhost -u oneadmin -p oneadmin -d opennebula
MySQL dump stored in /var/lib/one/mysql_localhost_opennebula.sql
Use 'onedb restore' or restore the DB using the mysql command:
mysql -u user -h server -P port db_name < backup_file

Total errors found: 0
```

3.8.8 Resume the Federation

This section applies only to environments working in a Federation.

For the **master zone**: This step is not necessary.

For a **slave zone**: The MySQL replication must be resumed now.

- First, add 3 new tables, `vdc_pool`, `marketplace_pool` and `marketplaceapp_pool` to the replication configuration.

Warning: Do not copy the server-id from this example, each slave should already have a unique ID.

```
# vi /etc/my.cnf
[mysqld]
server-id          = 100
replicate-do-table = opennebula.user_pool
replicate-do-table = opennebula.group_pool
replicate-do-table = opennebula.vdc_pool
replicate-do-table = opennebula.zone_pool
replicate-do-table = opennebula.db_versioning
replicate-do-table = opennebula.acl
replicate-do-table = opennebula.marketplace_pool
replicate-do-table = opennebula.marketplaceapp_pool

# service mysqld restart
```

- Start the **slave MySQL** process and check its status. It may take a while to copy and apply all the pending commands.

```
mysql> START SLAVE;
mysql> SHOW SLAVE STATUS\G
```

The `SHOW SLAVE STATUS` output will provide detailed information, but to confirm that the slave is connected to the master MySQL, take a look at these columns:

```
Slave_IO_State: Waiting for master to send event
Slave_IO_Running: Yes
Slave_SQL_Running: Yes
```

3.8.9 Update the Drivers

You should be able now to start OpenNebula as usual, running ‘one start’ as oneadmin. At this point, execute `onehost sync` to update the new drivers in the hosts.

Warning: Doing `onehost sync` is important. If the monitorization drivers are not updated, the hosts will behave erratically.

3.8.10 Create the Security Group ACL Rule

There is a new kind of resource introduced in 4.12: Security Groups. If you want your existing users to be able to create their own Security Groups, create the following ACL Rule:

```
$ oneacl create "* SECGROUP/* CREATE *"
```

Note: For environments in a Federation: This command needs to be executed only once in the master zone, after it is upgraded to 5.4.

3.8.11 Create the Virtual Router ACL Rule

There is a new kind of resource introduced in 5.0: Virtual Routers. If you want your existing users to be able to create their own Virtual Routers, create the following ACL Rule:

```
$ oneacl create "* VROUTER/* CREATE *"
```

Note: For environments in a Federation: This command needs to be executed only once in the master zone, after it is upgraded to 5.4.

3.8.12 Testing

OpenNebula will continue the monitoring and management of your previous Hosts and VMs.

As a measure of caution, look for any error messages in `oned.log`, and check that all drivers are loaded successfully. After that, keep an eye on `oned.log` while you issue the `onevm`, `onevnet`, `oneimage`, `oneuser`, `onehost list` commands. Try also using the **show** subcommand for some resources.

3.8.13 Restoring the Previous Version

If for any reason you need to restore your previous OpenNebula, follow these steps:

- With OpenNebula 5.4 still installed, restore the DB backup using ‘`onedb restore -f`’
- Uninstall OpenNebula 5.4, and install again your previous version.
- Copy back the backup of `/etc/one` you did to restore your configuration.

3.8.14 Known Issues

If the MySQL database password contains special characters, such as @ or #, the onedb command will fail to connect to it.

The workaround is to temporarily change the oneadmin's password to an ASCII string. The `set password` statement can be used for this:

```
$ mysql -u oneadmin -p
mysql> SET PASSWORD = PASSWORD('newpass');
```

3.9 Upgrading from OpenNebula 4.4.x

This section describes the installation procedure for systems that are already running a 4.4.x OpenNebula. The upgrade to OpenNebula 5.4 can be done directly following this section, you don't need to perform intermediate version upgrades. The upgrade will preserve all current users, hosts, resources and configurations; for both Sqlite and MySQL backends.

Read the Compatibility Guide for 4.6, 4.8, 4.10, 4.12, 4.14, 5.0 and 5.4, and the Release Notes to know what is new in OpenNebula 5.4.

3.9.1 Preparation

Before proceeding, make sure you don't have any VMs in a transient state (prolog, migr, epil, save). Wait until these VMs get to a final state (runn, suspended, stopped, done). Check the Managing Virtual Machines guide for more information on the VM life-cycle.

Warning: In 4.14 the FAILED state disappears. You need to delete all the VMs in this state **before** the new version is installed.

The network drivers since OpenNebula 5.0 are located in the Virtual Network, rather than in the host. The upgrade process may ask you questions about your existing VMs, Virtual Networks and hosts, and as such it is wise to have the following information saved beforehand, since in the upgrade process OpenNebula will be stopped.

```
$ onevnet list -x > networks.txt
$ onehost list -x > hosts.txt
$ onevm list -x > vms.txt
```

The list of valid network drivers since 5.0 Wizard are:

- 802.1Q
- dummy
- ebttables
- fw
- ovswitch
- vxlan

Stop OpenNebula and any other related services you may have running: EC2, OCCI, and Sunstone. As oneadmin, in the front-end:

```
$ sunstone-server stop
$ oneflow-server stop
$ econe-server stop
$ occi-server stop
$ one stop
```

3.9.2 Backup

Backup the configuration files located in `/etc/one`. You don't need to do a manual backup of your database, the `onedb` command will perform one automatically.

3.9.3 Installation

Follow the *Platform Notes* and the Installation guide, taking into account that you will already have configured the passwordless ssh access for `oneadmin`.

Make sure to run the `install_gems` tool, as the new OpenNebula version may have different gem requirements.

It is highly recommended **not to keep** your current `oned.conf`, and update the `oned.conf` file shipped with OpenNebula 5.4 to your setup. If for any reason you plan to preserve your current `oned.conf` file, read the *Compatibility Guide* and the complete `oned.conf` reference for 4.4 and 5.0 versions.

3.9.4 Database Upgrade

The database schema and contents are incompatible between versions. The OpenNebula daemon checks the existing DB version, and will fail to start if the version found is not the one expected, with the message 'Database version mismatch'.

You can upgrade the existing DB with the 'onedb' command. You can specify any Sqlite or MySQL database. Check the `onedb` reference for more information.

Warning: Make sure at this point that OpenNebula is not running. If you installed from packages, the service may have been started automatically.

Note: If you have a `MAC_PREFIX` in `oned.conf` different than the default `02:00`, open `/usr/lib/one/ruby/onedb/local/4.5.80_to_4.7.80.rb` and change the value of the `ONEDCONF_MAC_PREFIX` constant.

After you install the latest OpenNebula, and fix any possible conflicts in `oned.conf`, you can issue the 'onedb upgrade -v' command. The connection parameters have to be supplied with the command line options, see the `onedb manpage` for more information. Some examples:

```
$ onedb upgrade -v --sqlite /var/lib/one/one.db
```

```
$ onedb upgrade -v -S localhost -u oneadmin -p oneadmin -d opennebula
```

If everything goes well, you should get an output similar to this one:

```
$ onedb upgrade -v -u oneadmin -d opennebula
MySQL Password:
Version read:
Shared tables 4.4.0 : OpenNebula 4.4.0 daemon bootstrap
Local tables 4.4.0 : OpenNebula 4.4.0 daemon bootstrap

>>> Running migrators for shared tables
  > Running migrator /usr/lib/one/ruby/onedb/shared/4.4.0_to_4.4.1.rb
  > Done in 0.00s

  > Running migrator /usr/lib/one/ruby/onedb/shared/4.4.1_to_4.5.80.rb
  > Done in 0.75s

Database migrated from 4.4.0 to 4.5.80 (OpenNebula 4.5.80) by onedb command.

>>> Running migrators for local tables
Database already uses version 4.5.80
Total time: 0.77s
```

Note: Make sure you keep the backup file. If you face any issues, the onedb command can restore this backup, but it won't downgrade databases to previous versions.

3.9.5 Check DB Consistency

After the upgrade is completed, you should run the command `onedb fsck`.

First, move the 4.4 backup file created by the upgrade command to a safe place.

```
$ mv /var/lib/one/mysql_localhost_opennebula.sql /path/for/one-backups/
```

Then execute the following command:

```
$ onedb fsck -S localhost -u oneadmin -p oneadmin -d opennebula
MySQL dump stored in /var/lib/one/mysql_localhost_opennebula.sql
Use 'onedb restore' or restore the DB using the mysql command:
mysql -u user -h server -P port db_name < backup_file

Total errors found: 0
```

3.9.6 Update the Drivers

You should be able now to start OpenNebula as usual, running 'one start' as oneadmin. At this point, execute `onehost sync` to update the new drivers in the hosts.

Warning: Doing `onehost sync` is important. If the monitorization drivers are not updated, the hosts will behave erratically.

3.9.7 Create the Security Group ACL Rule

There is a new kind of resource introduced in 4.12: Security Groups. If you want your existing users to be able to create their own Security Groups, create the following ACL Rule:

```
$ oneacl create "* SECGROUP/* CREATE *"
```

3.9.8 Create the Virtual Router ACL Rule

There is a new kind of resource introduced in 5.0: Virtual Routers. If you want your existing users to be able to create their own Virtual Routers, create the following ACL Rule:

```
$ oneacl create "* VROUTER/* CREATE *"
```

Note: For environments in a Federation: This command needs to be executed only once in the master zone, after it is upgraded to 5.4.

3.9.9 Testing

OpenNebula will continue the monitoring and management of your previous Hosts and VMs.

As a measure of caution, look for any error messages in `oned.log`, and check that all drivers are loaded successfully. After that, keep an eye on `oned.log` while you issue the `onevm`, `onevnet`, `oneimage`, `oneuser`, `onehost` **list** commands. Try also using the **show** subcommand for some resources.

3.9.10 Restoring the Previous Version

If for any reason you need to restore your previous OpenNebula, follow these steps:

- With OpenNebula 5.4 still installed, restore the DB backup using `'onedb restore -f'`
- Uninstall OpenNebula 5.4, and install again your previous version.
- Copy back the backup of `/etc/one` you did to restore your configuration.

3.9.11 Known Issues

If the MySQL database password contains special characters, such as `@` or `#`, the `onedb` command will fail to connect to it.

The workaround is to temporarily change the `oneadmin`'s password to an ASCII string. The `set password` statement can be used for this:

```
$ mysql -u oneadmin -p
mysql> SET PASSWORD = PASSWORD('newpass');
```

3.10 Upgrading from OpenNebula 4.2

This section describes the installation procedure for systems that are already running a 4.2 OpenNebula. The upgrade to OpenNebula 5.4 can be done directly following this section, you don't need to perform intermediate version upgrades. The upgrade will preserve all current users, hosts, resources and configurations; for both Sqlite and MySQL backends.

Read the [Compatibility Guide](#) for 4.4, 4.6, 4.8, 4.10, 4.12, 4.14, 5.0 and 5.4, and the [Release Notes](#) to know what is new in OpenNebula 5.4.

Warning: With the new multi-system DS functionality, it is now required that the system DS is also part of the cluster. If you are using System DS 0 for Hosts inside a Cluster, any VM saved (stop, suspend, undeploy) **will not be able to be resumed after the upgrade process.**

Warning: Two drivers available in 4.0 are now discontinued: **ganglia** and **iscsi**.

- **iscsi** drivers have been moved out of the main OpenNebula distribution and are available (although not supported) as an [addon](#).
- **ganglia** drivers have been moved out of the main OpenNebula distribution and are available (although not supported) as an [addon](#).

3.10.1 Preparation

Before proceeding, make sure you don't have any VMs in a transient state (prolog, migr, epil, save). Wait until these VMs get to a final state (runn, suspended, stopped, done). Check the [Managing Virtual Machines](#) guide for more information on the VM life-cycle.

Warning: In 4.14 the `FAILED` state disappears. You need to delete all the VMs in this state **before** the new version is installed.

The network drivers since OpenNebula 5.0 are located in the Virtual Network, rather than in the host. The upgrade process may ask you questions about your existing VMs, Virtual Networks and hosts, and as such it is wise to have the following information saved beforehand, since in the upgrade process OpenNebula will be stopped.

```
$ onevnet list -x > networks.txt
$ onehost list -x > hosts.txt
$ onevm list -x > vms.txt
```

The list of valid network drivers since 5.0 Wizard are:

- 802.1Q
- dummy
- ebttables
- fw
- ovswitch
- vxlan

Stop OpenNebula and any other related services you may have running: EC2, OCCI, and Sunstone. As `oneadmin`, in the front-end:

```
$ sunstone-server stop
$ oneflow-server stop
$ econe-server stop
$ occi-server stop
$ one stop
```

3.10.2 Backup

Backup the configuration files located in `/etc/one`. You don't need to do a manual backup of your database, the `onedb` command will perform one automatically.

3.10.3 Installation

Follow the *Platform Notes* and the Installation guide, taking into account that you will already have configured the passwordless ssh access for `oneadmin`.

Make sure to run the `install_gems` tool, as the new OpenNebula version may have different gem requirements.

It is highly recommended **not to keep** your current `oned.conf`, and update the `oned.conf` file shipped with OpenNebula 5.4 to your setup. If for any reason you plan to preserve your current `oned.conf` file, read the *Compatibility Guide* and the complete `oned.conf` reference for 4.2 and 5.0 versions.

3.10.4 Database Upgrade

The database schema and contents are incompatible between versions. The OpenNebula daemon checks the existing DB version, and will fail to start if the version found is not the one expected, with the message 'Database version mismatch'.

You can upgrade the existing DB with the 'onedb' command. You can specify any Sqlite or MySQL database. Check the `onedb` reference for more information.

Warning: Make sure at this point that OpenNebula is not running. If you installed from packages, the service may have been started automatically.

Note: If you have a `MAC_PREFIX` in `oned.conf` different than the default `02:00`, open `/usr/lib/one/ruby/onedb/local/4.5.80_to_4.7.80.rb` and change the value of the `ONEDCONF_MAC_PREFIX` constant.

After you install the latest OpenNebula, and fix any possible conflicts in `oned.conf`, you can issue the 'onedb upgrade -v' command. The connection parameters have to be supplied with the command line options, see the `onedb` manpage for more information. Some examples:

```
$ onedb upgrade -v --sqlite /var/lib/one/one.db
```

```
$ onedb upgrade -v -S localhost -u oneadmin -p oneadmin -d opennebula
```

If everything goes well, you should get an output similar to this one:

```
$ onedb upgrade -v -u oneadmin -d opennebula
MySQL Password:
Version read:
Shared tables 4.4.0 : OpenNebula 4.4.0 daemon bootstrap
Local tables 4.4.0 : OpenNebula 4.4.0 daemon bootstrap

>>> Running migrators for shared tables
  > Running migrator /usr/lib/one/ruby/onedb/shared/4.4.0_to_4.4.1.rb
  > Done in 0.00s

  > Running migrator /usr/lib/one/ruby/onedb/shared/4.4.1_to_4.5.80.rb
  > Done in 0.75s

Database migrated from 4.4.0 to 4.5.80 (OpenNebula 4.5.80) by onedb command.

>>> Running migrators for local tables
Database already uses version 4.5.80
Total time: 0.77s
```

If you receive the message “ATTENTION: manual intervention required”, read the section *Manual Intervention Required* below.

Note: Make sure you keep the backup file. If you face any issues, the onedb command can restore this backup, but it won't downgrade databases to previous versions.

3.10.5 Check DB Consistency

After the upgrade is completed, you should run the command `onedb fsck`.

First, move the 4.2 backup file created by the upgrade command to a safe place.

```
$ mv /var/lib/one/mysql_localhost_opennebula.sql /path/for/one-backups/
```

Then execute the following command:

```
$ onedb fsck -S localhost -u oneadmin -p oneadmin -d opennebula
MySQL dump stored in /var/lib/one/mysql_localhost_opennebula.sql
Use 'onedb restore' or restore the DB using the mysql command:
mysql -u user -h server -P port db_name < backup_file

Total errors found: 0
```

3.10.6 Update the Drivers

You should be able now to start OpenNebula as usual, running ‘one start’ as oneadmin. At this point, execute `onehost sync` to update the new drivers in the hosts.

Warning: Doing `onehost sync` is important. If the monitorization drivers are not updated, the hosts will behave erratically.

3.10.7 Create the Security Group ACL Rule

There is a new kind of resource introduced in 4.12: Security Groups. If you want your existing users to be able to create their own Security Groups, create the following ACL Rule:

```
$ oneacl create "* SECGROUP/* CREATE *"
```

3.10.8 Create the Virtual Router ACL Rule

There is a new kind of resource introduced in 5.0: Virtual Routers. If you want your existing users to be able to create their own Virtual Routers, create the following ACL Rule:

```
$ oneacl create "* VROUTER/* CREATE *"
```

Note: For environments in a Federation: This command needs to be executed only once in the master zone, after it is upgraded to 5.4.

3.10.9 Setting new System DS

With the new multi-system DS functionality, it is now required that the system DS is also part of the cluster. If you are using System DS 0 for Hosts inside a Cluster, any VM saved (stop, suspend, undeploy) **will not be able to be resumed after the upgrade process**.

You will need to have at least one system DS in each cluster. If you don't already, create new system DS with the same definition as the system DS 0 (TM_MAD driver). Depending on your setup this may or may not require additional configuration on the hosts.

You may also try to recover saved VMs (stop, suspend, undeploy) following the steps described in this [thread of the users mailing list](#).

3.10.10 Testing

OpenNebula will continue the monitoring and management of your previous Hosts and VMs.

As a measure of caution, look for any error messages in oned.log, and check that all drivers are loaded successfully. After that, keep an eye on oned.log while you issue the onevm, onevnet, oneimage, oneuser, onehost **list** commands. Try also using the **show** subcommand for some resources.

3.10.11 Restoring the Previous Version

If for any reason you need to restore your previous OpenNebula, follow these steps:

- With OpenNebula 5.4 still installed, restore the DB backup using 'onedb restore -f'
- Uninstall OpenNebula 5.4, and install again your previous version.
- Copy back the backup of /etc/one you did to restore your configuration.

3.10.12 Known Issues

If the MySQL database password contains special characters, such as @ or #, the onedb command will fail to connect to it.

The workaround is to temporarily change the oneadmin's password to an ASCII string. The `set password` statement can be used for this:

```
$ mysql -u oneadmin -p
mysql> SET PASSWORD = PASSWORD('newpass');
```

3.10.13 Manual Intervention Required

If you have a datastore configured to use a tm driver not included in the OpenNebula distribution, the onedb upgrade command will show you this message:

```
ATTENTION: manual intervention required

The Datastore <id> <name> is using the
custom TM MAD '<tm_mad>'. You will need to define new
configuration parameters in oned.conf for this driver, see
http://opennebula.org/documentation:rel4.4:upgrade
```

Since OpenNebula 4.4, each `tm_mad` driver has a `TM_MAD_CONF` section in `oned.conf`. If you developed the driver, it should be fairly easy to define the required information looking at the existing ones:

```
# The configuration for each driver is defined in TM_MAD_CONF. These
# values are used when creating a new datastore and should not be modified
# since they define the datastore behaviour.
# name      : name of the transfer driver, listed in the -d option of the
#             TM_MAD section
# ln_target : determines how the persistent images will be cloned when
#             a new VM is instantiated.
#     NONE: The image will be linked and no more storage capacity will be used
#     SELF: The image will be cloned in the Images datastore
#     SYSTEM: The image will be cloned in the System datastore
# clone_target : determines how the non persistent images will be
#             cloned when a new VM is instantiated.
#     NONE: The image will be linked and no more storage capacity will be used
#     SELF: The image will be cloned in the Images datastore
#     SYSTEM: The image will be cloned in the System datastore
# shared : determines if the storage holding the system datastore is shared
#             among the different hosts or not. Valid values: "yes" or "no"

TM_MAD_CONF = [
  name      = "lvm",
  ln_target = "NONE",
  clone_target= "SELF",
  shared    = "yes"
]
```

3.11 Upgrading from OpenNebula 4.0.x

This section describes the installation procedure for systems that are already running a 4.0.x OpenNebula. The upgrade to OpenNebula 5.4 can be done directly following this section, you don't need to perform intermediate version upgrades. The upgrade will preserve all current users, hosts, resources and configurations; for both Sqlite and MySQL backends.

Read the Compatibility Guide for 4.2, 4.4, 4.6, 4.8, 4.10, 4.12, 4.14, 5.0 and 5.4, and the Release Notes to know what is new in OpenNebula 5.4.

Warning: With the new multi-system DS functionality, it is now required that the system DS is also part of the cluster. If you are using System DS 0 for Hosts inside a Cluster, any VM saved (stop, suspend, undeploy) **will not be able to be resumed after the upgrade process.**

Warning: Two drivers available in 4.0 are now discontinued: **ganglia** and **iscsi**.

- **iscsi** drivers have been moved out of the main OpenNebula distribution and are available (although not supported) as an [addon](#).
- **ganglia** drivers have been moved out of the main OpenNebula distribution and are available (although not supported) as an [addon](#).

Warning: There are combinations of **VMware storage** no longer supported (see the VMFS Datastore guide for the supported configurations).

If you want to upgrade and you are using SSH, NFS or VMFS without SSH-mode, you will need to manually migrate your images to a newly created VMFS with SSH-mode datastore. To do so implies powering off all the VMs with images in any of the deprecated datastores, upgrade OpenNebula, create a VMFS datastore and then manually register the images from those deprecated datastores into the new one. [Let us know](#) if you have doubts or problems with this process.

3.11.1 Preparation

Before proceeding, make sure you don't have any VMs in a transient state (prolog, migr, epil, save). Wait until these VMs get to a final state (runn, suspended, stopped, done). Check the Managing Virtual Machines guide for more information on the VM life-cycle.

Warning: In 4.14 the FAILED state disappears. You need to delete all the VMs in this state **before** the new version is installed.

The network drivers since OpenNebula 5.0 are located in the Virtual Network, rather than in the host. The upgrade process may ask you questions about your existing VMs, Virtual Networks and hosts, and as such it is wise to have the following information saved beforehand, since in the upgrade process OpenNebula will be stopped.

```
$ onevnet list -x > networks.txt
$ onehost list -x > hosts.txt
$ onevm list -x > vms.txt
```

The list of valid network drivers since 5.0 Wizard are:

- 802.1Q
- dummy
- ebttables
- fw
- ovswitch
- vxlan

Stop OpenNebula and any other related services you may have running: EC2, OCCl, and Sunstone. As `oneadmin`, in the front-end:

```
$ sunstone-server stop
$ oneflow-server stop
$ econe-server stop
$ occi-server stop
$ one stop
```

3.11.2 Backup

Backup the configuration files located in `/etc/one`. You don't need to do a manual backup of your database, the `onedb` command will perform one automatically.

3.11.3 Installation

Follow the *Platform Notes* and the Installation guide, taking into account that you will already have configured the passwordless ssh access for `oneadmin`.

Make sure to run the `install_gems` tool, as the new OpenNebula version may have different gem requirements.

It is highly recommended **not to keep** your current `oned.conf`, and update the `oned.conf` file shipped with OpenNebula 5.4 to your setup. If for any reason you plan to preserve your current `oned.conf` file, read the *Compatibility Guide* and the complete `oned.conf` reference for 4.0 and 5.0 versions.

3.11.4 Database Upgrade

The database schema and contents are incompatible between versions. The OpenNebula daemon checks the existing DB version, and will fail to start if the version found is not the one expected, with the message 'Database version mismatch'.

You can upgrade the existing DB with the 'onedb' command. You can specify any Sqlite or MySQL database. Check the `onedb` reference for more information.

Warning: Make sure at this point that OpenNebula is not running. If you installed from packages, the service may have been started automatically.

Note: If you have a `MAC_PREFIX` in `oned.conf` different than the default `02:00`, open `/usr/lib/one/ruby/onedb/local/4.5.80_to_4.7.80.rb` and change the value of the `ONEDCONF_MAC_PREFIX` constant.

After you install the latest OpenNebula, and fix any possible conflicts in `oned.conf`, you can issue the `'onedb upgrade -v'` command. The connection parameters have to be supplied with the command line options, see the `onedb` manpage for more information. Some examples:

```
$ onedb upgrade -v --sqlite /var/lib/one/one.db
```

```
$ onedb upgrade -v -S localhost -u oneadmin -p oneadmin -d opennebula
```

If everything goes well, you should get an output similar to this one:

```
$ onedb upgrade -v -u oneadmin -d opennebula
MySQL Password:
Version read:
Shared tables 4.4.0 : OpenNebula 4.4.0 daemon bootstrap
Local tables  4.4.0 : OpenNebula 4.4.0 daemon bootstrap

>>> Running migrators for shared tables
  > Running migrator /usr/lib/one/ruby/onedb/shared/4.4.0_to_4.4.1.rb
  > Done in 0.00s

  > Running migrator /usr/lib/one/ruby/onedb/shared/4.4.1_to_4.5.80.rb
  > Done in 0.75s

Database migrated from 4.4.0 to 4.5.80 (OpenNebula 4.5.80) by onedb command.

>>> Running migrators for local tables
Database already uses version 4.5.80
Total time: 0.77s
```

If you receive the message “ATTENTION: manual intervention required”, read the section *Manual Intervention Required* below.

Note: Make sure you keep the backup file. If you face any issues, the `onedb` command can restore this backup, but it won't downgrade databases to previous versions.

3.11.5 Check DB Consistency

After the upgrade is completed, you should run the command `onedb fsck`.

First, move the 4.0 backup file created by the upgrade command to a safe place.

```
$ mv /var/lib/one/mysql_localhost_opennebula.sql /path/for/one-backups/
```

Then execute the following command:

```
$ onedb fsck -S localhost -u oneadmin -p oneadmin -d opennebula
MySQL dump stored in /var/lib/one/mysql_localhost_opennebula.sql
Use 'onedb restore' or restore the DB using the mysql command:
mysql -u user -h server -P port db_name < backup_file

Total errors found: 0
```

3.11.6 Update the Drivers

You should be able now to start OpenNebula as usual, running ‘one start’ as oneadmin. At this point, execute `onehost sync` to update the new drivers in the hosts.

Warning: Doing `onehost sync` is important. If the monitorization drivers are not updated, the hosts will behave erratically.

3.11.7 Setting new System DS

With the new multi-system DS functionality, it is now required that the system DS is also part of the cluster. If you are using System DS 0 for Hosts inside a Cluster, any VM saved (stop, suspend, undeploy) **will not be able to be resumed after the upgrade process**.

You will need to have at least one system DS in each cluster. If you don't already, create new system DS with the same definition as the system DS 0 (TM_MAD driver). Depending on your setup this may or may not require additional configuration on the hosts.

You may also try to recover saved VMs (stop, suspend, undeploy) following the steps described in this [thread of the users mailing list](#).

3.11.8 Create the Security Group ACL Rule

There is a new kind of resource introduced in 4.12: Security Groups. If you want your existing users to be able to create their own Security Groups, create the following ACL Rule:

```
$ oneacl create "*" SECGROUP/* CREATE *
```

3.11.9 Create the Virtual Router ACL Rule

There is a new kind of resource introduced in 5.0: Virtual Routers. If you want your existing users to be able to create their own Virtual Routers, create the following ACL Rule:

```
$ oneacl create "*" VROUTER/* CREATE *
```

Note: For environments in a Federation: This command needs to be executed only once in the master zone, after it is upgraded to 5.4.

3.11.10 Testing

OpenNebula will continue the monitoring and management of your previous Hosts and VMs.

As a measure of caution, look for any error messages in `oned.log`, and check that all drivers are loaded successfully. After that, keep an eye on `oned.log` while you issue the `onevm`, `onevnet`, `oneimage`, `oneuser`, `onehost list` commands. Try also using the **show** subcommand for some resources.

3.11.11 Restoring the Previous Version

If for any reason you need to restore your previous OpenNebula, follow these steps:

- With OpenNebula 5.4 still installed, restore the DB backup using 'onedb restore -f'
- Uninstall OpenNebula 5.4, and install again your previous version.
- Copy back the backup of /etc/one you did to restore your configuration.

3.11.12 Known Issues

If the MySQL database password contains special characters, such as @ or #, the onedb command will fail to connect to it.

The workaround is to temporarily change the oneadmin's password to an ASCII string. The `set password` statement can be used for this:

```
$ mysql -u oneadmin -p
mysql> SET PASSWORD = PASSWORD('newpass');
```

3.11.13 Manual Intervention Required

If you have a datastore configured to use a tm driver not included in the OpenNebula distribution, the onedb upgrade command will show you this message:

```
ATTENTION: manual intervention required

The Datastore <id> <name> is using the
custom TM MAD '<tm_mad>'. You will need to define new
configuration parameters in oned.conf for this driver, see
http://opennebula.org/documentation:rel4.4:upgrade
```

Since OpenNebula 4.4, each tm_mad driver has a TM_MAD_CONF section in oned.conf. If you developed the driver, it should be fairly easy to define the required information looking at the existing ones:

```
# The configuration for each driver is defined in TM_MAD_CONF. These
# values are used when creating a new datastore and should not be modified
# since they define the datastore behaviour.
# name      : name of the transfer driver, listed in the -d option of the
#             TM_MAD section
# ln_target : determines how the persistent images will be cloned when
#             a new VM is instantiated.
#             NONE: The image will be linked and no more storage capacity will be used
#             SELF: The image will be cloned in the Images datastore
#             SYSTEM: The image will be cloned in the System datastore
# clone_target : determines how the non persistent images will be
#                 cloned when a new VM is instantiated.
#             NONE: The image will be linked and no more storage capacity will be used
#             SELF: The image will be cloned in the Images datastore
#             SYSTEM: The image will be cloned in the System datastore
# shared    : determines if the storage holding the system datastore is shared
#                 among the different hosts or not. Valid values: "yes" or "no"

TM_MAD_CONF = [
```

```

name          = "lvm",
ln_target     = "NONE",
clone_target  = "SELF",
shared       = "yes"
]

```

3.12 Upgrading from OpenNebula 3.8.x

This section describes the installation procedure for systems that are already running a 3.8.x OpenNebula. The upgrade to OpenNebula 5.4 can be done directly following this section, you don't need to perform intermediate version upgrades. The upgrade will preserve all current users, hosts, resources and configurations; for both Sqlite and MySQL backends.

Read the Compatibility Guide for 4.0, 4.2, 4.4, 4.6, 4.8, 4.10, 4.12, 4.14, 5.0 and 5.4, and the Release Notes to know what is new in OpenNebula 5.4.

Warning: With the new multi-system DS functionality, it is now required that the system DS is also part of the cluster. If you are using System DS 0 for Hosts inside a Cluster, any VM saved (stop, suspend, undeploy) **will not be able to be resumed after the upgrade process**.

Warning: Two drivers available in 3.8 are now discontinued: **ganglia** and **iscsi**.

- **iscsi** drivers have been moved out of the main OpenNebula distribution and are available (although not supported) as an [addon](#).
- **ganglia** drivers have been moved out of the main OpenNebula distribution and are available (although not supported) as an [addon](#).

Warning: There are combinations of **VMware storage** no longer supported (see the VMFS Datastore guide for the supported configurations).

If you want to upgrade and you are using SSH, NFS or VMFS without SSH-mode, you will need to manually migrate your images to a newly created VMFS with SSH-mode datastore. To do so implies powering off all the VMs with images in any of the deprecated datastores, upgrade OpenNebula, create a VMFS datastore and then manually register the images from those deprecated datastores into the new one. [Let us know](#) if you have doubts or problems with this process.

3.12.1 Preparation

Before proceeding, make sure you don't have any VMs in a transient state (prolog, migr, epil, save). Wait until these VMs get to a final state (runn, suspended, stopped, done). Check the Managing Virtual Machines guide for more information on the VM life-cycle.

Warning: In 4.14 the FAILED state disappears. You need to delete all the VMs in this state **before** the new version is installed.

The network drivers since OpenNebula 5.0 are located in the Virtual Network, rather than in the host. The upgrade process may ask you questions about your existing VMs, Virtual Networks and hosts, and as such it is wise to have the following information saved beforehand, since in the upgrade process OpenNebula will be stopped.

```
$ onevnet list -x > networks.txt
$ onehost list -x > hosts.txt
$ onevm list -x > vms.txt
```

The list of valid network drivers since 5.0 Wizard are:

- 802.1Q
- dummy
- ebttables
- fw
- ovswitch
- vxlan

Stop OpenNebula and any other related services you may have running: EC2, OCCI, and Sunstone. As `oneadmin`, in the front-end:

```
$ sunstone-server stop
$ oneflow-server stop
$ econe-server stop
$ occi-server stop
$ one stop
```

3.12.2 Backup

Backup the configuration files located in `/etc/one`. You don't need to do a manual backup of your database, the `onedb` command will perform one automatically.

3.12.3 Installation

Follow the *Platform Notes* and the Installation guide, taking into account that you will already have configured the passwordless ssh access for `oneadmin`.

Make sure to run the `install_gems` tool, as the new OpenNebula version may have different gem requirements.

It is highly recommended **not to keep** your current `oned.conf`, and update the `oned.conf` file shipped with OpenNebula 5.4 to your setup. If for any reason you plan to preserve your current `oned.conf` file, read the *Compatibility Guide* and the complete `oned.conf` reference for 3.8 and 5.0 versions.

3.12.4 Database Upgrade

The database schema and contents are incompatible between versions. The OpenNebula daemon checks the existing DB version, and will fail to start if the version found is not the one expected, with the message 'Database version mismatch'.

You can upgrade the existing DB with the 'onedb' command. You can specify any Sqlite or MySQL database. Check the `onedb` reference for more information.

Warning: Make sure at this point that OpenNebula is not running. If you installed from packages, the service may have been started automatically.

Note: If you have a `MAC_PREFIX` in `oned.conf` different than the default `02:00`, `open /usr/lib/one/ruby/onedb/local/4.5.80_to_4.7.80.rb` and change the value of the `ONEDCONF_MAC_PREFIX` constant.

After you install the latest OpenNebula, and fix any possible conflicts in `oned.conf`, you can issue the `'onedb upgrade -v'` command. The connection parameters have to be supplied with the command line options, see the `onedb manpage` for more information. Some examples:

```
$ onedb upgrade -v --sqlite /var/lib/one/one.db
```

```
$ onedb upgrade -v -S localhost -u oneadmin -p oneadmin -d opennebula
```

If everything goes well, you should get an output similar to this one:

```
$ onedb upgrade -v -u oneadmin -d opennebula
MySQL Password:
Version read:
Shared tables 3.8.0 : OpenNebula 3.8.0 daemon bootstrap
Local tables  3.8.0 : OpenNebula 3.8.0 daemon bootstrap

MySQL dump stored in /var/lib/one/mysql_localhost_opennebula.sql
Use 'onedb restore' or restore the DB using the mysql command:
mysql -u user -h server -P port db_name < backup_file

>>> Running migrators for shared tables
> Running migrator /usr/lib/one/ruby/onedb/shared/3.8.0_to_3.8.1.rb
> Done in 0.36s

> Running migrator /usr/lib/one/ruby/onedb/shared/3.8.1_to_3.8.2.rb
> Done in 0.00s

> Running migrator /usr/lib/one/ruby/onedb/shared/3.8.2_to_3.8.3.rb
> Done in 0.00s

> Running migrator /usr/lib/one/ruby/onedb/shared/3.8.3_to_3.8.4.rb
> Done in 0.56s

> Running migrator /usr/lib/one/ruby/onedb/shared/3.8.4_to_3.8.5.rb
> Done in 0.00s

> Running migrator /usr/lib/one/ruby/onedb/shared/3.8.5_to_3.9.80.rb

ATTENTION: manual intervention required
Virtual Machine deployment files have been moved from /var/lib/one to
/var/lib/one/vms. You need to move these files manually:

    $ mv /var/lib/one/[0-9]* /var/lib/one/vms

> Done in 1.10s

> Running migrator /usr/lib/one/ruby/onedb/shared/3.9.80_to_3.9.90.rb
```

ATTENTION: manual intervention required

IM and VM MADS have been renamed in oned.conf. To keep your existing hosts working, you need to duplicate the drivers with the old names.

For example, for kvm you will have IM_MAD "kvm" and VM_MAD "kvm", so you need to add IM_MAD "im_kvm" and VM_MAD "vmm_kvm"

```
IM_MAD = [
  name       = "kvm",
  executable = "one_im_ssh",
  arguments  = "-r 0 -t 15 kvm" ]
```

```
IM_MAD = [
  name       = "im_kvm",
  executable = "one_im_ssh",
  arguments  = "-r 0 -t 15 kvm" ]
```

```
VM_MAD = [
  name       = "kvm",
  executable = "one_vmm_exec",
  arguments  = "-t 15 -r 0 kvm",
  default    = "vmm_exec/vmm_exec_kvm.conf",
  type       = "kvm" ]
```

```
VM_MAD = [
  name       = "vmm_kvm",
  executable = "one_vmm_exec",
  arguments  = "-t 15 -r 0 kvm",
  default    = "vmm_exec/vmm_exec_kvm.conf",
  type       = "kvm" ]
```

> Done in 0.41s

> Running migrator /usr/lib/one/ruby/onedb/shared/3.9.90_to_4.0.0.rb
> Done in 0.00s

> Running migrator /usr/lib/one/ruby/onedb/shared/4.0.0_to_4.0.1.rb
> Done in 0.00s

> Running migrator /usr/lib/one/ruby/onedb/shared/4.0.1_to_4.1.80.rb
> Done in 0.09s

> Running migrator /usr/lib/one/ruby/onedb/shared/4.1.80_to_4.2.0.rb
> Done in 0.00s

> Running migrator /usr/lib/one/ruby/onedb/shared/4.2.0_to_4.3.80.rb
> Done in 0.68s

> Running migrator /usr/lib/one/ruby/onedb/shared/4.3.80_to_4.3.85.rb
> Done in 0.00s

> Running migrator /usr/lib/one/ruby/onedb/shared/4.3.85_to_4.3.90.rb
> Done in 0.00s

> Running migrator /usr/lib/one/ruby/onedb/shared/4.3.90_to_4.4.0.rb

```
> Done in 0.00s

> Running migrator /usr/lib/one/ruby/onedb/shared/4.4.0_to_4.4.1.rb
> Done in 0.00s

> Running migrator /usr/lib/one/ruby/onedb/shared/4.4.1_to_4.5.80.rb
> Done in 0.39s

Database migrated from 3.8.0 to 4.5.80 (OpenNebula 4.5.80) by onedb command.

>>> Running migrators for local tables
Database already uses version 4.5.80

Total time: 3.60s
```

Note: Make sure you keep the backup file. If you face any issues, the onedb command can restore this backup, but it won't downgrade databases to previous versions.

3.12.5 Check DB Consistency

After the upgrade is completed, you should run the command `onedb fsck`.

First, move the 4.0 backup file created by the upgrade command to a safe place.

```
$ mv /var/lib/one/mysql_localhost_opennebula.sql /path/for/one-backups/
```

Then execute the following command:

```
$ onedb fsck -S localhost -u oneadmin -p oneadmin -d opennebula
MySQL dump stored in /var/lib/one/mysql_localhost_opennebula.sql
Use 'onedb restore' or restore the DB using the mysql command:
mysql -u user -h server -P port db_name < backup_file

Total errors found: 0
```

3.12.6 Virtual Machine Directories

Note: Only for OpenNebula versions < 3.8.3

If you are upgrading from a version **lower than 3.8.3**, you need to move the Virtual Machine deployment files from `/var/lib/one` to `/var/lib/one/vms`:

```
$ mv /var/lib/one/[0-9]* /var/lib/one/vms
```

3.12.7 Driver Names

OpenNebula default driver names have changed in the configuration file. Now the names of the vmm and im drivers are not prepended by the type of driver:

- `vmm_kvmm` → `kvmm`

- vmm_xen → xen
- vmm_vmware → vmware
- vmm_ec2 → ec2
- vmm_dummy → dummy
- im_kvm → kvm
- im_xen → xen
- im_vmware → vmware
- im_ec2 → ec2
- im_ganglia → ganglia
- im_dummy → dummy

To keep your existing hosts working, you need to duplicate the drivers with the old names.

For example, for kvm you will have IM_MAD `kvm` and VM_MAD `kvm`, so you need to add IM_MAD `im_kvm` and VM_MAD `vmm_kvm`

```

IM_MAD = [
  name       = "kvm",
  executable = "one_im_ssh",
  arguments  = "-r 3 -t 15 kvm" ]

IM_MAD = [
  name       = "im_kvm",
  executable = "one_im_ssh",
  arguments  = "-r 3 -t 15 kvm" ]

VM_MAD = [
  name       = "kvm",
  executable = "one_vmm_exec",
  arguments  = "-t 15 -r 0 kvm",
  default    = "vmm_exec/vmm_exec_kvm.conf",
  type      = "kvm" ]

VM_MAD = [
  name       = "vmm_kvm",
  executable = "one_vmm_exec",
  arguments  = "-t 15 -r 0 kvm",
  default    = "vmm_exec/vmm_exec_kvm.conf",
  type      = "kvm" ]

```

3.12.8 Manual Intervention Required

Note: Ignore this section if onedb didn't output the following message

If you have a datastore configured to use a tm driver not included in the OpenNebula distribution, the onedb upgrade command will show you this message:

```

ATTENTION: manual intervention required

The Datastore <id> <name> is using the

```

```
custom TM MAD '<tm_mad>'. You will need to define new
configuration parameters in oned.conf for this driver, see
http://opennebula.org/documentation:rel4.4:upgrade
```

Since OpenNebula 4.4, each `tm_mad` driver has a `TM_MAD_CONF` section in `oned.conf`. If you developed the driver, it should be fairly easy to define the required information looking at the existing ones:

```
# The configuration for each driver is defined in TM_MAD_CONF. These
# values are used when creating a new datastore and should not be modified
# since they define the datastore behaviour.
# name      : name of the transfer driver, listed in the -d option of the
#            TM_MAD section
# ln_target : determines how the persistent images will be cloned when
#            a new VM is instantiated.
#           NONE: The image will be linked and no more storage capacity will be used
#           SELF: The image will be cloned in the Images datastore
#           SYSTEM: The image will be cloned in the System datastore
# clone_target : determines how the non persistent images will be
#              cloned when a new VM is instantiated.
#           NONE: The image will be linked and no more storage capacity will be used
#           SELF: The image will be cloned in the Images datastore
#           SYSTEM: The image will be cloned in the System datastore
# shared    : determines if the storage holding the system datastore is shared
#            among the different hosts or not. Valid values: "yes" or "no"

TM_MAD_CONF = [
    name      = "lvm",
    ln_target  = "NONE",
    clone_target= "SELF",
    shared    = "yes"
]
```

3.12.9 Update the Drivers

You should be able now to start OpenNebula as usual, running 'one start' as `oneadmin`. At this point, execute `onehost sync` to update the new drivers in the hosts.

Warning: Doing `onehost sync` is important. If the monitorization drivers are not updated, the hosts will behave erratically.

3.12.10 Setting new System DS

With the new multi-system DS functionality, it is now required that the system DS is also part of the cluster. If you are using System DS 0 for Hosts inside a Cluster, any VM saved (stop, suspend, undeploy) **will not be able to be resumed after the upgrade process**.

You will need to have at least one system DS in each cluster. If you don't already, create new system DS with the same definition as the system DS 0 (TM_MAD driver). Depending on your setup this may or may not require additional configuration on the hosts.

You may also try to recover saved VMs (stop, suspend, undeploy) following the steps described in this [thread of the users mailing list](#).

3.12.11 Create the Security Group ACL Rule

There is a new kind of resource introduced in 4.12: Security Groups. If you want your existing users to be able to create their own Security Groups, create the following ACL Rule:

```
$ oneacl create "*" SECGROUP/* CREATE *
```

3.12.12 Create the Virtual Router ACL Rule

There is a new kind of resource introduced in 5.0: Virtual Routers. If you want your existing users to be able to create their own Virtual Routers, create the following ACL Rule:

```
$ oneacl create "*" VROUTER/* CREATE *
```

Note: For environments in a Federation: This command needs to be executed only once in the master zone, after it is upgraded to 5.4.

3.12.13 Testing

OpenNebula will continue the monitoring and management of your previous Hosts and VMs.

As a measure of caution, look for any error messages in `oned.log`, and check that all drivers are loaded successfully. After that, keep an eye on `oned.log` while you issue the `onevm`, `onevnet`, `oneimage`, `oneuser`, `onehost` **list** commands. Try also using the **show** subcommand for some resources.

3.12.14 Restoring the Previous Version

If for any reason you need to restore your previous OpenNebula, follow these steps:

- With OpenNebula 5.4 still installed, restore the DB backup using `'onedb restore -f'`
- Uninstall OpenNebula 5.4, and install again your previous version.
- Copy back the backup of `/etc/one` you did to restore your configuration.

3.12.15 Known Issues

If the MySQL database password contains special characters, such as `@` or `#`, the `onedb` command will fail to connect to it.

The workaround is to temporarily change the `oneadmin`'s password to an ASCII string. The `set password` statement can be used for this:

```
$ mysql -u oneadmin -p
mysql> SET PASSWORD = PASSWORD('newpass');
```

3.13 vCenter upgrade 5.2 to 5.4

3.13.1 Pre-migration phase

OpenNebula provides a script that must be run **before** it is upgraded using the **oneadmin** user account. This script can be downloaded from https://downloads.opennebula.org/packages/opennebula-5.4.0/vcenter_one54_pre.rb.

Warning: If you are using the vCenter drivers, there is a manual intervention required in order to upgrade to OpenNebula 5.4. Note that **upgrading from OpenNebula < 5.2 to OpenNebula 5.4 is NOT supported**. You need to upgrade first to OpenNebula 5.2, and then upgrade to OpenNebula 5.4.

Warning: The pre-migration phase may fail if there are resources in error, please clean resources in failed state prior to continue with this process.

The script will perform the following tasks:

- Establish a connection to every vCenter instance known by OpenNebula
- Retrieve information about clusters, virtual machines, templates, datastores and port groups.
- New information will be added to the OpenNebula resources.
- Some manual intervention may be required.
- For each IMAGE datastore found, a SYSTEM datastore will be created.
- Templates and wild VMs will be inspected in order to discover virtual hard disks and network interface cards that are invisible.
- All Datastores that hosts those virtual hard disks will be imported into OpenNebula.
- OpenNebula images and virtual networks will be created so the invisible disks and nics become visible after upgrade.
- The virtual networks that represent port groups found inside existing templates will have an Ethernet address range with 255 MACs in the pool. You may want to change or increase this address range after the pre-migrator tool finishes.
- OpenNebula hosts, networks and datastores will be grouped into OpenNebula clusters. Each vCenter cluster will be assigned to an OpenNebula cluster.
- XML files will be generated under `/var/tmp` directory. They will be used in the migration phase.

Important: Read carefully the instructions of the Phase 0. It involves modifying `/etc/one/oned.conf` and `/var/lib/one/remotes/datastore/vcenter/rm` and restarting OpenNebula. **DON'T FORGET TO DO SO.**

```

=====
/ _ \ _ _ _ _ _ _ _ _ | \ | | _ _ | | _ _ _ | | _ _
| | | | ' _ \ / _ \ ' _ \ | \ | / _ \ ' _ \ | | | | / _ \
| | | | | ) | _ / | | | \ | | _ / | ) | | | | ( | |
 \ _ / | . _ / \ _ | | | | \ | \ _ | . _ / \ _ | | | \ _ /
  | |
-----

```

```
vCenter pre-migrator tool for OpenNebula 5.4 - Version: 1.0
=====
=====
PHASE 0 - Before running the script please read the following notes
=====

- Please check that you have set PERSISTENT_ONLY="NO" and REQUIRED_ATTRS=""
  in you DS_MAD_CONF vcenter inside the /etc/one/oned.conf and that you have
  restarted your OpenNebula services to apply the new configuration before
  launching the script.

- Edit the file /var/lib/one/remotes/datastore/vcenter/rm and replace the
  following lines:

vi_client.delete_virtual_disk(img_src,
                              ds_name)

with the following lines:

if drv_action["/DS_DRIVER_ACTION_DATA/IMAGE/TEMPLATE/VCENTER_IMPORTED"] != "YES"
  vi_client.delete_virtual_disk(img_src,ds_name)
end

in order to avoid that you accidentally remove a virtual hard disk from a template
or wild VM when you delete an image.

- Note that this script may take some time to perform complex tasks so please be
↪patient.

- Although this scripts will do its best to be fully automated there may be situations
  where a manual intervention is needed, in that case a WARNING will be shown.

- The virtual networks that represent port groups found inside existing templates
  will have an Ethernet address range with 255 MACs in the pool. You may want to
  change or increase this address range after the pre-migrator tool finishes.

- It's advisable to disable the Sunstone user interface before launching this script
  in order to avoid that OpenNebula objects created by users while
  the script is running are not pre-migrated by the tool.

- This script can be executed as many times as you wish. It will update previous
  results and XML template will be always overwritten.

Don't forget to restart OpenNebula if you have made changes!

Do you want to continue? ([y]/n):
```

In short, you need to replace the following in /etc/one/oned.conf:

```
DS_MAD_CONF = [
-   NAME = "vcenter", REQUIRED_ATTRS = "VCENTER_CLUSTER", PERSISTENT_ONLY = "YES",
+   NAME = "vcenter", REQUIRED_ATTRS = "", PERSISTENT_ONLY = "NO",
  MARKETPLACE_ACTIONS = "export"
]
```

And the following change in /var/lib/one/remotes/datastore/vcenter/rm:

```
-vi_client.delete_virtual_disk(img_src,  
-                               ds_name)  
+if drv_action["/DS_DRIVER_ACTION_DATA/IMAGE/TEMPLATE/VCENTER_IMPORTED"] != "YES"  
+  vi_client.delete_virtual_disk(img_src, ds_name)  
+end
```

Note: It's advisable to disable the Sunstone user interface while the pre-migrator script is run in order to avoid that OpenNebula objects created by users while the script is run are not pre-migrated.

This script can be executed as many times as wished. It will update previous results and XML template will be always overwritten.

In order to execute the script you need to download from https://downloads.opennebula.org/packages/opennebula-5.4.0/vcenter_one54_pre.rb and run it manually **as onedadmin**.

```
$ curl -skLO https://downloads.opennebula.org/packages/opennebula-5.4.0/vcenter_one54_  
→pre.rb  
$ ruby vcenter_one54_pre.rb
```

3.13.2 OpenNebula Upgrade

Important: Now you need to continue upgrading the software following the steps described in the upgrade guide.

Follow the *Upgrade OpenNebula software*.

3.13.3 Migration phase

Once OpenNebula packages have been upgraded, you need to execute the pre migration tool for vCenter.

Warning: The migration tool must be run **before** a onedb upgrade command is executed.

The migration tool is launched using the `onedb vcenter-one54` command, and it must be run from the same machine where the pre-migrator tool was executed as it requires some XML templates files stored in the `/var/tmp` directory.

```
$ onedb vcenter-one54 -v -u <dbuser> -p <dbpass> -d <dbname> -S <dbhost>
```

The migration tool will update some OpenNebula's database tables using the XML files that were created in the pre-migration phase. This is the list of affected tables:

- `template_pool`
- `vm_pool`
- `host_pool`
- `datastore_pool`
- `network_pool`
- `image_pool`

In the following sections you will need to execute `onedb fsck`. Note that you might get the following error: `[UNREPAIRED] VM XX has a lease from VNet XX, but it could not be matched to any AR`. This is expected for previously invisible NIC interfaces in VMs added in the pre-migration phase.

Continue the upgrade by moving on to the *next section*.