



vNIOSTM Installation Guide for KVM Hypervisor and KVM-based OpenStack

1. Infoblox Installation Guide vNIOS™ for KVM Hypervisor and KVM-based OpenStack	3
1.1 About vNIOS Virtual Appliance for KVM	4
1.1.1 vNIOS for KVM Virtual Appliance Models	5
1.1.2 Supported vNIOS Versions	7
1.2 Deploying vNIOS for KVM	8
1.2.1 Installing vNIOS Virtual Appliance in the KVM Environment	9
1.2.2 Installing vNIOS for KVM in the OpenStack Environment	13
1.3 Deploying vNIOS for KVM in OpenStack Using Elastic Scaling	18
1.3.1 Provisioning the Grid Master and Grid Members	19
1.4 Setting Up a Grid	22
1.4.1 Configuring vNIOS Appliances as Grid Masters	23
1.4.2 Configuring vNIOS Appliances as Grid Members	25
1.4.3 Verifying and Monitoring	27
1.5 Auto Scaling for Virtual DNS Cache Acceleration	28
1.6 Known Limitations	30

Copyright Statements

© 2019, Infoblox Inc.— All rights reserved.

The contents of this document may not be copied or duplicated in any form, in whole or in part, without the prior written permission of Infoblox, Inc.

The information in this document is subject to change without notice. Infoblox, Inc. shall not be liable for any damages resulting from technical errors or omissions which may be present in this document, or from use of this document.

This document is an unpublished work protected by the United States copyright laws and is proprietary to Infoblox, Inc. Disclosure, copying, reproduction, merger, translation, modification, enhancement, or use of this document by anyone other than authorized employees, authorized users, or licensees of Infoblox, Inc. without the prior written consent of Infoblox, Inc. is prohibited.

For Open Source Copyright information, refer to the *Infoblox NIOS Documentation*.

Trademark Statements

Infoblox, the Infoblox logo, Grid, NIOS, bloxTools, NetMRI, Network Automation, and PortIQ are trademarks or registered trademarks of Infoblox Inc.

All other trademarked names used herein are the properties of their respective owners and are used for identification purposes only.

Company Information

<http://www.infoblox.com/contact>

Warranty Information

Your purchase includes a 90-day software warranty and a one year limited warranty on the Infoblox appliance, plus an Infoblox Warranty Support Plan and Technical Support. For more information about Infoblox Warranty information, refer to Infoblox Web site, or contact Infoblox Technical Support.

About vNIOS Virtual Appliance for KVM

Infoblox vNIOS for KVM is a virtual appliance designed for KVM (Kernel-based Virtual Machine) hypervisor and KVM-based OpenStack deployments. The Infoblox vNIOS for KVM enables you to deploy large, robust, manageable and cost effective Infoblox Grids. For information about Infoblox Grids, refer to the *Infoblox NIOS Documentation*. Note that the vNIOS for KVM functions as a hardware virtual machine guest on the Linux system.

Infoblox NIOS provides core network services and a framework for integrating all the components of the modular Infoblox solution. NIOS provides integrated, secure, and easy-to-manage DNS (Domain Name System), DHCP (Dynamic Host Configuration Protocol) and IPAM (IP address management) services. In addition to DNS, DHCP and IPAM, NIOS also provides TFTP, HTTP, and FTP file transfer services. Infoblox vNIOS for KVM provides most of the features supported by the NIOS, with some limitations described in [Known Limitations](#).

You can configure most of the vNIOS appliances as independent or HA (high availability) Grid Masters, Grid Master Candidates, and Grid members (or a reporting member). [Table 1.1](#) lists the supported vNIOS appliance models and their specifications.

vNIOs for KVM Virtual Appliance Models

vNIOs Appliances	Overall Disk (GB)	# of CPU Cores	Memory Allocation (GB)	Supported as Grid Master and Grid Master Candidate
IB-VM-805 (Reporting only)	250 (+ 1 TB user defined reporting storage)	2	32	No
IB-VM-815	250	2	16	Yes
IB-VM-825	250	2	16	Yes
IB-VM-1405 (Reporting only)	250 (+ 1.2 TB user defined reporting storage)	4	32	No
IB-VM-1415	250	4	32	Yes
IB-VM-1425	250	4	32	Yes
IB-VM-2205 (Reporting only)	250 (+ 2.4 TB user defined reporting storage)	8	64	No
IB-VM-2215	250	8	64	Yes
IB-VM-2225	250	8	64	Yes

Network Insight Virtual Appliances	Overall Disk (GB)	# of CPU Cores	Memory Allocation	Supported as Grid Master and Grid Master Candidate
ND-V805 *1	250	2	32	No
ND-V1405 *1	250	4	32	No
ND-V2205 *1	250	8	32	No

Cloud Platform Virtual Appliances	Overall Disk (GB)	# of CPU Cores	Memory Allocation	Supported as Grid Master and Grid Master Candidate
CP-V805 *2	250	2	16	No
CP-V1405 *2	250	4	32	No
CP-V2205 *2	250	8	64	No

Note

NIOs for KVM is supported in the following environments: OpenStack, Red Hat Enterprise Linux 7.4, 7.5, and Ubuntu 16.04. To achieve best performance on your virtual appliances, follow the recommended specifications and allocate your resources within the limits of the licenses being installed on the appliances.

Only IB-V1405 as a Reporting server has been qualified for OpenStack.

NIOs virtual appliance for Hyper-V is not recommended as a Grid Master or Grid Master Candidate.

¹ ND virtual appliances are designed for Network Insight only. Discovery is supported in OpenStack only with SRIOV enabled.

² CP-V805, CP-V1405, and CP-V2205 do not support downgrading from NIOs 8.4.x to any earlier NIOs releases and is only supported in NIOs 8.4 and later releases.

Requirements

The following are required for the installation of the vNIOs for KVM virtual appliance:

- The qcow2 file for the specified vNIOs virtual appliance model. This is the vNIOs software package that you download from the Infoblox Technical Support site. To download the package, you must have a valid login account on the Infoblox Support site. Register your product at <https://support.infoblox.com> if you do not already have an account. Make sure that you download the file with a name that

corresponds to the appliance model number. For example, you can download this file: `nios-7.3.0-314352-2016-01-29-05-02-02-160G-1410-disk1.qcow2`. For more information about the supported vNIOS for KVM appliance models, see [Table 1.1](#). The current qcow2 file format is compatible with KVM/QEMU version 2.1.3 or higher.

Note

To configure a reporting server, ensure that you download two qcow2 files: one with "nios...disk1.qcow2" and the other "nios...disk2.qcow2" extensions.

- Any of the following Linux distributions:
 - RHEL 7.4 for KVM-based OpenStack
 - RHEL 7.5 for KVM-based OpenStack
 - Ubuntu 16.04 for KVM-based OpenStack

Supported vNIOS Versions

vNIOS for KVM are supported for NIOS versions 7.2.x, 7.3.x, and 8.x. Note that NIOS 7.2.x DOES NOT support the following features:

- Elastic scaling for the Grid Master
- Elastic scaling with temporary licenses
- Elastic scaling for vNIOS for KVM in OpenStack with DHCP enabled
- Reporting and the ND virtual appliances
- IPv6
- HA (High Availability)

Note

For NIOS 7.2.x, you must configure the LAN1 IP address and network properties through cloud-init or manually through the Infoblox console CLI after the vNIOS virtual appliance has started.

You can configure the vNIOS virtual appliance as a Grid member or a reporting member. vNIOS for KVM supports the following reporting appliances for NIOS 8.0.0 version: IB-V805, IB-V1405, IB-V2205, and IB-V4005. When you set up a vNIOS reporting virtual appliance, it is used solely for reporting purposes. You cannot add licenses to run other services, such as DNS and DHCP, on a reporting appliance. The Infoblox reporting solution automates the collection, analysis, and presentation of core network service data that assists you in planning and mitigating network outage risks so you can manage your networks more efficiently. For more information about Infoblox Grids and reporting solution, refer to the *Infoblox NIOS Documentation*. Note that you must create an instance that has a second disk associated with it. For more information see, [Setting Up vNIOS OpenStack Flavors](#).

Deploying vNIOs for KVM

You can deploy the vNIOs appliance from a remote web server or a local file system accessible from your management system. Instructions in this section assume that you have configured the server on your network, and you are able to connect to it from your management system. To deploy a vNIOs appliance for KVM, complete the following:

1. Install relevant package for the operating system, either CentOS for KVM or RHEL (Red Hat Enterprise Linux) for KVM-based OpenStack, that you use:
 - In your CentOS environment, verify that you have already installed the device-mapper package. If not, run the following commands to install the package:

```
yum install device-mapper
service libvirtd restart
```
 - In your RHEL environment, verify that you have already installed the gnome package. If not, you can run the following commands to install the package:

```
yum group list
```

The above command lists all available installation groups on the RHEL environment. Next, run the following command to install gnome package:

```
yum groupinstall 'Server with GUI'
```

Alternatively, you can use the following command to install only the core gnome package:

```
yum groupinstall 'X Window System' 'GNOME'
```
2. Download the qcow2 file from the Infoblox Technical Support site. For information, see [Requirements](#).
3. Install the vNIOs virtual appliance in KVM-only or KVM-based OpenStack environment, as described in the following sections:
 - For KVM, see [Installing vNIOs Virtual Appliance in the KVM Environment](#).
 - For KVM-based OpenStack, see [Installing vNIOs for KVM in the OpenStack Environment](#).
4. Configure and start the vNIOs virtual appliance, as described in the following sections:
 - For KVM, see [Configuring the vNIOs Instance in the KVM Environment](#).
 - For KVM-based OpenStack, see [Starting a vNIOs Instance in OpenStack Environment](#).

Note

Infoblox recommends that you back up your existing configuration before deploying vNIOs for KVM.

Installing vNIOS Virtual Appliance in the KVM Environment

1. Depending on which KVM Hypervisor you are using, upload the qcow2 file(s) for the specified vNIOS virtual appliance model to the KVM/libvirt environment. If you are deploying a vNIOS reporting server (which requires two disk images), you must upload both the xxx-disk1.qcow2 and xxx-disk2.qcow2 files, and then create an XML file to define the vNIOS instance in KVM. Note that you can download sample XML files from the **Downloads** page for specific vNIOS versions, available on the Infoblox Support web site. For supported appliance models, see [Table 1.1](#).
2. Set up the following networks for your KVM environment. You may need to manually modify certain files to suit your environment. Refer to the respective documentation for your KVM Hypervisor.
 - MGMT_network
 - LAN1_network
 - LAN2_network
3. Create XML files for the vNIOS virtual appliance you want to deploy, as described in [Defining XML files for vNIOS Appliances](#); and then open the KVM console and execute the following commands to define and start the vNIOS virtual appliance. You may need to include the paths for the XML files if you save them in a different directory.

```
virsh net-define <Name of the MGMT XML file>
virsh net-define <Name of the LAN1 XML file>
virsh define <Name of the vnios XML file>
virsh start <VM Name>
```
4. Optionally, if you are deploying a vNIOS reporting server, create an XML file as described in [Defining an XML File for Reporting Servers](#), and then open the KVM console and execute the following commands to define and start the vNIOS reporting instance:

```
chown -R qemu.qemu <the directory to which you uploaded the qcow2 files> (For example, if you store the qcow2 files in /storage/vm/reporting800, enter /storage/vm/reporting800.)
virsh define <XML file> (You may need to include the path for the XML file if you save it in a different directory.)
virsh start <VM Name>
```
5. Configure the vNIOS instance as described in [Configuring the vNIOS Instance in the KVM Environment](#).

Defining XML files for vNIOS Appliances

Instead of deploying a vNIOS virtual appliance through the GUI, you can create the following XML files for the appliance and then execute commands to define and start the appliance in KVM. For information about the commands, see [Installing vNIOS Virtual Appliance in the KVM Environment](#).

- XML file for defining the MGMT interface, as shown in [Sample XML File for MGMT](#).
- XML file for defining the LAN1 interface, as shown in [Sample XML File for LAN1](#).
- XML file for defining vNIOS information such as the name of the VM, memory size, number of CPUs, location of the qcow2 files, file format, and others, as shown in [Sample XML File for a vNIOS Appliance](#).

Sample XML File for MGMT

```
<network>
  <name>MGMT</name>
  <forward mode='bridge' />
  <bridge name='virbr0' />
</network>
```

Sample XML File for LAN1

```
<network>
  <name>LAN1</name>
  <forward mode='bridge' />
  <bridge name='virbr0' />
</network>
```

Sample XML File for a vNIOS Appliance

Following is a sample XML file for defining a vNIOS virtual appliance in KVM. Note that the VM name, memory, vCPU, and location of the qcow2 file (highlighted in **red** in the following example) may vary. You can change these parameters according to your deployment.

```
<domain type='kvm'>
  <name> Infoblox-TE-820 </name>
  <memory unit='KiB'> 2097152 </memory>
  <vcpu placement='static'> 2 </vcpu>
  <os>
    <type arch='x86_64' machine=' pc ' >hvm</type>
```

```

    <boot dev='hd' />
  </os>
<features>
  <acpi />
  <apic />
  <pae />
</features>
<clock offset='utc' />
<on_poweroff>destroy</on_poweroff>
<on_reboot>restart</on_reboot>
<on_crash>restart</on_crash>
  <devices>
    <emulator> /usr/bin/qemu-system-x86_64 </emulator>
    <disk type='file' device='disk'>
      <driver name='qemu' type='qcow2' cache='none' />
      <source
file=' /var/lib/libvirt/images/nios-7.3.2-316478-2016-02-17-19-34-52-55G-820-disk1.q cow2 ' />
      <target dev='vda' bus='virtio' />
      <address type='pci' domain='0x0000' bus='0x00' slot='0x04' function='0x0' />
    </disk>
    <interface type='network'>
      <rom bar='off' />
      <source network='MGMT' />
      <model type='virtio' />
      <address type='pci' domain='0x0000' bus='0x00' slot='0x03' function='0x0' />
    </interface>
    <interface type='network'>
      <source network='LAN1' />
      <rom bar='off' />
      <model type='virtio' />
      <address type='pci' domain='0x0000' bus='0x00' slot='0x06' function='0x0' />
    </interface>
    <serial type='pty'>
      <target port='0' />
    </serial>
    <console type='pty'>
      <target type='serial' port='0' />
    </console>
    <memballoon model='virtio'>
      <address type='pci' domain='0x0000' bus='0x00' slot='0x05' function='0x0' />
    </memballoon>
  </devices>
</domain>

```

Defining an XML File for Reporting Servers

If you are deploying a reporting server in KVM, you must define an XML file that includes information such as the name of the VM, memory size, number of CPUs, location of the qcow2 files, file format, and others before you can spin up the vNIOS instance.

Following is a sample XML file for defining a reporting server in KVM. Note that the VM name, memory, vCPU, and location of the qcow2 files (highlighted in red in the following example) may vary. You can change these parameters according to your deployment.

Depending on the KVM tool you are using to deploy the reporting server, you can use the following sample file as a reference to create your own XML file.

```

<domain type='kvm'>
  <name> reporting800 </name>
  <memory unit='KiB'> 8388608 </memory>
  <vcpu placement='static'> 2 </vcpu>
  <os>
    <type arch='x86_64' machine=' pc ' >hvm</type>
    <boot dev='hd' />
  </os>
  <features>
    <acpi />
    <apic />
    <pae />
  </features>
  <clock offset='utc' />
  <on_poweroff>destroy</on_poweroff>

```

```

<on_reboot>restart</on_reboot>
<on_crash>restart</on_crash>
<devices>
<emulator> /usr/bin/qemu-system-x86_64 </emulator>
<disk type='file' device='disk'>
  <driver name='qemu' type='qcow2' cache='none' />
  <source
file= ' /storage/vm/reporting800/nios-7.2.6-316673-2016-02-18-14-00-10-300G-800-disk1.qcow2 ' />
  <target dev='vda' bus='virtio' />
  <address type='pci' domain='0x0000' bus='0x00' slot='0x04' function='0x0' />
</disk>
<disk type='file' device='disk'>
<driver name='qemu' type='qcow2' cache='none' />
<source
file= ' /storage/vm/reporting800/nios-7.2.6-36673-2016-02-18-14-00-10-300G-800-disk2.qcow2 ' />
  <target dev='vdb' bus='virtio' />
  <address type='pci' domain='0x0000' bus='0x00' slot='0x09' function='0x0' />
</disk>
<interface type='network'>
  <rom bar='off' />
  <source network='MGMT_network' />
  <model type='virtio' />
  <address type='pci' domain='0x0000' bus='0x00' slot='0x03' function='0x0' />
</interface>
<interface type='network'>
  <source network='LAN1_network' />
  <rom bar='off' />
  <model type='virtio' />
  <address type='pci' domain='0x0000' bus='0x00' slot='0x06' function='0x0' />
</interface>
<interface type='network'>
  <source network='HA_network' />
  <rom bar='off' />
  <model type='virtio' />
  <address type='pci' domain='0x0000' bus='0x00' slot='0x07' function='0x0' />
</interface>
<interface type='network'>
  <source network='LAN2_network' />
  <rom bar='off' />
  <model type='virtio' />
  <address type='pci' domain='0x0000' bus='0x00' slot='0x08' function='0x0' />
</interface>
<serial type='pty'>
  <target port='0' />
</serial>
<console type='pty'>
  <target type='serial' port='0' />
</console>
<memballoon model='virtio'>
  <address type='pci' domain='0x0000' bus='0x00' slot='0x05' function='0x0' />
</memballoon>
</devices>
</domain>

```

Configuring the vNIOS Instance in the KVM Environment

To configure the vNIOS instance:

1. From the KVM administration tool, such as Virtual Machine Manager, select the vNIOS instance.
2. Open the KVM console.
3. When the Infoblox login prompt appears, log in with the default user name and password.

```

login: admin
password: infoblox
The Infoblox prompt appears: Infoblox >

```
4. You must have valid licenses before you can configure the vNIOS appliance. To obtain permanent licenses, first use the `Infoblox > show version` command to obtain the serial number of the vNIOS appliance, and then visit the Infoblox Support web site at <https://support.infoblox.com>. Log in with the user ID and password you receive when you register your product online at <http://www.infoblox.com/sup>

[port/customer/evaluation-and-registration](#).

If the vNIOS virtual appliance does not have the Infoblox licenses required to run NIOS services and to join a Grid, you can use the `set temp_license` command to generate and install a temporary 60-day license.

5. From the list of licenses, select to add the **Grid**, **vNIOS**, and other relevant licenses for your vNIOS virtual appliance. For the vNIOS reporting appliance, you must also select the **Reporting** license.

Note

You must have both the **Grid** and **vNIOS** licenses for the vNIOS virtual appliance to join a Grid.

6. Use the CLI command `set network` to configure the network settings.

```
Infoblox > set network
```

```
NOTICE: All HA configurations are performed from the GUI. This interface is used only to configure a standalone node or to join a Grid.
```

```
Enter IP address: 10.1.1.22
```

```
Enter netmask: [Default: 255.255.255.0]: 255.255.255.0
```

```
Enter gateway address [Default: 10.1.1.1]: 10.1.1.1
```

```
Become Grid member? (y or n): n
```

Note

For an HA Grid Master, ensure that you specify these settings for both the active and passive nodes.

After you confirm your network settings, the Infoblox Grid Manager automatically restarts. You can then proceed to setting up a Grid, as described in [Setting Up a Grid](#).

Managing vNIOS Instances

You can use the following commands to manage the vNIOS for KVM through a GUI virt-manager:

`virsh console <VM name>`: To access the VM console.

`virsh start <VM name>`: To start a VM.

`virsh list` and `virsh list --all`: To list VMs defined on the host. Note that the `--all` option includes VMs that are currently offline.

`virsh undefine <VM name>`: To remove the VM definition from the KVM environment. Use this command only after the VM has been shut down.

Installing vNIOS for KVM in the OpenStack Environment

Note

Before you issue commands with Nova, ensure that your environment contains the necessary credentials. You can do this by sourcing the `keystone_admin` file that is created during the OpenStack installation. For more information, refer to the section *Getting Credentials for a CLI* in the *OpenStack CLI Guide*.

To install vNIOS for KVM in OpenStack:

1. In OpenStack, run `source keystone_admin` to set up the OpenStack environment.
2. Upload the `qcow2` file for the specified vNIOS for KVM model to OpenStack. For more information, see [Requirements](#).
3. Set up the OpenStack flavors, as described in [Setting Up vNIOS OpenStack Flavors](#).
4. Import the vNIOS instance into OpenStack, as described in [Importing vNIOS Instance into OpenStack](#).
5. Set up security groups, as described in [Setting Up Security Groups](#).
6. Start an instance, as described in [Starting a vNIOS Instance in OpenStack Environment](#).

Setting Up vNIOS OpenStack Flavors

After you upload the `qcow2` file, set up the OpenStack flavors for your vNIOS models. Each flavor corresponds to different vCPU, RAM, disk size, and functionality.

Infoblox enables you to choose the size of the virtual disk that you use for reporting. To do so, you must create a vNIOS instance and associate an additional disk with it. You can add an ephemeral disk in the flavor that is used to create an instance. Note that the value for the ephemeral disk must be a non-zero.

To set up the vNIOS OpenStack flavors, run the following command:

```
nova flavor-create --is-public true <name> <ID> <Memory> <disk> <cpu> --swap 0
--ephemeral 0
where
```

- `name` defines the name for the vNIOS for KVM instance. For reporting, mention the name of the reporting model.
- `ID` defines the unique OpenStack flavor ID for the KVM instance.
- `memory` `disk` and `cpu` specify the flavors of the vNIOS for KVM instance.
- `ephemeral` defines the additional disk that is required to configure the reporting model.

Following is a sample command:

```
nova flavor-create --is-public true vnios-1420.160 6 8192 160 4 --swap 0 --ephemeral 0
```

For list of vNIOS OpenStack flavors with their specifications, refer [vNIOS for KVM Virtual Appliance Models](#).

To manage the vNIOS OpenStack flavors, use the standard Nova API. For example, run `nova flavor-list` to show the created flavors.

Importing vNIOS Instance into OpenStack

Use the following command to create the vNIOS instance in OpenStack:

```
glance image-create --name <the name of the vNIOS image> --visibility public
--container-format bare --disk-format qcow2 --file <location of the image>
```

Example:

```
glance image-create --name vnios-820 --visibility public --container-format bare
--disk-format qcow2 --file
/tmp/nios-7.3.3-318825-2016-03-04-23-16-19-55G-820-disk1.qcow2
```

Setting Up Security Groups

When you set up your OpenStack environment, you can create an additional security group "vnios-sec-group" or add certain protocol rules to existing or default security groups to allow specific network traffic. You can configure basic settings as described in [Basic Configuration](#), or configure optional settings as described in [Advanced Configuration](#). These sections contain sample scripts you can use to establish specific protocol rules.

Basic Configuration

Creating security group "vnios-sec-group":

```
#vnios security group
neutron security-group-rule-create vnios-sec-group
```

HTTPS communications:

```
# https
neutron security-group-rule-create --protocol tcp --port-range-min 443 --port-range-max
443 --ethertype IPv4 vnios-sec-group
neutron security-group-rule-create --protocol tcp --port-range-min 443 --port-range-max
443 --ethertype IPv6 vnios-sec-group
```

Advanced Configuration

Grid communications:

```
#tunnels
neutron security-group-rule-create --protocol udp --port-range-min 1023
--port-range-max 1023 --ethertype IPv4 vnios-sec-group
neutron security-group-rule-create --protocol udp --port-range-min 1023
--port-range-max 1023 --ethertype IPv6 vnios-sec-group
neutron security-group-rule-create --protocol udp --port-range-min 1194
--port-range-max 1195 --ethertype IPv4 vnios-sec-group
neutron security-group-rule-create --protocol udp --port-range-min 1194
--port-range-max 1195 --ethertype IPv6 vnios-sec-group
neutron security-group-rule-create --protocol udp --port-range-min 2114
--port-range-max 2114 --ethertype IPv4 vnios-sec-group
neutron security-group-rule-create --protocol udp --port-range-min 2114
--port-range-max 2114 --ethertype IPv6 vnios-sec-group
neutron security-group-rule-create --protocol udp --port-range-min 802 --port-range-max 802
--ethertype IPv4 vnios-sec-group
neutron security-group-rule-create --protocol udp --port-range-min 802 --port-range-max 802
--ethertype IPv6 vnios-sec-group
```

Optional for other protocols:

```
# dhcp
neutron security-group-rule-create --protocol udp --port-range-min 67 --port-range-max
69 --ethertype IPv4 vnios-sec-group
neutron security-group-rule-create --protocol udp --port-range-min 67 --port-range-max
69 --ethertype IPv6 vnios-sec-group
neutron security-group-rule-create --protocol tcp --port-range-min 647 --port-range-max 647
--ethertype IPv4 vnios-sec-group
neutron security-group-rule-create --protocol tcp --port-range-min 647 --port-range-max 647
--ethertype IPv6 vnios-sec-group
neutron security-group-rule-create --protocol tcp --port-range-min 546 --port-range-max
547 --ethertype IPv4 vnios-sec-group
neutron security-group-rule-create --protocol tcp --port-range-min 546 --port-range-max
547 --ethertype IPv6 vnios-sec-group
neutron security-group-rule-create --protocol udp --port-range-min 546 --port-range-max
547 --ethertype IPv4 vnios-sec-group
neutron security-group-rule-create --protocol udp --port-range-min 546 --port-range-max
547 --ethertype IPv6 vnios-sec-group
# ntp
neutron security-group-rule-create --protocol tcp --port-range-min 123 --port-range-max
123 --ethertype IPv4 vnios-sec-group
neutron security-group-rule-create --protocol tcp --port-range-min 123 --port-range-max
123 --ethertype IPv6 vnios-sec-group
neutron security-group-rule-create --protocol udp --port-range-min 123 --port-range-max
123 --ethertype IPv4 vnios-sec-group
neutron security-group-rule-create --protocol udp --port-range-min 123 --port-range-max
123 --ethertype IPv6 vnios-sec-group
# dns
neutron security-group-rule-create --protocol tcp --port-range-min 53 --port-range-max
53 --ethertype IPv4 vnios-sec-group
neutron security-group-rule-create --protocol tcp --port-range-min 53 --port-range-max
```

```

53 --ethertype IPv6 vnios-sec-group
neutron security-group-rule-create --protocol udp --port-range-min 53 --port-range-max
53 --ethertype IPv4 vnios-sec-group
neutron security-group-rule-create --protocol udp --port-range-min 53 --port-range-max
53 --ethertype IPv6 vnios-sec-group

# ftp
neutron security-group-rule-create --protocol tcp --port-range-min 20 --port-range-max
21 --ethertype IPv4 vnios-sec-group
neutron security-group-rule-create --protocol tcp --port-range-min 20 --port-range-max
21 --ethertype IPv6 vnios-sec-group

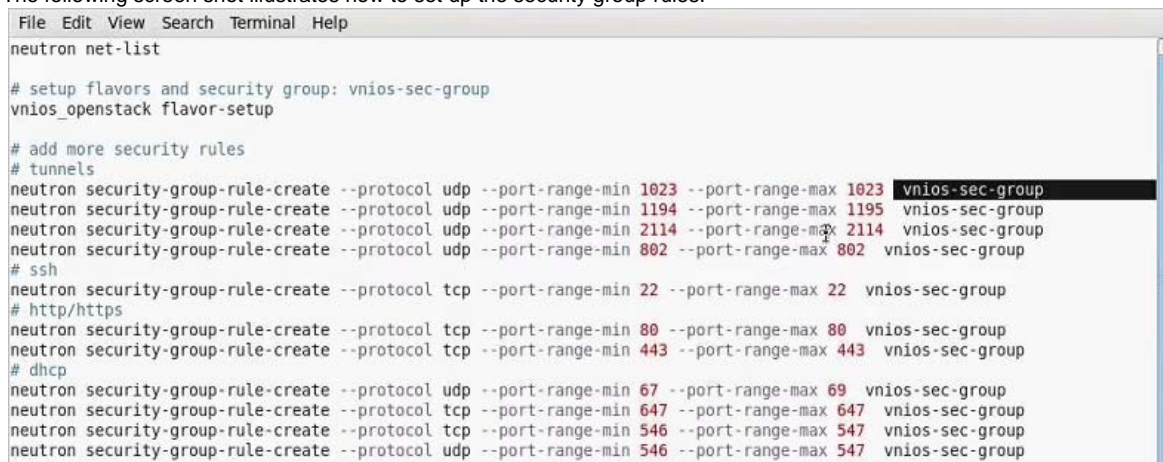
# syslog
neutron security-group-rule-create --protocol udp --port-range-min 514 --port-range-max
514 --ethertype IPv4 vnios-sec-group
neutron security-group-rule-create --protocol udp --port-range-min 514 --port-range-max
514 --ethertype IPv6 vnios-sec-group

# reporting
neutron security-group-rule-create --protocol tcp --port-range-min 9997
--port-range-max 9997 --ethertype IPv4 vnios-sec-group
neutron security-group-rule-create --protocol tcp --port-range-min 9997
--port-range-max 9997 --ethertype IPv6 vnios-sec-group

# ICMP
neutron security-group-rule-create --protocol icmp --ethertype IPv4 vnios-sec-group neutron
security-group-rule-create --protocol icmp --ethertype IPv6 vnios-sec-group

```

The following screen shot illustrates how to set up the security group rules.



Starting a vNIOS Instance in OpenStack Environment

To start a vNIOS instance:

1. Ensure that you have already specified the vNIOS flavors and provided a unique name for the instance you want to spin up. For the list of available flavors, see [vNIOS for KVM Virtual Appliance Models](#). Execute the `neutron port-create` command to create port IDs for the network interfaces (MGMT and LAN1/HA). You can use the `security-group` option to associate the vNIOS instance with the security group(s) you have created. For information, see [Setting Up Security Groups](#). (Optionally, you can associate the vNIOS instance with a security group when you execute the `nova boot` command.)

Note

Do not reuse the OpenStack neutron port of the deleted instances. When you reuse the neutron port of a deleted instance, a mismatch in the MAC address between the VM interface and the host VF might happen during NIC bonding. Also, the neutron port does not function properly when you reuse it repeatedly.

Following is a neutron example:

```
$ neutron port-create --security-group <name of the security group>
```

For an HA pair, you must also execute the `allowed-address-pairs` option to define the VIP port for the HA configuration, using the VRRP MAC address and the Virtual Router ID you use. Following is an example:

```
$ neutron port-create VIP --allowed-address-pairs list=true mac_address= 00:00:5e:00:01:c8 ip_address=10.0.0.22
```

2. For an HA pair configuration only, ensure that you set `allow_duplicate_networks=true` in the `nova.conf` file to remove the restriction of allowing only one interface for each network in OpenStack.
3. Run the `neutron port list` command to view the network and port IDs generated for all network interfaces so you can copy and paste them into the `nova boot` command.
4. Execute the `nova boot` command in OpenStack to spin up the vNIOS instance. (Note: Use the custom name you came up with when creating flavors).

Following is an example:

```
nova boot --config-drive False --image  
<nios-7.3.0-314352-2016-01-29-05-02-02-160G-1420-disk1.qcow2> --flavor <vnios1410.160>  
- security-groups <name of the security group> --nic net-id=<the network ID for the MGMT interface> -  
-nic net-id=<the network ID for the LAN1/HA interface only if you are configuring an HA pair> --nic  
port-id=<the IP address ID for the LAN1 interface>  
<my-vm-name>
```

where

- `image` defines the name of the software package you downloaded. For information about supported vNIOS for KVM models, see [vNIOS for KVM Virtual Appliance Models](#).
- `flavor` specifies the flavors of the vNIOS for KVM instance. For information about how to define flavors, see [Setting Up vNIOS OpenStack Flavors](#).
- `security-groups` defines the security group with which this vNIOS instance associates. For information about how to create rules for security groups, see [Setting Up Security Groups](#).
- `nic net-id` specifies the network ID for the MGMT interface. Note that when provisioning an HA pair, you must also specify the network ID for the LAN1/HA interface. For more information, see [Sample Commands for Provisioning an HA Pair](#).
- `nic port-id` specifies the IP address ID for the LAN1/HA interface.

Note

For the vNIOS appliance to run in OpenStack, you must specify at least two networks, MGMT and LAN1. To remove networks, use the `neutron net-delete` command. If some of the networks remain, use OpenStack Horizon to manually remove them.

- `my-vm-name` defines the unique name of the VM.
The vNIOS for KVM instance automatically spins up after the `nova boot` command is executed.

5. Go to OpenStack Horizon and select the previously launched instance.
6. Click the **Console** tab.
7. When the Infoblox login prompt appears, log in with the default user name and password.
login: **admin**
password: **infoblox**
The Infoblox prompt appears: `Infoblox >`
8. You must have valid licenses before you can configure the vNIOS appliance. To obtain permanent licenses, first use the `Infoblox > show version` command to obtain the serial number of the vNIOS appliance, and then visit the Infoblox Support web site at <https://support.infoblox.com>. Log in with the user ID and password you receive when you register your product online at <http://www.infoblox.com/support/customer/evaluation-and-registration>.
If the vNIOS virtual appliance does not have the Infoblox licenses required to run NIOS services and to join a Grid, you can use the `set temp_license` command to generate and install a temporary 60-day license.
9. From the list of licenses, select the Grid, vNIOS, and other relevant licenses for your vNIOS virtual appliance.

Note

You must have both the Grid and vNIOS licenses for the vNIOS virtual appliance to join a Grid (2 and 8 from the list).

10. In OpenStack Horizon, go to **InstanceOverview** and copy the floating IP address of the instance.
11. Go back to the console and run the `set network` command. Not required for Elastic Scaling.
12. Go to the Infoblox Grid Manager and enable the NAT mode for the Grid member:
 - a. Click **Grid** -> **Grid Manager** -> **Members** -> **Network**.
 - b. Select the Grid member and click **Edit**.
 - c. Click **Network** -> **Advanced**.
 - d. Click **Enable NAT Compatibility** and enter the floating IP address.
 - e. Click **Save & Close**.

Note

For an HA Grid Master, ensure that you specify these settings for both nodes.

After you confirm your network settings, the Infoblox Grid Manager automatically restarts. You can then proceed to set up a Grid, as described in [Setting Up a Grid](#).

Terminating vNIOS Instances

To terminate vNIOS instances, go to Horizon, select the instance and select "Terminate Instance" from the drop-down menu on the right hand side of the panel.

Deploying vNIOs for KVM in OpenStack Using Elastic Scaling

Before you provision vNIOs for KVM instances in OpenStack using the NIOS Elastic Scaling feature, ensure that you have the necessary feature licenses for each of your vNIOs for KVM instance, including dynamic licenses to support Elastic Scaling. For information about Elastic Scaling, refer to *Managing Licenses* in the *Infoblox NIOS Documentation*.

To use Elastic Scaling to provision Grid members, ensure that you understand how to compose user data files for provisioning the Grid Master and Grid members. A user data file includes configuration details such as enabling the remote console, installing licenses, defining network settings, and validating certificate and token for Elastic Scaling. For information about how to create a user data file, see [Defining User Data Settings for vNIOs for KVM Instances](#).

Provisioning the Grid Master and Grid Members

To set up the Grid Master and begin pre-provisioning vNIOS for KVM Grid members in OpenStack using Elastic Scaling, complete the following.

1. Log in to OpenStack and compose the user data file for the Grid Master using the vi editor. You can also compose the data file locally and move it to the OpenStack node later. Note that you can deploy the Grid Master VM using temporary licenses. Elastic Scaling is not supported on the Grid Master. For a sample of the Grid Master user data file, see [Sample User Data Files for the Grid Master](#).
2. Execute the `neutron port-create` command to create port IDs for the network interfaces (MGMT and LAN1/HA). You can use the `security-group` option to associate the vNIOS instance with the security group(s) you have created. For information, see [Setting Up Security Groups](#). (Optionally, you can associate the vNIOS instance with a security group when you execute the `nova boot` command.)

Following is an example:

```
$ neutron port-create --security-group <name of the security group>
```

For HA pairs, you must also execute the `allowed-address-pairs` option to define the VIP port for the HA configuration, using the VRRP MAC address and the Virtual Router ID you use. Following is an example:

```
$ neutron port-create VIP --allowed-address-pairs list=true mac_address= 00:00:5e:00:01:c8  
ip_address=10.0.0.22
```

3. For an HA pair configuration only, ensure that you set `allow_duplicate_networks=true` in the `nova.conf` file to remove the restriction of allowing only one interface for each network in OpenStack.
4. Run the `neutron port list` command to view the network and port IDs generated for all network interfaces so you can copy and paste them into the `nova boot` command.
5. Execute the `nova boot` command in OpenStack to spin up the Grid Master VM. (Note: Use the custom name you came up with when creating flavors).

Following is an example:

```
nova boot --config-drive False --image  
<nios-7.3.0-314352-2016-01-29-05-02-02-160G-1420-disk1.qcow2> --flavor <vnios1410.160>  
- security-groups <name of the security group> --nic net-id=<the network ID for the MGMT interface> -  
-nic net-id=<the network ID for the LAN1/HA interface only if you are configuring an HA pair> --nic  
port-id=<the IP address ID for the LAN1 interface>  
<my-vm-grid-master>
```

where

- a. `image` defines the name of the software package you downloaded. For information about supported vNIOS for KVM models, see [vNIOS for KVM Virtual Appliance Models](#).
- b. `flavor` specifies the flavors of the vNIOS for KVM instance. For information about how to define flavors, see [Setting Up vNIOS OpenStack Flavors](#).
- c. `security-groups` defines the security group with which this vNIOS instance associates. For information about how to create rules for security groups, see [Setting Up Security Groups](#).
- d. `nic net-id` specifies the network ID for the MGMT interface. Note that when provisioning an HA pair, you must also specify the network ID for the LAN1/HA interface. For more information, see [Sample Commands for Provisioning an HA Pair](#).
- e. `nic port-id` specifies the IP address ID for the LAN1/HA interface.

Note

For the vNIOS appliance to run in OpenStack, you must specify at least two networks, MGMT and LAN1/HA. To remove networks, use the `neutron net-delete` command. If some of the networks remain, use OpenStack Horizon to manually remove them.

- f. `my-vm-grid-master` defines the unique name of the VM.

The vNIOS for KVM instance automatically spins up after the `nova boot` command is executed.

6. Log in to the NIOS GUI (Grid Manager) and do the following,
 - Create offline Grid members you plan to join the Grid.
 - Pre-provision these Grid members.
 - Generate a token for each member. Copy this token and save it for use in each Grid member user data file. For detailed instructions on how to pre-provision a member, refer to the [Infoblox NIOS Documentation](#).
7. Log in to OpenStack and compose a user data file for each Grid member you plan to join the Grid. For a sample Grid member user data file, see [Sample User data File for Grid Members](#).
8. Execute the `nova boot` command in OpenStack to spin up each Grid member VM, as follows:

```
nova boot --config-drive False --image nios-7.2.4-1410-160.qcow2 --flavor vnios1410.160  
--nic net-id=9db90ecf-83e8-44c5-930d-7e3548ff4a02 --nic  
port-id=620d9fba-2f2d-4b81-9e51-eeefee551c15 --user-data ./user-data-2 my-vm-grid-member
```

Note

Use `config-drive True` to tell OpenStack to use the virtual CD-ROM drive transport mechanism for the user data file. For information about user data files, see [Defining User Data Settings for vNIOS for KVM Instances](#).

- `image` defines the name of the software package you downloaded. For information about supported vNIOS for KVM models, see [vNIOS for KVM Virtual Appliance Models](#).
- `flavor` specifies the flavors of the vNIOS for KVM instance. For information about how to define flavors, see [Setting Up vNIOS OpenStack Flavors](#).

- `nic net-id` specifies the MGMT interface.
- `nic port-id` specifies the LAN1 interface.
- `user-data` specifies the name of the user data file.
- `my-vm-grid-member` defines the name of the VM.

After you execute the `nova boot` command and launch the vNIOS for KVM instances, the Grid members automatically join the Grid. Further communications with the instance take place through Grid Manager and the NIOS CLI.

Defining User Data Settings for vNIOS for KVM Instances

When you provision appliances using Elastic Scaling, vNIOS for KVM instances in an OpenStack environment require different user data settings. In OpenStack, you compose the user data file using the `vi` editor in plain text format.

OpenStack supports two transport mechanisms for the user data file: one through a virtual CDROM drive (also known as ISO parameter injection) and the other through the metadata network service. You can provide either or both transport mechanisms. OpenStack will make the user data file available to the vNIOS instance using the transport you configure in the user data file.

You use the following data fields in the user data files for provisioning new instances using Elastic Scaling:

- `remote_console_enabled`: Set this field to "true" to enable the remote console.
- `default_admin_password`: Enter "infoblox" as the default admin password.
- `temp_license` or `license`: Enter the name of the licenses you plan to install on the Grid Master or Grid members. For example, you can enter "vnios,enterprise,dns". You install temporary licenses on the Grid Master when deploying it in the OpenStack environment.
- `lan1`: Specify the following LAN1 parameters only if you have disabled DHCP in OpenStack for the LAN1 network. Remove the entire `lan1` section from the user data file if DHCP is enabled in OpenStack. Otherwise, if the LAN1 parameters are the same as those of DHCP, the interpretation of the parameters stops at the LAN1 interface until you restart NIOS manually. If the LAN1 parameters are different from those of DHCP, NIOS restarts automatically and the LAN1 parameters are overwritten by the new DHCP response. Note that IPv6 on LAN1 is not supported in this release.
 - `v4_addr`: Use this field to specify the IP address of the Grid Master or Grid member instance.
 - `v4_netmask`: Use this field to specify the netmask of the Grid Master or Grid member instance.
 - `v4_gw`: Use this field to specify the gateway address of the Grid Master or Grid member instance.
- `gridmaster`: This field remains blank. It signifies to OpenStack that the instance is a Grid Master. Use this only for the Grid member user data file.
- `certificate`: Use this field only for the Grid member user data file. Copy and paste the string for the Infoblox NIOS certificate credential generated for the Grid member token. Note that the certificate string must not contain any space(s); otherwise, it might cause issues during the provisioning process.
- `token`: This field provides the string for the generated token for the new instance. Here, you copy and paste the security token generated for the new vNIOS for KVM instance.

Note

The certificate and token are generated on the Grid Master when you pre-provision the vNIOS instance. The certificate and token values are valid only for a period of time. For information, see *About Elastic Scaling* in the *Infoblox NIOS Documentation*.

- `ip_addr`: This is applicable only to Grid members. This field specifies the IP address for the Grid Master.

For sample user data files, see [Sample User Data Files for the Grid Master](#) and [Sample User data File for Grid Members](#).

Sample User Data Files for the Grid Master

Following is a sample user data file for deploying the Grid Master in the OpenStack environment:

```
#infoblox-config
remote_console_enabled: true default_admin_password: infoblox temp_license: vnios,enterprise,dns

lan1:
v4_addr: 10.2.0.9
v4_netmask: 255.255.255.0
v4_gw: 10.2.0.1
```

Sample User data File for Grid Members

Following is a sample user data file for deploying Grid members in the OpenStack environment:

```
#infoblox-config remote_console_enabled: true
default_admin_password: infoblox license: vnios,enterprise,dns
lan1:
v4_addr: 10.2.0.9
```

```

v4_netmask: 255.255.255.0
v4_gw: 10.2.0.1
gridmaster:
certificate: ----BEGIN CERTIFICATE----
MIIDdzCCAl8CEBdLzTDHhS3SgcInykFe/qUwDQYJKoZIhvcNAQEFBQAwEjELMAkGALUEBhMCVVMxEzARBgNVBA
gTCKNhbG1mb3JuaWExEjAQBgNVBAcTCVN1bm55dmFsZTERMA8GA1UEChMISW5mb2Jsb3gxZDASBgNVBAstC0Vu
Z2luZWVyaW5nMRkwFwYDVQQDExB3d3cuaW5mb2Jsb3guY29tMjB4XDTElMTAxNTIzNDgzMVoXDTE2MTAxNDIzND
gzMVoEjELMAkGALUEBhMCVVMxEzARBgNVBAgTCKNhbG1mb3JuaWExEjAQBgNVBAcTCVN1bm55dmFsZTERMA8G
A1UEChMISW5mb2Jsb3gxZDASBgNVBAstC0VuZ2luZWVyaW5nMRkwFwYDVQQDExB3d3cuaW5mb2Jsb3guY29tMI
IBIjANBgkqhkiG9w0BAQEFAAOCAQ8AMIIBCgKCAQEArGBBrcJls7UfPRY40lsW3+JSkX4Uw04ssx1lQQgJEI3X
Ia335wC5fP37wtGgeCflJwDMhF6Z3a7nLgx6RZN2cPeDHdLQp45+P6Xi4I6J1gXPL/TPhtrrDfsX3Lq337eUi5
3D30qTfz+NwMgrJU6SRzxXUOkt+Tx6VTwFkCThrVktXURhg4Ik8frVBI8qFTFDRIs+z106E09LzoScNjOQXKZP
b2uqPwBhSQ7PYgX+vbyXy2CAEzmsDyW7TQFzeZ+8xA/sqlQUXZa8AQtyLlMnx2T9upP9g9e00+UBpZrKhdVZP
VYaWgpM71S0wv1a0NLjwVyGgj0igIughaqz5A54QIDAQABMA0GCsGSIb3DQEBBQUAA4IBAQBf3pAAxFiJ3ADg
3Nc36e0MxZAv2TFHgg190PjexBQHyDs9jU+Q1dnSeiVVwgYBSwxLfnEJthnelGHG1mgN92TRDehCpTiIFRnBH8
RnedtHQqQ/cNLHgHpcUW0eJXSR6kCSGHTSCFWQI/ie4RHhg3vXuOXA4ZkOAAgCO+korRUKcRc2kjIM1vZnRf9H
rDci+HLCGGTH/dHdQNIjasPWYnSQa3RKEHb153THfGEJXF5VBYIsu321WPqhaMIlTg7Rj0C+4pD1XQ3Z2Qa7TP
JDeNKfy+VZVcNQGD4hCeBRGhhmj7X7TqnCdKea1PftEwug1X24x84tVn1Xpw6GmwrLWpi----END CERTIFICATE----
token: 6VPPn51m46cw0RI/9F1o3D1cVR0/dogB_ip_addr: 10.2.0.7

```

Sample User Data Files for the IB-FLEX Grid Master

Following is a sample user data file for deploying an IB-FLEX Grid Master using the Flex Grid Activation license in the OpenStack environment:

```

#infoblox-config
remote_console_enabled: y
hardware_type: IB-FLEX
temp_license: flex_grid
lan1:

    v4_addr: 10.39.51.33
    v4_netmask: 255.255.255.0
    v4_gw: 10.39.51.1

mgmt:

    v4_addr: 10.39.50.22
    v4_netmask: 255.255.255.0
    v4_gw: 10.39.50.1

lan2:

    nic_bonding_enabled: Y
    bonding_failback_interface: lan1

mac:

    mgmt: fa:16:3e:14:3a:ae
    lan1: fa:16:3e:01:29:0b
    ha: fa:16:3e:25:43:8a
    lan2: fa:16:3e:8e:26:4c

```

Sample Commands for Provisioning an HA Pair

Following are sample nova boot commands for deploying an HA pair in the OpenStack environment: Execute the following command to spin up the active node:

```

nova boot --config-drive True --image nios-7.3.0-305525-1420-160.qcow2 --flavor vnios1420.160 --nic
net-id=776f3ea4-9412-464a-b923-cf0c79e579f0 --nic
port-id=81345eaf-49dd-4ad1-a31c-32159ef4d948 --nic
port-id=99c73fc5-71ca-457a-a505-f40e85f88207 --user-data ./user-data-ha HA_active1
Once the active node is up and running, change the token value in the user-data file and executed
following command to create the passive node:
nova boot --config-drive True --image nios-7.3.0-Alpha-305525-1420-160.qcow2 --flavor vnios1420.160
--nic net-id=776f3ea4-9412-464a-b923-cf0c79e579f0 --nic
port-id=ddd57167-2173-4f11-860b-c4efb9ddd3d6 --nic
port-id=07b9df9b-0bb3-4ddd-9ede-850bf3a27671 --user-data ./user-data-ha HA_passive1

```

Setting Up a Grid

An Infoblox Grid is a group of two or more NIOS and vNIOS virtual appliances that share sections of a common, distributed, built-in database and which you configure and monitor through a single, secure point of access: the Grid Master. A Grid consists of a master and at least one member. A Grid member can be a single appliance or an HA pair. For information about the Infoblox Grid and HA pairs, refer to the *Infoblox NIOS Documentation*.

To create a Grid, you must first set up a Grid Master and then add members. In a Grid, you can configure vNIOS virtual appliances as Grid Masters, Grid Master candidates, and Grid members (or a reporting member), depending on the vNIOS appliance specifications. Note that some of the vNIOS for KVM and OpenStack virtual appliances are supported as Grid members only. See [Table 1.1](#) for information about supported appliance models and their specifications.

To set up a Grid:

1. Configure the Grid Master. You can configure a single Grid Master or two vNIOS appliances as an HA Grid Master. For information, see [Configuring vNIOS Appliances as Grid Masters](#).
2. Provision Grid members on the Grid Master. Define Grid member settings on the Grid Master before you join the members to the Grid. For information, see [Provisioning vNIOS Members on the Grid Master](#).
3. Join members to the Grid. For information, see [Configuring and Joining vNIOS Grid Members](#).

Configuring vNIOS Appliances as Grid Masters

After you deploy a vNIOS appliance for KVM, you can configure it as a single or an HA Grid Master. To configure a vNIOS HA Grid Master, deploy two vNIOS appliances and define the network settings for each node. The procedure is the same as joining two physical appliances as an HA pair. You must configure a Grid Master and set up the Grid before you join Grid members.

You configure the vNIOS appliance as a Grid Master using the Infoblox Grid Manager. If you are configuring a vNIOS HA Grid Master, complete the configuration for node 1 as described below. To configure node 2, see [Configuring Node 2 for HA Grid Master](#).

To configure the single Grid Master or node 1 of the HA Grid Master:

1. On your management system, open a new browser window, and then connect to `https://ip_addr`, where `ip_addr` is the address of the single appliance or LAN1 port on node 1.
2. Log in to the Infoblox Grid Manager using the default user name **admin** and password **infoblox**.
3. Review the End-User License Agreement and click **I Accept**.
4. In the *Grid Setup* wizard, select **Configure a Grid Master** and click **Next**.
5. Complete the following to specify the Grid properties, and then click **Next**:
 - **Grid Name:** Enter a text string, such as **DaveyJones**, that the Grid Master and appliances joining the Grid use to authenticate each other when establishing a VPN tunnel between them. The default Grid name is **Infoblox**. Note that the VPN MTU (maximum transmission unit) size for any virtual appliance using OpenStack must be 1500 bytes to avoid any inconsistencies in the instance. For more information about MTU, see [Setting the MTU for VPN Tunnels](#).
 - **Shared Secret:** Enter a text string, such as **L0ck37**, that the Grid Master and appliances joining the Grid use as a shared secret to authenticate each other when establishing a VPN tunnel between them. The default shared secret is **test**.
 - **Show Password:** Select this to display the password. Clear the check box to conceal the password.
 - **Host name:** Enter a valid domain name for the appliance. You can use the name that you entered for the vNIOS appliance when you deployed it.
 - **Is the Grid Master an HA pair?:** Select **No** for the single Grid Master. Select **Yes** for an HA pair.
1. Complete the following to configure the network settings, and then click **Next**:
 - **HostName:** Enter a valid domain name for the appliance.
 - **IP Address:** Displays the IP address of the LAN port.
 - **Subnet Mask:** Displays the subnet mask of the LAN port.
 - **Gateway:** Displays the IP address of the gateway of the subnet on which the LAN port is set.
 - **Port Settings:** The default is **Automatic**. You cannot change port settings for vNIOS appliances.
2. For an HA pair, complete the following to specify the network properties and click **Next**:
 - **VirtualRouterID:** Enter the VRID (virtual router ID). This must be a unique VRID number—from 1 to 255—for this subnet.
 - **RequiredPortsandAddresses:** Enter information about the following virtual interfaces: VIP, Node 1 HA and LAN ports, Node 2 HA and LAN ports. The VIP address and the IP addresses for all the ports must be in the same subnet. Enter the IP address of the gateway for the subnet on which the interfaces are set. This is the same for all interfaces. All fields are required. Note that you cannot change the port settings.
3. Optionally, enter a new password and click **Next**. The password must be a single hexadecimal string (no spaces) that is at least four characters long.
4. Select the time zone of the Grid Master and indicate whether the Grid Master synchronizes its time with an NTP (Network Time Protocol) server, and then click **Next**.
 - If you choose to enable NTP, click the Add icon and enter the IP address of an NTP server. You can enter IP addresses for multiple NTP servers.
 - If you choose to disable NTP, set the date and time for the appliance.
5. The last screen displays the settings you specified in the previous panels of the wizard. Verify that the information is correct and click **Finish**. The application restarts after you click **Finish**.

Note: The *GridSetup* wizard provides options such as not changing the default password and manually entering the time and date. However, changing the password and using an NTP server improve security and accuracy (respectively), and so these choices are presented here.

Record and retain this information in a safe place. If you forget the shared secret, you need to contact Infoblox Technical Support for help. When you add an appliance to the Grid, you must configure it with the same Grid name, shared secret, and VPN port number that you configure on the Grid Master.

Configuring Node 2 for HA Grid Master

For an HA pair, complete the following to configure node 2:

1. On your management system, open a new browser window, and connect to `https://ip_addr`, where `ip_addr` is the address of the LAN1 port on node 2.
2. Log in to the Infoblox Grid Manager using the default user name and password **admin** and **infoblox**.
3. Review the End-User License Agreement and click **I Accept**.
4. In the *Grid Setup* wizard, select **Join Existing Grid** and click **Next**.
5. Complete the following to specify the Grid properties and click **Next**.
 - **Grid Name:** Enter the Grid name you entered for node 1.
 - **Grid Master's IP Address:** Enter the VIP you entered for node 1.
 - **Shared Secret:** Enter the shared secret you entered for node 1.

6. Verify the IP address settings of the member and click **Next**.
7. The last screen displays the settings you specified in the previous panels of the wizard. Verify that the information is correct and click **Finish**.

The setup of the HA Grid Master is complete. If the two nodes cannot join (it should not take more than a few seconds), check the IP addresses of Node 1 LAN and Node 1 HA (the Grid Master) and for Node 2 LAN and Node 2 HA (the node attempting to join the Grid Master to form the HA Pair). Ensure that the network IP address of node 2 is set to the same value as Node 2 LAN on the Grid Master.

As a method of verifying successful communication, open the console window for node 2. You should see a pair of messages as follows:

```
Contacting the Grid Master at 10.36.0.200....  
Synchronizing database with the Grid Master....
```

For more information about HA pair configurations, refer to the *Infoblox NIOS Documentation*.

Configuring vNIOS Appliances as Grid Members

You can configure a vNIOS appliance as a single Grid member, a reporting member, or two vNIOS appliances as a vNIOS HA Grid member. To configure a vNIOS HA Grid member, deploy two vNIOS appliances and define the network settings for each node. Connect to the Grid Master and specify the two vNIOS appliances as nodes in the HA pair. The procedure is the same as joining two physical appliances as an HA pair. You must configure a Grid Master and set up the Grid before you join Grid members. For information, see [Configuring vNIOS Appliances as Grid Masters](#).

To configure a vNIOS appliance as a Grid member or a reporting member:

1. Define the vNIOS appliance on the Grid Master, as described in [Provisioning vNIOS Members on the Grid Master](#).
2. Specify the initial settings and join the vNIOS appliance to the Grid, as described in [Configuring and Joining vNIOS Grid Members](#).

Provisioning vNIOS Members on the Grid Master

Before you configure the individual appliances that you want to add to the Grid, you must first define them on the Grid Master, as follows:

1. Log in to the Grid Master.
2. From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab, and then click **Add** -> **Add Grid Member** from the Toolbar.
3. In the *Add Grid Member* wizard, enter the following and click **Next**:
 - **Member Type**: Select **Virtual NIOS**.
 - **Host Name**: Type the FQDN (fully qualified domain name) of the vNIOS single or HA appliance that you want to add to the Grid.
 - **Time Zone**: If the vNIOS Grid member is in a different time zone from the Grid, click **Override** and select a time zone.
 - **Comment**: Enter useful information about the vNIOS appliance.
4. Enter the following information about the member that you want to add to the Grid and click **Next**:
 - For a single Grid Member:
 - **Standalone Member**: Select this option.
 - **Address**: Type the IP address of the vNIOS Grid member.
 - **Subnet Mask**: Choose the netmask.
 - **Gateway**: Type the IP address of the default gateway of the vNIOS Grid member.
 - **Port Settings**: The default value is **Automatic**. You cannot change port settings for vNIOS appliances.
 - For an HA Grid member:
 - **High Availability Pair**: Select this option.
 - **Virtual Router ID**: Enter a unique VRID number—from 1 to 255—for the local subnet.
 - **Required Ports and Addresses**: Enter information about the following virtual interfaces: VIP, Node 1 HA and LAN ports, Node 2 HA and LAN ports. The VIP address and the IP addresses for all the ports must be in the same subnet. Enter the IP address of the gateway for the subnet on which the interfaces are set. This is the same for all interfaces. All fields are required. Note that you cannot change the port settings.
5. Optionally, define extensible attributes. For information, refer to the *Infoblox NIOS Documentation*.
6. Save the configuration and click **Restart** if it appears at the top of the screen.

Configuring and Joining vNIOS Grid Members

After you successfully install the vNIOS virtual appliance and start the vNIOS appliance, connect to the NIOS CLI and specify the initial settings. If you are configuring a vNIOS HA Grid member, you must complete the following steps for each virtual node in the HA pair.

1. Connect to the Grid Master where you can add the vNIOS appliance to the Grid.
2. From the KVM Hypervisor, select the vNIOS instance.
3. Select the **Console** tab.
4. Click anywhere in the console screen to activate the console.
5. When the Infoblox login prompt appears, log in with the default user name and password.

```
login: admin
password: infoblox
```

The Infoblox prompt appears: `Infoblox >`
6. You must have valid licenses before you can configure the vNIOS appliance. To obtain permanent licenses, first use the **show version** command to obtain the serial number of the vNIOS appliance, and then visit the Infoblox Support web site at <https://support.infoblox.com>. Log in with the user ID and password you receive when you register your product online at: <http://www.infoblox.com/support/customer/evaluation-and-registration>. If the vNIOS appliance does not have the Infoblox licenses required to run NIOS services and to join a Grid, you can use the `set temp_license` command to generate and install a temporary 60-day license. The appliance lists the available licenses.
7. From the list of licenses, select the Grid, vNIOS, and other relevant licenses for your vNIOS virtual appliance. For the vNIOS reporting appliance, you must also select the **Reporting** license.

Note

You must have both the **Grid** and **vNIOS** licenses for the vNIOS appliance to join a Grid.

8. Set the network settings and join the vNIOS appliance to the Grid. Use the CLI command **set network** to configure the network settings and specify the Grid.

Infoblox > **set network**

NOTICE: All HA configurations are performed from the GUI. This interface is used only to configure a standalone node or to join a Grid.

Enter IP address: **10.1.1.11**

Enter netmask: [Default: 255.255.255.0]: **255.255.255.0**

Enter gateway address [Default: 10.1.1.1]: **10.1.1.1**

Become Grid member? (y or n): **y**

Enter Grid Master VIP: **10.1.1.22**

Enter Grid Shared Secret: **L0ck37**

Join Grid as member with attributes:

Join Grid Master VIP: **10.1.1.22**

Grid Name: **DaveyJones**

Grid Shared Secret: **L0ck37**

WARNING: Joining a Grid will replace all the data on this node!

Is this correct? (y or n): **y**

Are you sure? (y or n): **y**

The network settings have been updated.

Verifying and Monitoring

After you configure the vNIOS appliance, you can check its status on the Dashboard and in the **Grid -> Grid Manager -> Members** tab through Grid Manager, as shown in [Figure 1.1](#) and [Figure 1.2](#). For information about Grid Manager, refer to the *Infoblox NIOS Documentation*.

Figure 1.1 vNIOS Appliance Status on the Dashboard

The screenshot shows a window titled "Grid Status" with a sub-header "Infoblox" and a green status indicator. Below this, several services are listed with their status: DHCP (grey), DNS (green), TFTP (grey), HTTP (File Dist) (grey), FTP (grey), NTP (grey), and bloxTools (grey). A table below lists the members:

Member Name	IP Address	Status
infoblox.localdomain	10.34.10.41	Running
infoblox-vn-1050a.infoblox.com	10.39.6.60	Running

Figure 1.2 vNIOS Appliance Status in the Members Tab

The screenshot shows the "Members" tab in Grid Manager. At the top, there are tabs for "Members" and "Services". Below the tabs, there are controls for "Toggle Replication Status View" (Off), "Toggle Filter On" (Off), and "Show Filter". A toolbar with various icons is visible. The main content is a table with the following data:

Name	Status	IP Address	DHCP	DNS	TFTP	HTTP
ib-10-35-3-110	Running	10.35.3.110	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
vnios5.com	Running	10.39.7.106	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

Auto Scaling for Virtual DNS Cache Acceleration

Auto Scaling contains OpenStack packages that allow DNS Auto Scaling for virtual DNS cache acceleration using a NIOS Grid. Auto Scaling helps you ensure that you have required number of resources to handle the load in your application. You can create an Auto Scaling group and specify the minimum, maximum and desired number of resources for each group. Auto Scaling ensures that the group contains the required number of resources at all times, neither exceed the maximum limit nor fall short of resources. With Auto Scaling, you can adjust scaling to best meet the needs of your applications by automatically increasing or decreasing the computing capacity of the associated application. For more information about virtual DNS cache acceleration, refer to [Configuring DNS Cache Acceleration on IB-FLEX](#).

Note

Infoblox supports Auto Scaling for OpenStack only.

Infoblox supports Auto Scaling on the IB-FLEX and the following vSOT platforms: IB-V815, IB-V825, IB-V1415, IB-V1425, IB-V2215, IB-V2225, IB-V4015, and IB-V4025. You can install Auto Scaling on both SRIOV and Non-SRIOV servers, but only on the IPv4 interfaces. You can create an Auto Scaling group using the orchestration component, **Standard OpenStack Heat**, and specify the type of resource that must be scaled. You can also define policies that indicate when and how to scale the resource. For example, you can define an Auto Scale group of server resources, and configure it to trigger the new server instance using OpenStack when the aggregate CPU utilization of the entire group exceeds the CPU utilization threshold for a specified period of time. You can also specify to halt the number of servers if the CPU utilization has been low for a longer period. For more information about OpenStack heat, refer to <https://github.com/infobloxopen/heat-infoblox>. For virtual DNS cache acceleration, NIOS Grid members are triggered using OpenStack based on the queries per second. The new member acts as a secondary DNS. Note that Auto Scaling for virtual DNS cache acceleration feature consists of two components, **heat-infoblox** and **ceilometer-infoblox**, which are Python packages that must be installed with **pip**. For more information, refer to <https://github.com/infobloxopen/infoblox-netmri>.

The **heat-infoblox** package contains **Heat resources classes** for Grid members and name server group entries, along with the supporting code. The **Heat resources classes** enable you to add and remove Grid members, enable or disable DNS service for the Grid members, and add the Grid members as secondary servers in a name server group or remove the Grid members. Note that all these are orchestrated through the heat engine.

Note

You must restart the Heat engine after you install and configure the package.

The **ceilometer-infoblox** package contains code that enables SNMP polling of the Infoblox NIOS instances within OpenStack to gather DNS queries per second. After installing the **OpenStack Ceilometer**, you must install the **ceilometer-infoblox** package on each compute node. For more information about the configuration details, refer to <https://github.com/infobloxopen/ceilometer-infoblox>.

Note

You must configure OpenStack Ceilometer in the compute node and restart the compute polling agent.

Configuring NIOS Tenant

The templates directory within the `heat-infoblox/doc/templates/` package contains a `setup.sh` script that creates a tenant, user and appropriate images.

To configure a NIOS tenant:

1. Log in to the OpenStack **admin** user account, copy the images to the `heat-infoblox/doc/templates/directory` and execute the `setup.sh` script.
2. Next, switch to the OpenStack **nios** user account and launch a heat stack using the following command:

```
heat stack-create -f member-server.yaml -P"mgmt_network=management-net;lan1_network=service-net" member-server
```
3. To view the progress, log in to the **Orchestration -> Resource Types** section of the OpenStack Horizon UI or through the heat CLI utility.
4. Execute the `member.yaml` script to create a member in NIOS. Note that this script neither creates a server nor does it add the member to the name server group.
5. Execute the `member-server.yaml` script. This script creates a member in NIOS, adds the member to the name server group, and launches a Nova server and pushes parameters to the new server, so that it can automatically join the Grid.

Configuring NIOS Grid Master

To configure the Grid Master:

1. Configure a default name server group with the Grid Master as primary and name it as **default**.

2. Generate license pools for all the required licenses.
3. Get the Grid Master certificate and save it in `member-server.yaml`:

```
echo | openssls_client -connect gm-ip:443 2>/dev/null | openssl x509
```
4. Configure SNMP and set the community string to **public**.

If you set up a management interface on the Grid Master, then you must add an interface on the **router** and connect it to **management**. You must also provide the MGMT IP address in `member-server.yaml`.

Configuring Auto Scale

Auto Scaling uses groups and members automatically to scale up or scale down the resources based on the QPS alarm. Based on the threshold and the period that is defined in the respective QPS alarm, corresponding QPS alarm is generated either to scale up or scale down the resources. For example, if the `threshold` for `qps_alarm_high` is set to 5000 and the period is set to 120, then the `qps_alarm_high` triggers `scaleup_policy` when the threshold is more than 5000 continuously. The server adds the number of members to the Grid based on the value set for `scaling_adjustment` in the `scaleup_policy`. If the `scaling_adjustment` is set to 2, the server adds two members. Similarly, for `qps_alarm_low` the number of members are reduced from the Grid.

You must define the LAN1 port and Grid member details in the `autoscale-member.yaml`. The URL and the certificate mentioned in `autoscale-member.yaml` is used to join the member to the Grid. This also contains the name of the group to which the member belongs. Note that the queries generated for the respective member are directed through the Anycast Loopback address, so the server knows when to scale up when the load increases.

To install the vNIOS software package and configure Auto Scale in the OpenStack environment, complete the following:

1. Download the OpenStack server. Example: 10.39.50.13 (non-SRIOV server).
2. Install relevant package for the operating system, either CentOS for KVM or RHEL (Red Hat Enterprise Linux) for KVM-based OpenStack, that you use.
3. Download the `qcow2` file from the Infoblox Technical Support site. For information, see [Requirements](#). Upload the `qcow2` image file in to the OpenStack server. Example: `DCA_354869`.
4. Log in using the command `ssh root@openstack server ip`. Example: `ssh root@10.39.50.13`.
5. Ensure that the files mentioned below exist in the `heat-infoblox/doc/templates/` directory:
 - `[root@rhel72-10-39-50-13 ~]# source keystonerc_admin`
 - `[root@rhel72-10-39-50-13 ~(keystone_admin)]# cd /opt/templates/`
 - `config-gm.sh`
 - `autoscale-member.yaml`
 - `autoscale.yaml`
 - `gm.yaml`
6. Create and configure the Grid Master using the following command:

```
root@rhel72-10-39-50-13 templates(keystone_admin)]#openstack stack create -f yaml -t gm.yaml
--parameter "imageName=DCA_354869" GMaster
```

Ensure that you update the `imageName`.

7. Execute `config-gm.sh`. It automatically uses the floating IP or the LAN1 IP for SRIOV to generate the certificate. It then starts the DNS service, adds required records, like FQDN and A records, to the zones, and creates name server groups to which the member belongs. The following command generates an environment file `gm-10.39.52.162-env.yaml` where 10.39.52.162 is the floating IP of the Auto Scale member:

```
[root@rhel72-10-39-50-13 templates(keystone_admin)]# ./config-gm.sh 10.39.52.162
```

8. Next, create an Auto Scale stack and launch the Auto Scale using the following command:

```
[root@rhel72-10-39-50-13 templates(keystone_admin)]#openstack stack create -e
gm-10.39.52.162-env.yaml -f yaml -t autoscale.yaml autoscale
```

The IP address 10.39.52.162 is the IP address of the Grid Master that joins the Auto Scale member to the Grid.

Note that after you create the Auto Scale stack that uses `autoscale.yaml`, which in turn uses `autoscale-member.yaml` to direct queries when the query rate is more than the `threshold` and scale up resources.

Known Limitations

vNIOS for KVM supports most of the features of the Infoblox NIOS appliances, with the following limitations:

- vNIOS for KVM does not support the following features:
 - Configuration of port settings for MGMT, LAN, LAN2, and HA ports
 - The bloxTools environment
- When you configure an HA pair, both nodes in the HA pair must be vNIOS instances. You cannot configure a physical NIOS appliance and a vNIOS instance in an HA pair.
- vNIOS appliances run on virtual hardware. They do not have sensors to monitor the physical CPU temperature, fan speed, and system temperature.
- Changing the vNIOS appliance settings through KVM may violate the terms of the vNIOS licensing and support models. The vNIOS appliance may not join the Grid or function properly.

The following known issues are specific to vNIOS for KVM deployed in the KVM-only environment:

- A vNIOS instance may fail to start if it is deployed in KVM-only environment with Linux bridged networking enabled. You may need to modify certain files to fit your environment.
- If you get the message `error: Unable to read from monitor: Connection reset by peer` when starting a vNIOS instance in KVM hypervisor, check the memory of the hypervisor. Failure to start an instance may be due to the lack of memory.

The following known issues are specific to vNIOS for KVM deployed in the OpenStack environment:

- If the vNIOS instance fails to launch, the ports created for the instance are not deleted automatically. You need to delete them manually.
- When you start an FTP connection on OpenStack with the **Listen on port** field set to 2021, you need to manually add a rule to allow the 2021 port. This ensures that the connection is successful.
- Connection to the FTP service might fail when the virtual appliance enter the passive mode. To avoid this, do the following:
 - Use active mode instead of passive mode.
 - Modify the vnios-sec-group security group to open ports 1023 and above.
 - Use the FTP client inside the internal network.
- The vNIOS instances deployed in the OpenStack environment do not support HSM Safenet groups.
- IPv6 support in the OpenStack environment is limited by the Juno release and has not been verified by Infoblox.
- You cannot configure the reporting virtual appliance as an HA pair. You also cannot configure it as a Grid Master or Grid Master Candidate. You can use it only as a dedicated reporting server in the Grid.