



OpenStack Administration Guide

FortiOS 8.0



FORTINET DOCUMENT LIBRARY

<https://docs.fortinet.com>

FORTINET VIDEO LIBRARY

<https://video.fortinet.com>

FORTINET BLOG

<https://blog.fortinet.com>

CUSTOMER SERVICE & SUPPORT

<https://support.fortinet.com>

FORTINET TRAINING & CERTIFICATION PROGRAM

<https://www.fortinet.com/training-certification>

FORTINET TRAINING INSTITUTE

<https://training.fortinet.com>

FORTIGUARD LABS

<https://www.fortiguard.com>

END USER LICENSE AGREEMENT

<https://www.fortinet.com/doc/legal/EULA.pdf>

FEEDBACK

Email: techdoc@fortinet.com



April 21, 2026

FortiOS 8.0 OpenStack Administration Guide

01-80-1054263-20260421

TABLE OF CONTENTS

About FortiGate-VM on OpenStack	4
Preparing for deployment	4
Virtual environment	4
Management software	5
Connectivity	5
Configuring resources	5
Registering the FortiGate-VM	5
Downloading the FortiGate-VM virtual appliance deployment package	6
Deployment package contents for OpenStack	6
Deploying a FortiGate-VM instance in an OpenStack environment	7
Setting up networks in OpenStack	7
Deploying a FortiGate-VM instance into the configured networks	8
Creating a user_data file to pre-configure a FortiGate-VM instance	8
Disabling port security for the FortiGate-VM and CirrOS instances	9
Setting up the FortiGate-VM network configuration	11
Verifying internet access	11
Deploying two FortiGate-VM instances in an HA configuration in an OpenStack environment	12
Setting up networks in OpenStack	13
Deploying two FortiGate-VMs into the configured networks	14
Creating a user_data file to preconfigure FortiGate-VM instances	14
Disabling port security for the FortiGate-VM and CirrOS instances	16
Setting up the FortiGate-VM HA configuration	18
Completing the FortiGate-VM network configuration	18
Testing HA operation and failover	19
Verifying HA cluster status	20
Deploying a FortiGate-VM instance in an OpenStack environment using service insertion or chaining	22
Sample configuration	22
Optimizing FortiGate-VM performance	24
SR-IOV	24
FortiGate-VM interrupt affinity	27
FortiGate-VM affinity packet redistribution	28
FortiGate-VM MTU setting	29
Automatically updating dynamic addresses using an SDN connector	30
Configuring the OpenStack environment	30
Troubleshooting OpenStack Horizon SDN connector	32
Configuring OpenStack SDN connector with domain filter	32
Change log	35

About FortiGate-VM on OpenStack

FortiGate virtual appliances allow you to mitigate blind spots by implementing critical security controls within your virtual infrastructure. They also allow you to rapidly provision security infrastructure whenever and wherever it is needed. FortiGate virtual appliances feature all the security and networking services common to hardware-based FortiGate appliances. You can deploy a mix of FortiGate hardware and virtual appliances, operating together and managed from a common centralized management platform.

OpenStack-based clouds provide the environment needed for elastic on-demand multitenant applications. Networks are transitioning to new models more suited to the cloud with software-defined networking, NFV, and virtual network infrastructure, and their relationships between networking, security orchestration, and policy enforcement.

Fortinet's OpenStack Neutron solution embraces the software-defined security framework providing out-of-the-box integration so that you can apply advanced network security seamlessly in logical and dynamic environments.

This guide describes how to deploy a FortiGate virtual appliance in an OpenStack environment. You can install FGT-VM64-KVM and FOS-VM64-KVM firmware into an OpenStack environment.

The following OpenStack releases support FortiGate-VM virtualized network function deployment:

- Rocky
- Stein
- Train
- Ussuri
- Victoria
- Wallaby
- Xena

Preparing for deployment

This guide assumes that before deploying the FortiGate-VM virtual appliance on the OpenStack virtual platform, you have addressed the following requirements:

Virtual environment

The OpenStack software is installed on a physical server with sufficient resources to support the FortiGate-VM and all other VMs that you deploy on the platform.

If you configure the FortiGate-VM to operate in transparent mode, ensure that the OpenStack environment includes virtual switches configured to support the FortiGate-VM's operation before you create the FortiGate-VM.

If you set up multiple FortiGate-VMs in a FortiGate clustering protocol high availability (HA) cluster, do the following:

- Purchase identical FGT-VM64-KVM or FOS-VM64-KVM licenses
- Be prepared to set up a dedicated network in the OpenStack environment for the HA heartbeat. See [Verifying HA cluster status on page 20](#).

Management software

The VMware management software, OpenStack Horizon, is installed on a computer with network access to the OpenStack server.

Connectivity

The FortiGate-VM requires an Internet connection to contact FortiGuard to validate its license. If the FortiGate-VM is in a closed environment, it must be able to connect to a FortiManager to validate the FortiGate-VM license.

Configuring resources

Before you start the FortiGate-VM for the first time, configure the following resources as the FortiGate-VM license specifies:

- Disk sizes
- CPUs
- RAM: FortiOS does not have licensed RAM size restrictions. Having at least 4 GB of RAM for proper FortiGate-VM operation is recommended, especially if unified threat management, zero trust network access, or proxy is enabled.
- Network settings

To configure the resources for a FortiGate-VM deployed on OpenStack, use the OpenStack Horizon client.

You can apply a smaller FortiGate-VM license if you are OK with consuming less CPU than is present on your instance. See [FortiGate-VM virtual licenses and resources](#).

Registering the FortiGate-VM

Registering the FortiGate-VM with [Customer Service & Support](#) allows you to obtain the FortiGate-VM license file.

To register the FortiGate-VM:

1. Log in to the [Customer Service & Support site](#) using a support account, or select *REGISTER* to create an account.
2. In the main page, under *Asset*, select *Register/Activate*.
3. In the *Registration* page, enter the registration code that you received via email, and select *Register* to access the registration form.
4. Complete and submit the registration form.
5. In the registration acknowledgment page, click the *License File Download* link.
6. Save the license file (.lic) to your local computer.

Downloading the FortiGate-VM virtual appliance deployment package

You can find FortiGate-VM deployment packages on the [Customer Service & Support site](#).

To download the FortiGate-VM virtual appliance deployment package:

1. Log into the [Customer Service & Support site](#).
2. From the *Support* dropdown list, select *VM Images* to access the available VM deployment packages.
3. From the *Select Product* dropdown list, select *FortiGate*.
4. From the *Select Platform* dropdown list, select *KVM*.
5. Select the desired FortiOS version. There are files available for download: files required to upgrade from an earlier version and files required for a new deployment.
6. Click the *Download* and save the file.

See the [FortiGate Virtual Appliances datasheet](#).

You can also download the following resources for the firmware version:

- [FortiOS Release Notes](#)
- FORTINET-FORTIGATE MIB file
- FSSO images
- SSL VPN client

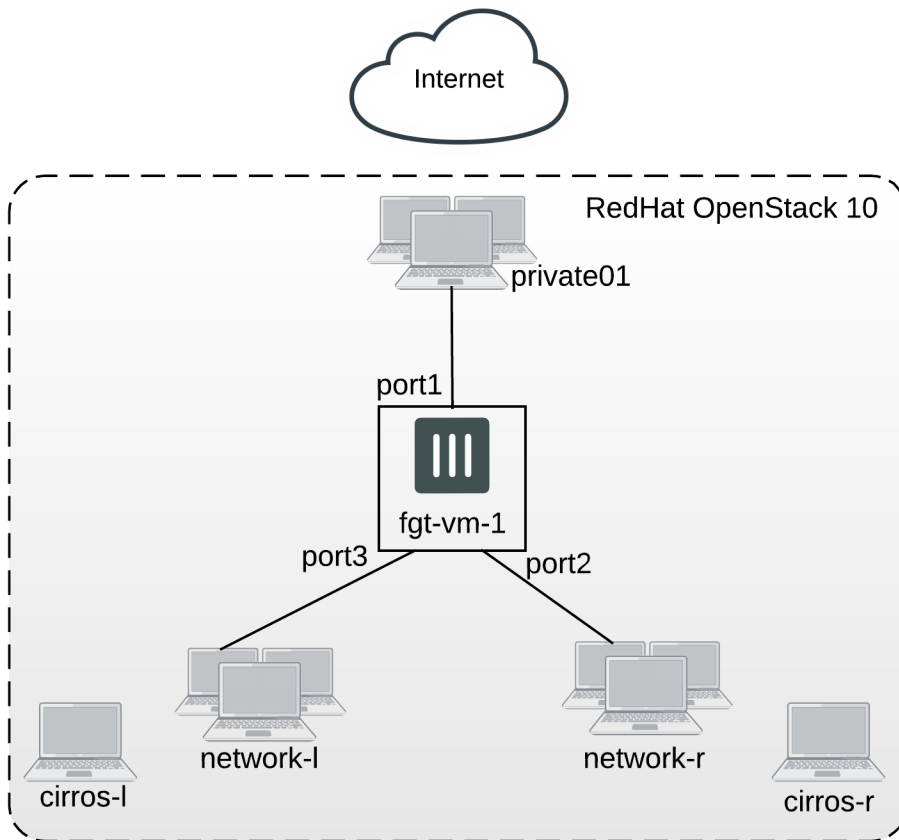
Deployment package contents for OpenStack

The FORTINET.out.kvm.zip contains only fortios.qcow2, the FortiGate-VM system hard disk in qcow2 format. You must manually:

- Create a 32 GB log disk
- Specify the virtual hardware settings

Deploying a FortiGate-VM instance in an OpenStack environment

This example shows how to set up FortiGate-VM instance in an OpenStack 10 environment. The FortiGate-VM instance is connected to a private network (private01) and protects two networks (network-l and network-r). Each network includes a Cirros instance (cirros-l and cirros-r) for testing.



Setting up networks in OpenStack

From the OpenStack environment command line, enter the following commands to create network-r and network-l:

```
$ source overcloudrc_tenant01
$ openstack network create network-r
$ openstack subnet create subnet-r --network network-r --subnet-range 172.32.0.0/24 --dns-nameserver
  208.91.112.53
$ openstack network create network-l
```

```
$ openstack subnet create subnet-1 --network network-1 --subnet-range 172.33.0.0/24 --dns-nameserver 208.91.112.53
```

Add the CirrOS instances to network-r and network-l:

```
$ openstack server create --flavor m1.tiny --image cirros035 --security-group web --nic net-id=network-r cirros-r
$ openstack server create --flavor m1.tiny --image cirros035 --security-group web --nic net-id=network-l cirros-l
```

Deploying a FortiGate-VM instance into the configured networks

From the OpenStack command line, enter the following commands to deploy the fgt-vm-1 FortiGate-VM instance. These commands use the standard license file you receive when you register your FortiGate-VM instance (in this example, FGVM080000103268.lic):

```
$ openstack server create --flavor m1.fortigate --image fgtb1486 --user-data /home/stack/openstack/cloud-init/user_data --config-drive=true --file license=/home/stack/FG-licenses/FGVM080000103268.lic --security-group web --nic net-id=private01 --nic net-id=network-r --nic net-id=network-l --nic net-id=ha-sync fgt-vm-1
```

Creating a user_data file to pre-configure a FortiGate-VM instance

The following example user_data file sets up a FortiGate-VM instance (fgt-vm-1) with a basic default configuration customized for your environment and requirements. This example configures interfaces, adds a DNS server, and configures two firewall policies that allow devices in network-l and network-r to access the private01 network and the internet through the private01 network.

The following example user_data file could be used for fgt-vm-1:

```
#FGT VM Config File
config sys global
set hostname fgt-vm-1
end
config system interface
edit port1
set mode dhcp
set allowaccess http https ssh ping
next
edit port2
set mode dhcp
set defaultgw disable
set allowaccess http https ssh ping
next
edit port3
set mode dhcp
```

```
set defaultgw disable
set allowaccess http https ssh ping
next
end
config system dns
set primary 208.91.112.53
end
config firewall policy
edit 1
set name "network-l internet access"
set dstintf "port3"
set srcintf "port1"
set srcaddr "all"
set dstaddr "all"
set action accept
set schedule "always"
set service "ALL"
set nat enable
next
edit 2
set name "network-r internet access"
set dstintf "port2"
set srcintf "port1"
set srcaddr "all"
set dstaddr "all"
set action accept
set schedule "always"
set service "ALL"
set nat enable
end
config system central-management
set include-default-servers disable
set type fortimanager
set fmg 10.210.8.25
config server-list
edit 1
set server-type update rating
set server-address 10.210.8.25
end
end
```

Disabling port security for the FortiGate-VM and CirrOS instances

In OpenStack, the networking component (called Neutron) only allows traffic with known IP/MAC address combinations. This makes the network very secure. However, normal firewall traffic contains very many IP/MAC address combinations, and it is not practical to add them all to the configuration. Instead, to allow normal firewall traffic, you must disable port security for your FortiGate-VM instance. See [Managing port level security in OpenStack](#).

Use the Horizon *Instances* view to verify the IP addresses of the FortiGate-VM instance, the CirrOS instances, and the networks that the interfaces are connected to. For example:

Instances

Instance Name	Image Name	IP Address
		private01 • 172.31.0.10 Floating IPs: • 10.210.9.10
fgt-vm-1	fgtb1486	network-r • 172.32.0.11 network-l • 172.33.0.5
cirros-l	cirros035	• 172.33.0.9
cirros-r	cirros035	• 172.32.0.12

From the OpenStack command line, run the following bash script to disable port security on the FortiGate-VM interfaces:

```
#!/bin/bash
echo
echo 'Disable port_security on fgt-vm-1'
echo
echo
`source /home/stack/overcloudrc_tenant01`
FGT='fgt-vm-1'
IPADDR=`openstack server show $FGT | grep addresses | awk -F "|" '{print $3}' | awk -F "=" '{print $2}' | awk -F ";" '{print $1}'`
PORTID=`openstack port list | grep $IPADDR | awk -F "|" '{print $2}'`
neutron port-update $PORTID --no-security-groups --port_security_enabled=False
echo
echo $IPADDR
echo `openstack port show $PORTID`
echo
IPADDR=`openstack server show $FGT | grep addresses | awk -F "|" '{print $3}' | awk -F "=" '{print $3}' | awk -F ";" '{print $1}'`
PORTID=`openstack port list | grep $IPADDR | awk -F "|" '{print $2}'`
`neutron port-update $PORTID --no-security-groups --port_security_enabled=False`
echo
echo $IPADDR
echo `openstack port show $PORTID`
echo
IPADDR=`openstack server show $FGT | grep addresses | awk -F "|" '{print $3}' | awk -F "=" '{print $4}' | awk -F ";" '{print $1}'`
PORTID=`openstack port list | grep $IPADDR | awk -F "|" '{print $2}'`
`neutron port-update $PORTID --no-security-groups --port_security_enabled=False`
echo
echo $IPADDR
echo `openstack port show $PORTID`
echo
IPADDR=`openstack server show $FGT | grep addresses | awk -F "|" '{print $3}' | awk -F "=" '{print $5}' | awk -F ";" '{print $1}'`
PORTID=`openstack port list | grep $IPADDR | awk -F "|" '{print $2}'`
`neutron port-update $PORTID --no-security-groups --port_security_enabled=False`
echo
echo $IPADDR
```

```
echo `openstack port show $PORTID`  
echo
```

From the OpenStack command line, associate floating IP addresses to the FortiGate-VM by entering the `openstack server add floating ip fgt-vm-1 10.210.9.10` command.

Setting up the FortiGate-VM network configuration

From the FortiGate-VM instance CLI, enter the following commands to change the FortiGate-VM interfaces from DHCP to static and add IP addresses. The IP addresses assigned to the interfaces must be on the subnets of the networks that the interfaces are connected to.

```
config system interface  
  edit "port1"  
    set mode static  
    set ip 172.31.0.3 255.255.255.0  
    set allowaccess ping https ssh http  
  next  
  edit "port2"  
    set mode static  
    set ip 172.32.0.9 255.255.255.0  
    set allowaccess ping https ssh http  
  next  
  edit "port3"  
    set mode static  
    set ip 172.33.0.4 255.255.255.0  
    set allowaccess ping https ssh http  
end
```

Enter the following command to add a static route:

```
config router static  
  edit 1  
    set gateway 172.31.0.1  
    set device "port1"  
end
```


Verifying internet access

Log into the CirrOS instances on each network, and attempt to access an address on the internet or on the private01 network if the private01 network cannot access the internet. One way to do this is to ping an IP address on the private01 network or on the internet.

Deploying two FortiGate-VM instances in an HA configuration in an OpenStack environment

This example shows how to set up FortiGate clustering protocol high availability (HA) with two FortiGate-VM instances in an OpenStack 10 environment. The FortiGate-VMs are connected to a private network (private01) and protect two networks (network-l and network-r). Each network includes a CirrOS instance (cirros-l and cirros-r) for testing.

To support HA heartbeat communication, the OpenStack environment also includes a network named ha-sync configured with the subnet used by the HA heartbeat interfaces (169.254.0.0/24).



Setting up networks in OpenStack

From the OpenStack environment command line, enter the following commands to create network-r and network-l:

```
$ source overcloudrc_tenant01
$ openstack network create network-r
$ openstack subnet create subnet-r --network network-r --subnet-range 172.32.0.0/24 --dns-nameserver
  208.91.112.53
$ openstack network create network-l
$ openstack subnet create subnet-l --network network-l --subnet-range 172.33.0.0/24 --dns-nameserver
  208.91.112.53
```

Add the CirrOS instances to network-r and network-l:

```
$ openstack server create --flavor m1.tiny --image cirros035 --security-group web --nic net-
  id=network-r cirros-r
$ openstack server create --flavor m1.tiny --image cirros035 --security-group web --nic net-
  id=network-l cirros-l
```

To set up the ha-sync network for the HA heartbeat:

From the OpenStack environment command line, enter the following commands to create the HA-sync network to use for high availability (HA) heartbeat communication.

```
$ openstack network create ha-sync
$ openstack subnet create subnet-ha --network ha-sync --subnet-range 169.254.0.0/24 --dns-nameserver
  208.91.112.53
```

To verify the MTU assigned to the ha-sync network:

You can use the OpenStack Horizon *Networks* view to verify the MTU assigned to the ha-sync network.

The screenshot shows the OpenStack Horizon interface for the 'ha-sync' network. At the top, there is a breadcrumb trail: 'Project / Network / Networks / ha-sync'. Below this, the title 'ha-sync' is displayed. There are three tabs: 'Overview' (selected), 'Subnets', and 'Ports'. Under the 'Overview' tab, the 'Network Overview' section is visible, containing a table of network details.

Name	ha-sync
ID	386bbd55-b8e6-4b60-bd99-7c43e653eab3
Project ID	9c454a837fa347478c8aaafc4417c71d
Status	Active
Admin State	UP
Shared	No
External Network	No
MTU	1446
Provider Network	Network Type: vxlan Physical Network: - Segmentation ID: 34

Deploying two FortiGate-VMs into the configured networks

From the OpenStack command line, enter the following commands to deploy two FortiGate-VM instances (fgt-vm-1 and fgt-vm-2). These commands use the standard license files you receive when you register your FortiGate-VMs (in this example, FGVM080000103268.lic and FGVM080000109643.lic).

```
$ openstack server create --flavor m1.fortigate --image fgtb1486 --user-data
/home/stack/openstack/cloud-init/user_data --config-drive=true --file license=/home/stack/FG-
licenses/FGVM080000103268.lic --security-group web --nic net-id=private01 --nic net-
id=network-r --nic net-id=network-l --nic net-id=ha-sync fgt-vm-1
$ openstack server create --flavor m1.fortigate --image fgtb1486 --user-data
/home/stack/openstack/cloud-init/user_data --config-drive=true --file license=/home/stack/FG-
licenses/FGVM080000109643.lic --security-group web --nic net-id=private01 --nic net-
id=network-r --nic net-id=network-l --nic net-id=ha-sync fgt-vm-2
```

Creating a user_data file to preconfigure FortiGate-VM instances

The following example user_data file sets up a FortiGate-VM instance with a basic default configuration customized for your environment and requirements. This example configures interfaces, and adds a DNS server and two firewall policies that allow any traffic to pass between the port2 and port3 interfaces. These policies make it easier to test high availability (HA) failover.

In addition, the port4 interface MTU is set to be compatible with the OpenStack 10 environment, which by default, has an MTU of 1446. (In the user_data file, the MTU of port4 is set to 1400.) Using the same MTU setting as the OpenStack 10 environment enables the HA heartbeat interfaces to communicate effectively over the ha-sync network.

See the following for information on RedHat OpenStack networks and MTU values:

- [MTU for VLAN networks is by default 1496 Bytes in Red Hat OpenStack Platform 10](#)
- [Configure MTU Settings](#)

You may use the following example user_data file for fgt-vm-1. The user_data file for fgt-vm-2 would be the same except for the hostname:

```
#FGT VM Config File
config sys global
set hostname fgt-vm-1
end
config system interface
edit port1
set mode dhcp
set allowaccess http https ssh ping
next
edit port2
set mode dhcp
set defaultgw disable
```

```
set allowaccess http https ssh ping
next
edit port3
set mode dhcp
set defaultgw disable
set allowaccess http https ssh ping
next
edit port4
set mtu-override enable
set mtu 1400
next
end
config system dns
set primary 208.91.112.53
end
config firewall policy
edit 1
set name "Allow port2 to port3"
set dstintf "port2"
set srcintf "port3"
set srcaddr "all"
set dstaddr "all"
set action accept
set schedule "always"
set service "ALL"
set nat enable
next
edit 2
set name "Allow port3 to port2"
set dstintf "port3"
set srcintf "port2"
set srcaddr "all"
set dstaddr "all"
set action accept
set schedule "always"
set service "ALL"
set nat enable
end
config system central-management
set include-default-servers disable
set type fortimanager
set fmg 10.210.8.25
config server-list
edit 1
set server-type update rating
set server-address 10.210.8.25
end
end
```

Disabling port security for the FortiGate-VM and CirrOS instances

In OpenStack, the networking component (called Neutron) only allows traffic with known IP/MAC address combinations. This secures the network. However, normal firewall traffic contains many IP/MAC address combinations, and adding them all to the configuration is impractical. Instead, to allow normal firewall traffic, you must disable port security for your FortiGate-VM instance. See [Managing port level security in OpenStack](#).

Use the Horizon *Instances* view to verify the IP addresses of the FortiGate-VM instances, the CirrOS instances, and the networks that the interfaces are connected to. The following shows an example:

Instances

Instance Name	Image Name	IP Address
<input type="checkbox"/>		private01 • 172.31.0.6 Floating IPs: • 10.210.9.15
<input type="checkbox"/> fgt-vm-2	fgtb1486	network-r • 172.32.0.5 ha-sync • 169.254.0.13 network-l • 172.33.0.11
<input type="checkbox"/>		private01 • 172.31.0.10 Floating IPs: • 10.210.9.10
<input type="checkbox"/> fgt-vm-1	fgtb1486	network-r • 172.32.0.11 network-l • 172.33.0.5 ha-sync • 169.254.0.10
<input type="checkbox"/> cirros-l	cirros035	• 172.33.0.9
<input type="checkbox"/> cirros-r	cirros035	• 172.32.0.12

From the OpenStack command line, run the following bash script to disable port security on the FortiGate-VM interfaces:

```
#!/bin/bash
echo
echo 'Disable port_security on fgt-vm-1'
echo
echo
`source /home/stack/overcloudrc_tenant01`
FGT='fgt-vm-1'
IPADDR=`openstack server show $FGT | grep addresses | awk -F "|" '{print $3}' | awk -F "=" '{print $2}' | awk -F ";" '{print $1}`
PORTID=`openstack port list | grep $IPADDR | awk -F "|" '{print $2}`
neutron port-update $PORTID --no-security-groups --port_security_enabled=False
echo
echo $IPADDR
echo `openstack port show $PORTID`
```

```
echo
IPADDR=`openstack server show $FGT | grep addresses | awk -F "|" '{print $3}' | awk -F "=" '{print $3}' | awk -F ";" '{print $1}'`
PORTID=`openstack port list | grep $IPADDR | awk -F "|" '{print $2}'`
`neutron port-update $PORTID --no-security-groups --port_security_enabled=False`
echo
echo $IPADDR
echo `openstack port show $PORTID`
echo
IPADDR=`openstack server show $FGT | grep addresses | awk -F "|" '{print $3}' | awk -F "=" '{print $4}' | awk -F ";" '{print $1}'`
PORTID=`openstack port list | grep $IPADDR | awk -F "|" '{print $2}'`
`neutron port-update $PORTID --no-security-groups --port_security_enabled=False`
echo
echo $IPADDR
echo `openstack port show $PORTID`
echo
IPADDR=`openstack server show $FGT | grep addresses | awk -F "|" '{print $3}' | awk -F "=" '{print $5}' | awk -F ";" '{print $1}'`
PORTID=`openstack port list | grep $IPADDR | awk -F "|" '{print $2}'`
`neutron port-update $PORTID --no-security-groups --port_security_enabled=False`
echo
echo $IPADDR
echo `openstack port show $PORTID`
echo
echo 'Disable port-security on fgt-vm-2'
echo
FGT='fgt-vm-2'
IPADDR=`openstack server show $FGT | grep addresses | awk -F "|" '{print $3}' | awk -F "=" '{print $2}' | awk -F ";" '{print $1}'`
PORTID=`openstack port list | grep $IPADDR | awk -F "|" '{print $2}'`
neutron port-update $PORTID --no-security-groups --port_security_enabled=False
echo
echo $IPADDR
echo `openstack port show $PORTID`
echo
IPADDR=`openstack server show $FGT | grep addresses | awk -F "|" '{print $3}' | awk -F "=" '{print $3}' | awk -F ";" '{print $1}'`
PORTID=`openstack port list | grep $IPADDR | awk -F "|" '{print $2}'`
`neutron port-update $PORTID --no-security-groups --port_security_enabled=False`
echo
echo $IPADDR
echo `openstack port show $PORTID`
echo
IPADDR=`openstack server show $FGT | grep addresses | awk -F "|" '{print $3}' | awk -F "=" '{print $4}' | awk -F ";" '{print $1}'`
PORTID=`openstack port list | grep $IPADDR | awk -F "|" '{print $2}'`
`neutron port-update $PORTID --no-security-groups --port_security_enabled=False`
echo
echo $IPADDR
echo `openstack port show $PORTID`
echo
IPADDR=`openstack server show $FGT | grep addresses | awk -F "|" '{print $3}' | awk -F "=" '{print $5}' | awk -F ";" '{print $1}'`
PORTID=`openstack port list | grep $IPADDR | awk -F "|" '{print $2}'`
`neutron port-update $PORTID --no-security-groups --port_security_enabled=False`
echo
echo $IPADDR
```

```
echo `openstack port show $PORTID`  
echo
```

From the OpenStack command line, associate floating IP addresses to the two FortiGate-VMs by entering the following commands:

```
openstack server add floating ip fgt-vm-1 10.210.9.10  
openstack server add floating ip fgt-vm-2 10.210.9.14
```

Setting up the FortiGate-VM HA configuration

From the CLI of each FortiGate-VM instance, configure both FortiGate-VMs for high availability (HA). Both FortiGate-VM instances must have the same HA configuration, for example:

```
config system ha  
  set group-name "group-01"  
  set mode a-p  
  set password <password>  
  set hbdev "port4" 50  
  set override disable  
  set monitor "port2"  
end
```

Completing the FortiGate-VM network configuration

From each FortiGate-VM instance CLI, enter the following commands to change the FortiGate-VM interfaces from DHCP to static, add IP addresses, and add a static route. The IP addresses assigned to the interfaces must be on the subnets of the networks that the interfaces are connected to.

The example shows the fgt-vm-1 configuration. The fgt-vm-2 configuration would be the same except for the interface IP addresses:

```
config system interface  
  edit "port1"  
    set mode static  
    set ip 172.31.0.3 255.255.255.0  
    set allowaccess ping https ssh http  
  next  
  edit "port2"  
    set mode static  
    set ip 172.32.0.9 255.255.255.0  
    set allowaccess ping https ssh http  
  next  
  edit "port3"  
    set mode static  
    set ip 172.33.0.4 255.255.255.0  
    set allowaccess ping https ssh http  
end
```

```
config router static
  edit 1
    set gateway 172.31.0.1
    set device "port1"
end
```

Testing HA operation and failover

This section describes how to verify that a FortiGate-VM high availability (HA) cluster in an OpenStack environment operates normally and fails over successfully.

On the cirros-l instance console (see the diagram in [Deploying two FortiGate-VM instances in an HA configuration in an OpenStack environment on page 12](#)), start a continuous ping to the IP address of cirros-r. On the cirros-r instance console, start a continuous ping to the IP address of cirros-l:

```
$ ping 172.32.0.11
PING 172.32.0.11 (172.32.0.11): 56 data bytes
64 bytes from 172.32.0.11: seq=0 ttl=63 time=0.402 ms
64 bytes from 172.32.0.11: seq=0 ttl=63 time=0.433 ms
64 bytes from 172.32.0.11: seq=0 ttl=63 time=0.502 ms
64 bytes from 172.32.0.11: seq=0 ttl=63 time=0.408 ms
64 bytes from 172.32.0.11: seq=0 ttl=63 time=0.362 ms
```

On both FortiGate-VMs, use the following `diagnose sniffer packet` command to sniff ICMP packets. You should only see packets going through the primary unit.

```
fgt-vm-1 # diagnose sniffer packet any 'icmp' 4
interfaces = [any]
filters= [icmp]
109.413710 port_ha in 169.251.0.1 - > 169.251.0.2: icmp: 169.251.0.1 udp port 53
unreachable
111.797651 port2 in 172.32.0.11 - > 172.33.0.12: icmp: echo request
111.797676 port3 out 172.33.0.1 - > 172.33.0.12: icmp: echo request
111.797932 port3 in 172.33.0.12 - > 172.33.0.1: icmp: echo reply
111.797910 port2 out 172.33.0.12 - > 172.32.0.11: icmp: echo reply
112.372066 port3 in 172.33.0.12 - > 172.32.0.11: icmp: echo request
112.372081 port2 out 172.32.0.9 - > 172.32.0.11: icmp: echo request
112.372225 port2 in 172.32.0.11 - > 172.32.0.9: icmp: echo reply
112.372232 port3 out 172.32.0.11 - > 172.33.0.12: icmp: echo reply
112.797831 port2 in 172.32.0.11 - > 172.33.0.12: icmp: echo request
112.797839 port3 out 172.33.0.1 - > 172.33.0.12: icmp: echo request
112.798019 port3 in 172.33.0.12 - > 172.33.0.1: icmp: echo reply
112.798021 port2 out 172.33.0.12 - > 172.32.0.11: icmp: echo reply
```

Shut down the primary unit. You can do this from the OpenStack Horizon *Instances* list.

After failover, enter the following `diagnose` command from the new primary unit to verify that the pings are now going through that unit:

```
fgt-vm-2 # diagnose sniffer packet any ' icmp' 4
interfaces= [any]
filter s= [icmp]
0.360973 port3 in 172.33.0.12 - > 172.32.0.11: icmp: echo request
0.360983 port2 out 172.32.0.9 - > 172.32.0.11: icmp: echo request
0.361220 port2 in 172.32.0.11 - > 172.32.0.9: icmp: echo reply
```

```
0.361222 port3 out 172.32.0.11 - > 172.33.0.12: icmp: echo reply
0.785522 port2 in 172.32.0.11 - > 172.33.0.12: icmp: echo request
0.785527 port3 out 172.33.0.4 - > 172.33.0.12: icmp: echo request
0.785688 port3 in 172.33.0.12 - > 172.33.0.4: icmp: echo reply
0.785690 port2 out 172.33.0.12 - > 172.32.0.11: icmp: echo reply
1.360860 port3 in 172.33.0.12 - > 172.32.0.11: icmp: echo request
1.360864 port2 out 172.32.0.9 - > 172.32.0.11: icmp: echo request
1.361025 port2 in 172.32.0.11 - > 172.32.0.9: icmp: echo reply
1.361027 port3 out 172.32.0.11 - > 172.33.0.12: icmp: echo reply
```

Restart the FortiGate-VM instance that you shut down. After a short while it should re-join the cluster.

Verifying HA cluster status

On a FortiGate-VM in an high availability (HA) cluster, you can use the following command to verify the cluster status:

```
fgt-vm # diagnose sys ha status
HA information
Statistics
  traffic.local = s:0 p:42311 b:9008646
  traffic.total = s:0 p:42316 b:9009528
  activity.fdb = c:0 q:0
Model=80008, Mode=2 Group=0 Debug=0
nvcluster=1, ses_pickup=0, delay=0
[Debug_Zone HA information]
HA group member information: is_manage_master=1.
FGVM080000109643: Master, serialno_prio=0, usr_priority=128, hostname=fgt-vm
FGVM080000103268: Slave, serialno_prio=1, usr_priority=128, hostname=fgt-vm
[Kernel HA information]
vcluster 1, state=work, master_ip=169.254.0.1, master_id=0:
FGVM080000109643: Master, ha_prio/o_ha_prio=0/0
FGVM080000103268: Slave, ha_prio/o_ha_prio=1/1
```

The following command shows similar information:

```
fgt-vm # get system ha status
HA Health Status: OK
Model: FortiGate-VM64-KVM
Mode: HA A-P
Group: 0
Debug: 0
Cluster Uptime: 0 days 02:04:26
Cluster state change time: 2017-09-01 03:08:19
Master selected using:
  <2017/09/01 03:08:19> FGVM080000109643 is selected as the master because it has the largest value
  of serialno.
ses_pickup: disable
override: disable
Configuration Status:
  FGVM080000109643(updated 2 seconds ago): in-sync
  FGVM080000103268(updated 0 seconds ago): out-of-sync
System Usage stats:
  FGVM080000109643(updated 2 seconds ago):
    sessions=4, average-cpu-user/nice/system/idle=0%/0%/0%/100%, memory=55%
```

```
FGVM080000103268(updated 0 seconds ago):
  sessions=0, average-cpu-user/nice/system/idle=0%/0%/0%/100%, memory=54%
HBDEV stats:
FGVM080000109643(updated 2 seconds ago):
  port4: physical/10000full, up, rx-bytes/packets/dropped/errors=15043566/61878/0/0,
  tx=158364378/146977/0/0
FGVM080000103268(updated 0 seconds ago):
  port4: physical/10000full, up, rx-bytes/packets/dropped/errors=29442835/61625/49/0,
  tx=25246662/68626/0/0
MONDEV stats:
FGVM080000109643(updated 2 seconds ago):
  port2: physical/10000full, up, rx-bytes/packets/dropped/errors=1892/8/0/0, tx=173710/307/0/0
FGVM080000103268(updated 0 seconds ago):
  port2: physical/10000full, up, rx-bytes/packets/dropped/errors=174390/306/0/0, tx=2352/13/0/0
Master: fgt-vm , FGVM080000109643
Slave : fgt-vm , FGVM080000103268
number of vcluster: 1
vcluster 1: work 169.254.0.1
Master:0 FGVM080000109643
Slave :1 FGVM080000103268
```

The command `diagnose system ha checksum show` shows whether the configurations of the FortiGate-VMs in the cluster are synchronized. If the configurations are synchronized, both sets of checksums should match.

```
fgt-vm # diagnose sys ha checksum show
is_manage_master()=1, is_root_master()=1
debugzone
global: 33 6f ee 5b 78 a5 22 84 39 ec 36 d3 1c 54 7c 78
root: 40 0d fb 04 12 41 df ad f1 64 14 03 ff ec f5 01
all: d3 2f 6f bb a6 e7 77 db 27 75 81 b2 94 f3 fd 68
checksum
global: 33 6f ee 5b 78 a5 22 84 39 ec 36 d3 1c 54 7c 78
root: 40 0d fb 04 12 41 df ad f1 64 14 03 ff ec f5 01
all: d3 2f 6f bb a6 e7 77 db 27 75 81 b2 94 f3 fd 68
```

If the checksums do not match, you can use the `diagnose sys ha checksum show` and `diagnose sys ha checksum show global` commands to show more detailed checksum results. The following example shows the first few lines of output of the `diagnose sys ha checksum show global` command:

```
diagnose sys ha checksum show global
system.global: 2c79958c132639dfe61ab782a2f213ec
system.accprofile: 7d79452c78377be2616149264a18fd5c
system.vdom-link: 00000000000000000000000000000000
wireless-controller.inter-controller: 00000000000000000000000000000000
wireless-controller.global: 00000000000000000000000000000000
wireless-controller.vap: 00000000000000000000000000000000
system.switch-interface: 00000000000000000000000000000000
system.interface: 8690699bc33c7c15b20e017876cf1e37
...
```

If the configurations are synchronized, all the checksums displayed using these commands from both FortiGate-VMs should match. If they do not, you can use the output to see what parts of the configuration are not synchronized.

Deploying a FortiGate-VM instance in an OpenStack environment using service insertion or chaining

This version provides NSH chaining support for virtual wire pair, TP mode networks. FortiOS receives and unwraps the NSH packets and re-encapsulates them before sending them out. Firewall policies process the inner packet.

NSH support in FortiGate is basically unwrapping the packet on ingress and putting the NSH header back on before sending it out. FortiOS does not yet support other parts of NSH (SI is left unchanged).

There is no CLI or GUI change. The only change is to show `ext_header=nsh` in NSH session information when listing sessions.

Sample configuration

To configure virtual wire pair and firewall policy using the CLI:

```
config system virtual-wire-pair
  edit "test-vw"
    set member "port1" "mgmt2"
  next
end
config firewall policy
  edit 99
    set uuid 241710a0-3ac6-51e9-10e9-9dd3eb65e708
    set srcintf "mgmt2"
    set dstintf "port1"
    set srcaddr "all"
    set dstaddr "all"
    set action accept
    set schedule "always"
    set service "ALL"
    set logtraffic all
  next
end
```

Sample results of configuring a wire pair and policy between port1 and mgmt2. Packets with NSH are processed and the session list shows `ext_header=nsh`.

```
A (vdom1) # diag sys session list
session info: proto=6 proto_state=01 duration=10 expire=3595 timeout=3600 flags=00000000
  sockflag=00000000 sockport=0 av_idx=0 use=4
origin-shaper=
reply-shaper=
per_ip_shaper=
class_id=0 ha_id=0 policy_dir=0 tunnel=/ vlan_cos=0/0
```

```
state=log may_dirty br src-vis dst-vis f00
statistic(bytes/packets/allow_err): org=112/2/1 reply=60/1/1 tuples=2
tx speed(Bps/kbps): 10/0 rx speed(Bps/kbps): 5/0
origin->sink: org pre->post, reply pre->post dev=4->9/9->4 gwy=0.0.0.0/0.0.0.0
hook=pre dir=org act=noop 172.16.200.11:46739->172.16.200.55:23(0.0.0.0:0)
hook=post dir=reply act=noop 172.16.200.55:23->172.16.200.11:46739(0.0.0.0:0)
pos/(before,after) 0/(0,0), 0/(0,0)
src_mac=00:00:11:11:11:11 dst_mac=00:00:22:22:22:22
misc=0 policy_id=99 auth_info=0 chk_client_info=0 vd=1
serial=0000094d tos=ff/ff app_list=0 app=0 url_cat=0
rpidb_link_id = 00000000
dd_type=0 dd_mode=0
npu_state=0x040001 no_offload
no_ofld_reason: mac-host-check disabled-by-policy non-npu-intf
ext_header_type=nsh
total session 1
```

Optimizing FortiGate-VM performance

The FortiGate-VM and OpenStack performance optimization techniques that this section describes can improve your FortiGate-VM performance by optimizing the hardware and the OpenStack host environment for network- and CPU-intensive performance FortiGate-VM requirements.

SR-IOV

FortiGate-VMs installed on OpenStack platforms support single root I/O virtualization (SR-IOV) to provide FortiGate-VMs with direct access to hardware devices. Enabling SR-IOV means that one PCIe device (CPU or network card) can function for a FortiGate-VM as multiple separate physical devices (CPUs or network devices). SR-IOV reduces latency and improves CPU efficiency by allowing network traffic to pass directly between a FortiGate-VM and network card without passing through the OpenStack kernel or using virtual switching.

FortiGate-VMs benefit from SR-IOV because SR-IOV optimizes network performance and reduces latency. FortiGate-VMs do not use OpenStack features that are incompatible with SR-IOV so you can enable SR-IOV without negatively affecting your FortiGate-VM.

SR-IOV hardware compatibility

SR-IOV requires that the hardware on which your OpenStack host is running has BIOS, physical NIC, and network driver support for SR-IOV.

To enable SR-IOV, your OpenStack platform must run on hardware that is compatible with SR-IOV and with FortiGate-VMs. FortiGate-VMs require network cards that are compatible with ixgbevf or i40evf drivers.

For optimal SR-IOV support, install the most up to date ixgbevf or i40evf network drivers.

To create SR-IOV VFs:

This section describes how to create virtual functions (VFs) for SR-IOV-compatible Intel network interfaces. An SR-IOV VF is a virtual PCIe device that you must add to OpenStack to allow your FortiGate-VM to use SR-IOV to communicate with a physical Ethernet interface or physical function (PF).

1. Enable SR-IOV in the host system's BIOS by enabling VT-d.
2. Enable IOMMU for Linux by adding `intel_iommu=on` to kernel parameters. Do this by adding the following line to the `/etc/default/grub` file:

```
GRUB_CMDLINE_LINUX_DEFAULT="nomdmonddf nomdmonisw intel_iommu=on"
```
3. Save your changes. From the Linux command line, enter the following commands to update grub and reboot the host device:

```
# update-grub  
# reboot
```
4. On each compute node, create VFs using the PCI SYS interface:

```
# echo '7' > /sys/class/net/eth3/device/sriov/numvfs
```

5. If the previous command produces a Device or resource busy error message, you must set `sriov_numvfs` to 0 before setting it to the new value.
6. Optionally determine the maximum number of VFs a PF can support:

```
# cat /sys/class/net/eth3/device/sriov_totalvfs
```
7. Enter the following command to ensure an SR-IOV interface is up and verify its status:

```
# ip link set eth3 up
# ip link show eth3
```
8. Enter the following command to verify that the VFs have been created:

```
# lspci | grep Ethernet
```
9. Enter the following command to ensure the VFs are recreated when the system reboots:

```
# echo "echo '7' > /sys/class/net/eth3/device/sriov_numvfs" >> /etc/rc.local
```

To allowlist PCI devices:

You must allowlist SR-IOV devices so their traffic can pass through OpenStack to the FortiGate-VM. The following example shows how to allowlist SR-IOV devices by modifying the nova-compute service. (You can also edit the `pci_passthrough_whitelist` parameter to add allowlisting.)

1. To modify the nova-compute service, open the `nova.comp` file and add the following line. This setting adds traffic from eth3 to the physnet2 physical network and allows physnet2 traffic to pass through OpenStack to your FortiGate-VM:

```
pci_passthrough_whitelist = { "devname": "eth3", "physical_network": "physnet2" }
```
2. After entering this command, restart the nova-compute service.

To configure neutron-server:

Use the following steps to configure OpenStack neutron-server to support SR-IOV:

1. Add the `sriovnicswitch` as mechanism driver, edit the `m12_conf.ini` file and add the following line:

```
mechanism_drivers = openvswitch,sriovnicswitch
```
2. Find the `vendor_id` and `product_id` of the VFs that you created. For example:

```
# lspci -nn | grep -i ethernet
87:00.0 Ethernet controller [0200]: Intel Corporation 82599 10 Gigabit Dual Port Backplane
      Connection [8086:10f8] (rev 01)
87:10.1 Ethernet controller [0200]: Intel Corporation 82599 Ethernet Controller Virtual Function
      [8086:10ed] (rev 01)
87:10.3 Ethernet controller [0200]: Intel Corporation 82599 Ethernet Controller Virtual Function
      [8086:10ed] (rev 01)
```
3. Add the following line to the `m12_conf_sriov.ini` on each controller:

```
supported_pci_vendor_devs = 8086:10edM
```

In this example the `vendor_id` is 8086 and the `product_id` is 10ed.
4. Add `m12_conf_sriov.ini` to the neutron-server daemon. Edit the initialization script to configure the neutron-server service to load the SR-IOV configuration file. Include the following lines:

```
--config-file /etc/neutron/neutron.conf --config-file /etc/neutron/plugin.ini
--config-file /etc/neutron/plugins/m12/m12_conf_sriov.ini
```
5. Restart the neutron-server service.

To configure the nova-schedule controller:

To complete this step, on controllers running the nova-scheduler service, add `PciPassthroughFilter` to the `scheduler_default_filters` parameter and add the following new line under the [DEFAULT] section in

nova.conf:

```
[DEFAULT]
scheduler_default_filters = RetryFilter, AvailabilityZoneFilter, RamFilter, ComputeFilter,
    ComputeCapabilitiesFilter, ImagePropertiesFilter, ServerGroupAntiAffinityFilter,
    ServerGroupAffinityFilter, PciPassthroughFilter
scheduler_available_filters = nova.scheduler.filters.all_filters
scheduler_available_filters = nova.scheduler.filters.pci_passthrough_filter.PciPassthroughFilter
```

Restart the nova-scheduler service.

To enable the Neutron sriov-agent process:

To enable the sriov-agent process, on each compute node, edit the sriov_agent.ini file and add the following:

Under [securitygroup] add:

```
firewall_driver = neutron.agent.firewall.NoopFirewallDriver
```

Under [sriov_nic] add:

```
physical_device_mappings = physnet2:eth3
exclude_devices =
```

The example physical_device_mappings setting includes one mapping between the physical network (physnet2) and one VF called eth3. If you have multiple VFs connected to the same physical network, you can add them all using the following syntax that shows how to add two VFs to physnet2.

```
physical_device_mappings = physnet2:eth3,physnet2:eth4
```

Also in the example, exclude_devices is empty and all VFs associated with eth3 may be configured by the agent. You can also use exclude_devices to exclude specific VFs, for example to exclude eth1 and eth2:

```
exclude_devices = eth1:0000:07:00.2; 0000:07:00.3, eth2:0000:05:00.1; 0000:05:00.2
```

Enter the following command to verify that the neutron sriov_agent runs successfully:

```
# neutron-sriov-nic-agent --config-file /etc/neutron/neutron.conf --config-file
    /etc/neutron/plugins/ml2/sriov_agent.ini
```

Finally, you should enable the neutron sriov_agent service.

To assign SR-IOV interfaces to a FortiGate-VM:

After SR-IOV has been added to your OpenStack host, you can now launch FortiGate-VM instances with neutron SR-IOV ports. Use the following steps:

Use the following command to display, the ID of the neutron network where you want the SR-IOV port to be created.

```
$ net_id=`neutron net-show net04 | grep "\ id\ " | awk '{ print $4 }'`
```

Use the following command to create the SR-IOV port. This command sets vnic_type=direct. Other options include normal, direct-physical, and macvtap:

```
$ port_id=`neutron port-create $net_id --name sriov_port --binding:vnic_type direct | grep "\ id\ "
    | awk '{ print $4 }'`
```

Create the VM. This example includes the SR-IOV port created in the previous step:

```
$ nova boot --flavor m1.large --image ubuntu_14.04 --nic port-id=$port_id test-sriov
```

FortiGate-VM interrupt affinity

In addition to enabling SR-IOV in the VM host, to fully take advantage of SR-IOV performance improvements you must configure interrupt affinity for your FortiGate-VM. Interrupt affinity (also called CPU affinity) maps FortiGate-VM interrupts to the CPUs that are assigned to your FortiGate-VM. You use a CPU affinity mask to define the CPUs that the interrupts are assigned to.

A common use of this feature would be to improve your FortiGate-VM's networking performance by doing the following:

- On the VM host:
 - Add multiple host CPUs to your FortiGate-VM.
 - Configure CPU affinity to specify the CPUs that the FortiGate-VM can use.
 - Configure other VM clients on the VM host to use other CPUs.
- On the FortiGate-VM, assign network interface interrupts to a CPU affinity mask that includes the CPUs that the FortiGate-VM can use.

In this way, the configuration uses all available CPU interrupts for the configured host CPUs to process traffic on your FortiGate interfaces. This configuration could lead to improved FortiGate-VM network performance because you have dedicated VM host CPU cycles processing your FortiGate-VM's network traffic.

You can use the following CLI command to configure interrupt affinity for your FortiGate-VM:

```
config system affinity-interrupt
  edit <index>
    set interrupt <interrupt-name>
    set affinity-cpumask <cpu-affinity-mask>
  next
```

The following defines the values for the commands:

Value	Description
<interrupt-name>	Name of the interrupt to associate with a CPU affinity mask. You can view your FortiGate-VM interrupts using the <code>diagnose hardware sysinfo interrupts</code> command. Usually you would associate all of the interrupts for a given interface with the same CPU affinity mask.
<cpu-affinity-mask>	CPU affinity mask for the CPUs that process the associated interrupt.

For example, consider the following configuration:

- The FortiGate-VM port2 and port3 interfaces send and receive most of the traffic.
- On the VM host, you have set up CPU affinity between your FortiGate-VM and four CPUs (CPU 0, 1, 2, and 3).
- SR-IOV is enabled and SR-IOV interfaces use the i40evf interface driver.

Output from `diagnose hardware sysinfo interrupts` shows that port2 has the following transmit and receive interrupts:

```
i40evf-port2-TxRx-0
i40evf-port2-TxRx-1
i40evf-port2-TxRx-2
i40evf-port2-TxRx-3
```

Output from `diagnose hardware sysinfo interrupts` shows that port3 has the following transmit and receive interrupts:

```
i40evf-port3-TxRx-0
i40evf-port3-TxRx-1
i40evf-port3-TxRx-2
i40evf-port3-TxRx-3
```

Use the following command to associate the port2 and port3 interrupts with CPU 0, 1, 2, and 3:

```
config system affinity-interrupt
  edit 1
    set interrupt "i40evf-port2-TxRx-0"
    set affinity-cpumask "0x0000000000000001"
  next
  edit 2
    set interrupt "i40evf-port2-TxRx-1"
    set affinity-cpumask "0x0000000000000002"
  next
  edit 3
    set interrupt "i40evf-port2-TxRx-2"
    set affinity-cpumask "0x0000000000000004"
  next
  edit 4
    set interrupt "i40evf-port2-TxRx-3"
    set affinity-cpumask "0x0000000000000008"
  next
  edit 1
    set interrupt "i40evf-port3-TxRx-0"
    set affinity-cpumask "0x0000000000000001"
  next
  edit 2
    set interrupt "i40evf-port3-TxRx-1"
    set affinity-cpumask "0x0000000000000002"
  next
  edit 3
    set interrupt "i40evf-port3-TxRx-2"
    set affinity-cpumask "0x0000000000000004"
  next
  edit 4
    set interrupt "i40evf-port3-TxRx-3"
    set affinity-cpumask "0x0000000000000008"
  next
end
```

FortiGate-VM affinity packet redistribution

With SR-IOV enabled on the VM host and interrupt affinity configured on your FortiGate-VM, there is an additional configuration that you can add that may improve performance. Most common network interface hardware restricts the number of RX/TX queues that it can process. This can result in some CPUs being much busier than others and the busy CPUs may develop extensive queues.

You can get around this potential bottleneck by configuring affinity packet redistribution to allow overloaded CPUs to redistribute packets they receive to other less busy CPUs. This may result in a more even distribution of packet processing to all available CPUs.

You configure packet redistribution for interfaces by associating an interface with an affinity CPU mask. This configuration distributes packets that that interface sets and receives to the CPUs that the CPU affinity mask associated with the interface defines.

You can use the following CLI command to configure affinity packet redistribution for your FortiGate-VM:

```
config system affinity-packet-redistribution
  edit <index>
    set interface <interface-name>
    set affinity-cpumask <cpu-affinity-mask>
  next
```

The following defines the values for the commands:

Value	Description
<interface-name>	Name of interface to associate with a CPU affinity mask.
<cpu-affinity-mask>	CPU affinity mask for the CPUs that process packets to and from the associated interface.

For example, you can improve the performance of the interrupt affinity example that the following command shows to allow packets that the port3 interface sends and receives to be redistributed to CPUs according to the 0xE CPU affinity mask.

```
config system affinity-packet-redistribution
  edit 1
    set interface port3
    set affinity-cpumask "0xE"
  next
end
```

FortiGate-VM MTU setting

For optimal performance, you can set your FortiGate-VM interfaces MTU to be compatible with the OpenStack 10 environment, which by default has an MTU of 1446. See [Creating a user_data file to preconfigure FortiGate-VM instances on page 14](#) for details.

Automatically updating dynamic addresses using an SDN connector

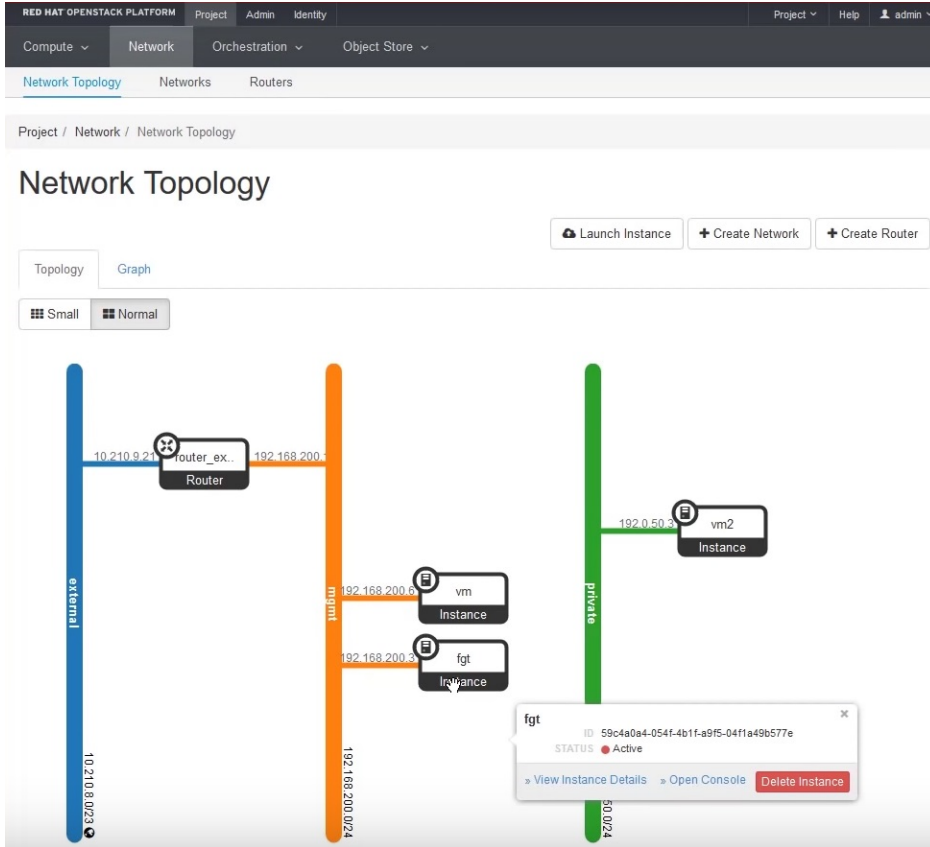
FortiOS supports the following OpenStack releases for this feature:

- Pike
- Queen
- Rocky
- Stein
- Train
- Ussuri
- Victoria
- Wallaby
- Xena

Configuring the OpenStack environment

This example assumes that the OpenStack environment is configured as follows:

- Three virtual machines (VM):
 - A FortiGate-VM
 - Two protected VMs
- Three networks



The following shows each VM IP address:

The screenshot shows the 'Instances' page in the OpenStack dashboard. It displays a table with the following data:

Instance Name	Image Name	IP Address	Size	Key Pair	Status	Availability Zone	Task	Power State	Time since created	Actions
<input type="checkbox"/> vm2	-	• 192.0.50.3	m1.nano	-	Active	nova	None	Running	37 minutes	Create Snapshot
<input type="checkbox"/> vm	-	• 192.168.200.6	m1.nano	-	Active	nova	None	Running	40 minutes	Create Snapshot
<input type="checkbox"/> fgt	-	• 192.168.200.3 Floating IPs: • 10.210.9.11	m1.fortigate	-	Active	nova	None	Running	14 hours, 56 minutes	Create Snapshot

Displaying 3 items

Troubleshooting OpenStack Horizon SDN connector

You can check if API calls are made successfully by running the following in the CLI:

```
diagnose debug enable
diagnose debug application openstackd -1
```

The console displays that debug messages are on for 30 minutes.

In the FortiOS GUI, toggle the SDN connector to disable it, then reenable it. If the SDN connector was configured correctly and can populate IP addresses, the CLI output resembles the following:

```
openstackd exit
openstackd start
openstackd request new tokens
openstackd request new token for project demo
openstackd request new token for project admin
openstackd list IP addresses for project admin successfully
openstackd list IP addresses for project demo successfully
openstackd finished address update
```

If something fails, the CLI shows the reason. The following example shows that the OpenStack identity's URL (IP address) may be incorrect:

```
FGVM123456 (global) # openstackd safeguard_fn()-1701
openstackd request new tokens
openstackd invalid url
openstackd failed to get unscoped token
openstackd failed to get token
```

Configuring OpenStack SDN connector with domain filter

You can select a domain attribute when configuring an OpenStack SDN connector in FortiOS. When a domain is configured for the OpenStack SDN connector, FortiOS resolves OpenStack dynamic firewall addresses from the specified OpenStack domain. If you do not specify a domain, FortiOS resolves the dynamic firewall addresses using the default OpenStack domain.

To configure OpenStack SDN connector with a domain filter using the GUI:

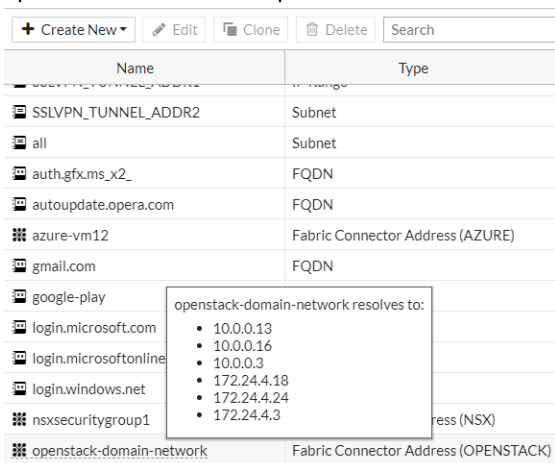
1. Configure the OpenStack SDN connector:
 - a. Go to *Security Fabric > External Connectors*.
 - b. Click *Create New*, and select *OpenStack (Horizon)*.
 - c. In the *Domain* field, enter the desired domain name from OpenStack. The SDN connector only resolves IP addresses for instances that belong to the specified domain.
 - d. Configure the connector, using the server IP address, username, and password for your deployment.

2. Create a dynamic firewall address for the configured OpenStack SDN connector:

- a. Go to *Policy & Objects > Addresses*.
- b. Click *Create new*.
- c. Configure the address:
 - i. From the *Type* dropdown list, select *Dynamic*.
 - ii. From the *Sub Type* dropdown list, select *Fabric Connector Address*.
 - iii. From the *SDN Connector* dropdown list, select the OpenStack connector.
 - iv. In the *Filter* dropdown list, configure a filter, such as *Network=<networkname>*. The OpenStack SDN connector automatically populates and updates IP addresses only for instances that belong to the specified domain and network.
 - v. Click *OK*.

3. Ensure that the OpenStack SDN connector resolves dynamic firewall IP addresses:

- a. Go to *Policy & Objects > Addresses*.
- b. Hover over the address created in step 2 to see a list of IP addresses for instances that belong to the specified domain and specified network as configured in steps 1 and 2:



To configure an OpenStack SDN connector with a domain filter using CLI commands:

- 1. Configure the OpenStack SDN connector. The SDN connector only resolves IP addresses for instances that belong to the specified domain:

```
config system sdn-connector
  edit "openstack-domain"
    set type openstack
    set server "http://172.16.165.86:5000"
    set username "example_username"
    set password xxxxx
    set domain "example_domain"
    set update-interval 30
  next
end
```

- 2. Create a dynamic firewall address for the configured OpenStack SDN connector with the supported OpenStack filter. The OpenStack SDN connector automatically populates and updates IP addresses only for instances that belong to the specified domain and the specified network:

```
config firewall address
  edit "openstack-domain-network"
```

```
        set type dynamic
        set sdn "openstack-domain"
        set filter "Network=example-net1"
    next
end
```

3. Confirm that the OpenStack SDN connector resolves dynamic firewall IP addresses using the configured domain and filter:

```
config firewall address
    edit "openstack-domain-network"
        set type dynamic
        set sdn "openstack-domain"
        set filter "Network=example-net1"
    config list
        edit "10.0.0.13"
        next
        edit "10.0.0.16"
        next
        edit "10.0.0.3"
        next
        edit "172.24.4.18"
        next
        edit "172.24.4.24"
        next
        edit "172.24.4.3"
        next
    end
next
end
```

Change log

Date	Change description
2026-04-21	Initial release.



www.fortinet.com

Copyright © 2026 Fortinet, Inc. All rights reserved. Fortinet®, FortiGate®, FortiCare® and FortiGuard®, and certain other marks are registered trademarks of Fortinet, Inc., and other Fortinet names herein may also be registered and/or common law trademarks of Fortinet. All other product or company names may be trademarks of their respective owners. Performance and other metrics contained herein were attained in internal lab tests under ideal conditions, and actual performance and other results may vary. Network variables, different network environments and other conditions may affect performance results. Nothing herein represents any binding commitment by Fortinet, and Fortinet disclaims all warranties, whether express or implied, except to the extent Fortinet enters a binding written contract, signed by Fortinet's Chief Legal Officer, with a purchaser that expressly warrants that the identified product will perform according to certain expressly-identified performance metrics and, in such event, only the specific performance metrics expressly identified in such binding written contract shall be binding on Fortinet. For absolute clarity, any such warranty will be limited to performance in the same ideal conditions as in Fortinet's internal lab tests. Fortinet disclaims in full any covenants, representations, and guarantees pursuant hereto, whether express or implied. Fortinet reserves the right to change, modify, transfer, or otherwise revise this publication without notice, and the most current version of the publication shall be applicable.