



# OpenStack Cloud Advanced Configuration Options

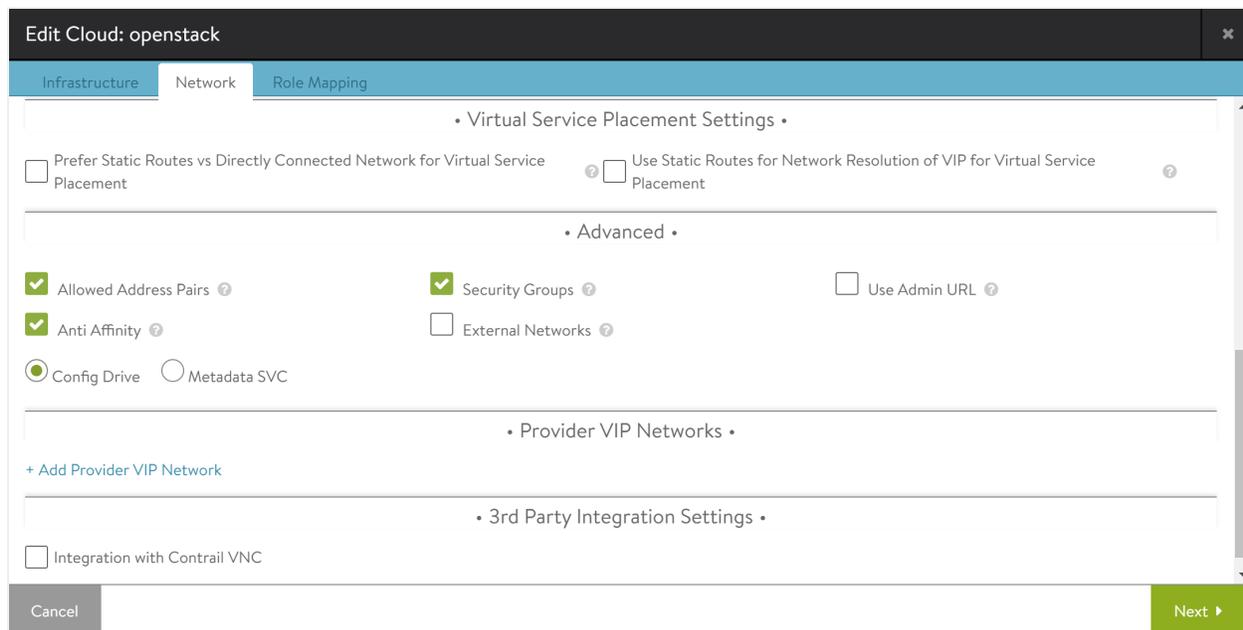
Avi Technical Reference (v20.1)

# OpenStack Cloud Advanced Configuration Options

[view online](#)

## Overview

This article explains the advanced configuration options relevant to the creation of an Avi Vantage OpenStack cloud. The advanced configuration options can be accessed from Step 3: Network tab of the Avi Vantage cloud editor wizard as shown in the image:



## Security-Groups

Default                      True

The security-groups Neutron extension supports specifying the allowed rules for both ingress and egress. Avi Vantage uses this extension to create one service group per Avi Service Engine (SE). This service group is created with all egress and Secure Shell (SSH) and Internet Control Message Protocol (ICMP) ingress. As virtual services (VS) are created and placed on this SE, the corresponding service ports are added to the service group. Similarly, when the virtual services are unplaced from the SE, the corresponding service ports are removed from the service group (if no longer used by any other VS on the same SE). When set to *True*, the security-group extension will be used. If the underlying network plugin does not support this feature, then Virtual IP (VIP) traffic will not work unless there are other means to achieve the same effect. This option can be turned off if the underlying network supports turning off of security filter rules on ports.

This example shows the security group of an SE with a virtual service with the service port '80' placed on it.

```
[root@sivacos ~(keystone_admin)]# neutron security-group-list
| 5544b75d-2a57-4f56-b1d0-ef68242293ba | avi-se-30af06c4-09c6-4c94-92be-f39d4dfddf91 | egress, IPv4
|                                     |                                     | egress, IPv6
```

```
|
|
| ingress, IPv4,22/tcp, remote_ip
| ingress, IPv4,80/tcp, remote_ip
| ingress, IPv4,icmp, remote_ip
```

## Anti-Affinity

Default True

Compute uses the nova-scheduler service to determine the host upon which to launch a virtual machine (VM), based on various criteria and filters. One such filter, `ServerGroupAntiAffinityFilter`, ensures that each instance in an anti-affinity group is on a different host of the group. Avi Vantage uses one anti-affinity group per SE group, thereby allowing each SE in the SE group to be placed on a different host. This provides better isolation of SEs in the event of host failures. If this option is set to `False`, anti-affinity filters will not be used. This option can be turned off if nova-compute has only one compute node.

This example shows an anti-affinity group, `serviceenginegroup-37dac996-7c88-4761-a920-6dc9d265c786` in a tenant with two SE VMs.

```
root@node-17:~# nova server-group-list

+-----+-----+
| ID                | Name                                     |
| Project ID       | User ID                                 |
| Policies         | Members                                 |
| Metadata         |                                         |
+-----+-----+
| c605a898-86fa-457f-80c8-f1db21dfb68a | avi-aasg-serviceenginegroup-37dac996-7c88-4761-a920-6dc9d265c786 |
| fefb594ef03e4670beaffe3305440e24    | aba3667db25e44afb5aff73f3f363027 |
| [u'anti-affinity']                   | [u'd7509390-6afe-4865-ade2-231e9a664421', u'1867c24e-8495-4cbf-80d0-06a2328656c6'] |
+-----+-----+

root@node-17:~# nova list | egrep "d7509390|1867c24e" | d7509390-6afe-4865-ade2-231e9a664421 | cc_os-se-ozmkj
| ACTIVE | - | Running | avimgmt=10.10.44.231 |
| 1867c24e-8495-4cbf-80d0-06a2328656c6 | cc_os-se-xmrzn
| ACTIVE | - | Running | network-80.21=10.80.21.13;avimgmt=10.10.44.230
```

## External-Networks

Default False

When set to `True`, this option enables selection of OpenStack networks marked 'external' for Avi management, VIP or data networks.

## Metadata Provisioning

Default Config-drive

OpenStack allows metadata to be passed on to VMs using: \* Config-drive: Metadata is written to a special configuration drive that attaches to the instance when it boots. The instance can mount this drive and read the data. Please refer to [this](#).

[OpenStack document](#) for further details. \* Metadata-service: Instances can access the metadata-service at <http://169.254.169.254> in most common configurations to retrieve instance-specific data. Avi Vantage supports both options, with `config-drive` preferred over `metadata-service`; and requires the OpenStack deployment to support one of these options.

## VIP Placement

Default      `allowed-address-pairs`

Avi Vantage supports multiple VIP placement methodologies using:

- Allowed-address-pairs

Default      `True`

This extension is a Neutron extension that allows traffic with specific Classless Inter-Domain Routing (CIDR) to egress from a port. Avi Vantage uses this to place VIPs on SE data ports, thereby allowing VIP traffic to egress these data ports. When set to `True`, the `allowed-address-pairs` extension will be used. If the underlying network plugin does not support this feature, then VIP traffic will not work unless there are other means to achieve the same effect. This option can be turned off if the underlying network supports turning off security, firewall, or spoof filter rules on ports. In this mode,

- An unbound VIP port is created.
- The VIP address is added as an Allowed-Address Pairs (AAPs) entry on the Service Engine's Data VNIC.
- The floating IP is associated to the unbound VIP port.

This option is ignored for Contrail integration (when `integration with Contrail VNC` is checked under 3rd Party Integration Settings).

If DVR is in use, in Mitaka and older versions of OpenStack, Floating IP doesn't work with AAPs.

There is a limit of 10 AAPs on a vNIC by default.

Going forward, with Contrail based deployments in conjunction with Avi Vantage, users would see `/32` entries being populated for AAP (IPv4) and `/128` prefixes for AAP (IPv6). This is done to make sure that the entire subnet for AAP ( prior to Avi Vantage Release 18.1 for IPv4 it was `/24`) is not subjected to DDoS.

For instance,

Prior to Release 18.1

```

root@contrail2:~# neutron port-show 87b01c63-f18d-477e-ad6c-c3689c89ad8f
+-----+-----+
| Field                | Value                |
+-----+-----+
| admin_state_up      | True                 |
|
|
|<b> allowed_address_pairs | {"ip_address": "b100::f/128", "mac_address": ""} | | {"ip_address": "192.168.124.21/24",
+----output truncated-----+

root@contrail2:~#
    
```

Release 18.1 onwards

```

root@contrail2:~# neutron port-show 26ecd42e-ddea-47ec-a749-cca3dae215
+-----+-----+
| Field          | Value          |
+-----+-----+
| admin_state_up | True           |
|<b> allowed_address_pairs | {"ip_address": "10.40.3.12", "mac_address": ""} | {"ip_address": "fd10:120:3::c", "mac_a
+----output truncated-----+
    
```

**Note:** To increase the scale of Avi Vantage Service Engines beyond 10 AAPs, one may increase the limit of 10 in the underlying OpenStack deployment. Refer to [this OpenStack.org page](#).

**Example:** On OVS with iptables, the "a" rule for 172.24.10.7 would be added to the Avi data port with UUID prefix 019ec61b.

```

[root@sivacos ~(keystone_admin)]# neutron port-show 019ec61b-3be2-4e25-a4a8-d48740ffa3a
+-----+-----+
| Field          | Value          |
+-----+-----+
| admin_state_up | True           |
|
| allowed_address_pairs | {"ip_address": "172.24.10.7", "mac_address": "fa:16:3e:47:a2:0e"}
| binding:vif_details | {"port_filter": true, "ovs_hybrid_plug": true} |
| device_id      | c32926e6-6c86-49a0-90a4-9e634a7ac6dd |
| device_owner   | compute:None   |
| fixed_ips      | {"subnet_id": "b679630f-f3d1-4a32-86ca-04ef04534adc", "ip_address": "172.24.10.11"}
| id             | 019ec61b-3be2-4e25-a4a8-d48740ffa3ad |
| mac_address    | fa:16:3e:47:a2:0e |

[root@sivacos ~(keystone_admin)]# iptables -S | grep -i fa:16:3e:47:a2:0e-A neutron-openvswi-s019ec61b-3 -s 172.24.10.7
    
```

### Map-admin-to-cloudadmin

Default            False

By default, the Avi admin tenant maps to OpenStack admin tenant. If set to *True*, then the Avi admin tenant maps to the admin\_tenant configured in the Avi cloud. This directly maps the load-balancer-related operations onto the corresponding tenant in OpenStack.

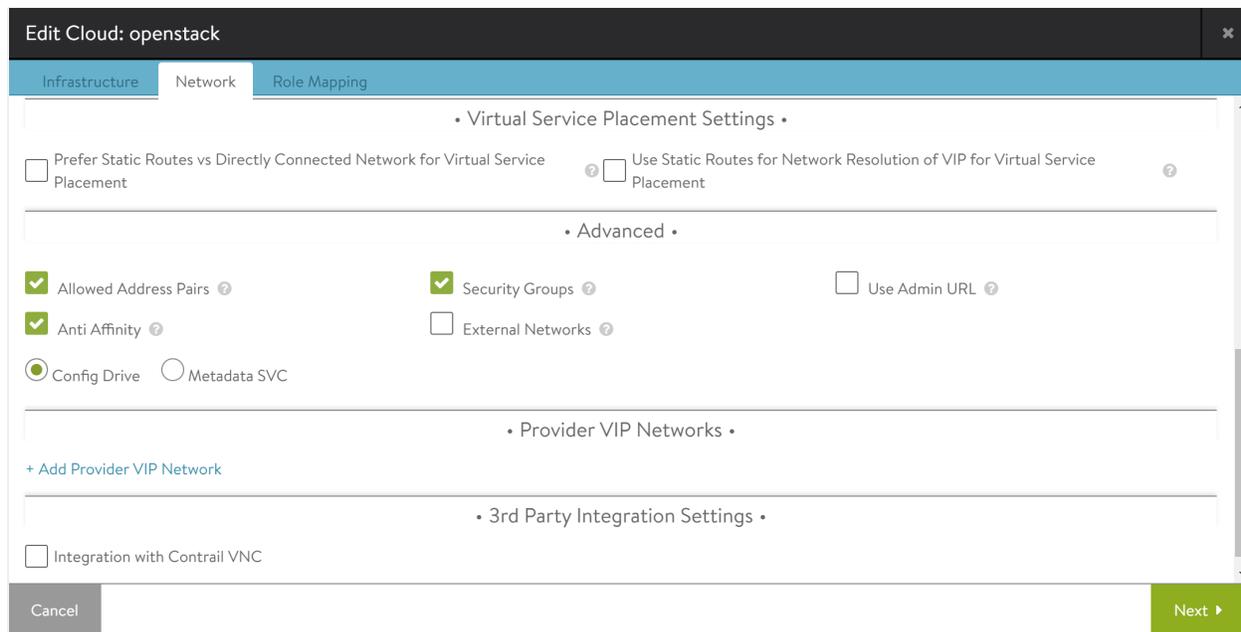
### Neutron-rbac

Default            True

By default, Avi Vantage consults the Neutron role-based-access-control (RBAC) rules to retrieve the 'usable' list of networks for a tenant. This list would normally include the tenant's own networks, any non-tenant networks widely shared with 'all', and any non-tenant networks explicitly shared with the tenant using RBAC. This flag is useful in a provider-mode SE configuration and, if set to *False*, the RBAC shared networks are not included in the 'usable' list.

## Provider VIP Networks

A tenant can normally use its own networks and any networks shared with it. In addition, this setting provides extra networks that are usable by tenants.

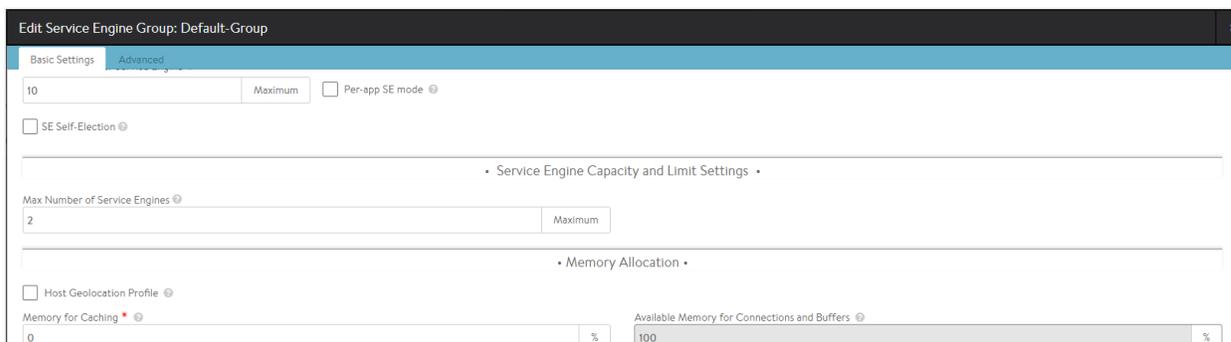


To provide extra networks that are usable by clients, click on Add Provider VIP Network button. You can specify the following details: \* Provider VIP Networks- Select the network provider details form the drop-down list. This is a mandatory field if you click Add Provider VIP Network button. \* Tenants- Select the networks for a tenant. This is a mandatory field if you click Add Provider VIP Network button.

## Configuring Service Engine Group Flavor Settings

You can configure Service Engine flavours as follows:

1. Click on Infrastructure > Service Engine Group tab. Choose the desired Service Engine Group and click on edit icon to set the Instance Flavor option.



The screenshot shows a configuration window with a slider for 'Connections and Buffers Memory Distribution' set to 50% for connections and 50% for buffers. Below the slider, there are radio buttons for 'Pay-as-you-go' (selected) and 'Bring your own license'. There are dropdown menus for 'License Type' (set to 'Cores') and 'License Tier' (set to 'Enterprise 18'). There is also a checkbox for 'Allow Burst' (unchecked) and a dropdown for 'Instance Flavor' (set to 'm1.se'). At the bottom, there are 'Cancel' and 'Save' buttons.

2. Click on Basic Settings and then click on Instance Flavor.

3. Select Instance Flavor for SE instance from the drop-down list.

The screenshot shows the 'Edit Service Engine Group: Default-Group' configuration page. It has tabs for 'Basic Settings' and 'Advanced'. Under 'Basic Settings', there are input fields for '10' and 'Maximum', a checkbox for 'Per-app SE mode', and a checkbox for 'SE Self-Election'. There is a section for 'Service Engine Capacity and Limit Settings' with a 'Max Number of Service Engines' set to '2' and 'Maximum'. Under 'Memory Allocation', there is a checkbox for 'Host Geolocation Profile', a 'Memory for Caching' field set to '0', and an 'Available Memory for Connections and Buffers' slider set to '100'. The 'Connections and Buffers Memory Distribution' slider is also shown. The 'License' section has radio buttons for 'Pay-as-you-go' (selected) and 'Bring your own license'. There are dropdowns for 'License Type' (set to 'Cores') and 'License Tier' (set to 'Enterprise 18'). There is a list of instance flavors: 'm1.controller', 'm1.se', and 'Auto select flavor'. At the bottom, there are 'Cancel' and 'Save' buttons.

**Note:** SE Group page takes Instance Flavor as configuration input. The input is flavour name not Flavor UUID. The list of flavours fetched from OpenStack are only those flavors matching the minimum requirement of 1 GB RAM and recommended minimum disk size of 2 times RAM (in GB) + 5 GB.

You can manually configure the flavor if you want to use flavors other than recommended flavor using Avi CLI as follows:

```
[admin:avi-controller]: > configure serviceenginegroup Default-Group
Updating an existing object. Currently, the object is:
+-----+-----+
| Field                | Value                                     |
+-----+-----+
| uuid                  | serviceenginegroup-052ac264-a5d8-43a6-b3ff-53eae9e29f54 |
```

```

| name | Default-Group |
| max_vs_per_se | 10 |
| min_scaleout_per_vs | 1 |
| max_scaleout_per_vs | 4 |
| max_se | 2 |
| active_standby | False |
| placement_mode | PLACEMENT_MODE_AUTO |
| instance_flavor | m1.se |
| auto_rebalance_interval | 300 sec |
| aggressive_failure_detection | False |
| vs_scaleout_timeout | 30 sec |
| vs_scalein_timeout | 30 sec |
| config_debugs_on_all_cores | False |
| accelerated_networking | True |
| vs_se_scaleout_ready_timeout | 25 sec |
| vs_se_scaleout_additional_wait_time | 0 sec |
| bgp_state_update_interval | 10 sec |
| max_memory_per_mempool | 64 mb |
+-----+
[admin:avi-controller]: serviceenginegroup> instance_flavor m1.se
Overwriting the previously entered value for instance_flavor
[admin:avi-controller]: serviceenginegroup> save
+-----+
| Field | Value |
+-----+
| uuid | serviceenginegroup-052ac264-a5d8-43a6-b3ff-53eae9e29f54 |
| name | Default-Group |
| max_vs_per_se | 10 |
| min_scaleout_per_vs | 1 |
| max_scaleout_per_vs | 4 |
| max_se | 2 |
| active_standby | False |
| placement_mode | PLACEMENT_MODE_AUTO |
| instance_flavor | m1.se |
| auto_rebalance_interval | 300 sec |
| aggressive_failure_detection | False |
| vs_scaleout_timeout | 30 sec |
| vs_scalein_timeout | 30 sec |
| config_debugs_on_all_cores | False |
| accelerated_networking | True |
| vs_se_scaleout_ready_timeout | 25 sec |
| vs_se_scaleout_additional_wait_time | 0 sec |
| bgp_state_update_interval | 10 sec |
| max_memory_per_mempool | 64 mb |
+-----+
[admin:avi-controller]: >

```

**Note:** The OpenStack flavour name should be specified and not the flavor ID or UUID.

## Configuring Multi-Queue in OpenStack

The following are steps to configure multiqueue property in Avi Cloud using CLI (Example cloud name is Default-Cloud):

**Note:** In SE-Group parameter `max_queues_per_vnic` is a boot-up property. This flag is a Service Engine Group boot-up property and if it is enabled or disabled; it will require all the Service Engines under that respective Service Engine Group to be rebooted.

```
[admin:avi-controller]: > configure cloud Default-Cloud
[admin:avi-controller]: cloud> openstack_configuration
[admin:avi-controller]: cloud:openstack_configuration> custom_se_image_properties
New object being created
[admin:avi-controller]: cloud:openstack_configuration:custom_se_image_properties> name hw_vif_multiqueue_enabled
[admin:avi-controller]: cloud:openstack_configuration:custom_se_image_properties> value true
[admin:avi-controller]: cloud:openstack_configuration:custom_se_image_properties> save
[admin:avi-controller]: cloud:openstack_configuration> save
[admin:avi-controller]: cloud> save
```

The following are the steps to remove the configuration:

```
[admin:avi-controller]: > configure cloud Default-Cloud
[admin:avi-controller]: cloud> openstack_configuration
[admin:avi-controller]: cloud:openstack_configuration> no custom_se_image_properties name hw_vif_multiqueue_enabled
Removed custom_se_image_properties with name=hw_vif_multiqueue_enabled
+-----+-----+
| Field          | Value          |
+-----+-----+
| username       | admin          |
| password       |                |
<sensitive>
| | admin_tenant | admin | | mgmt_network_name | public | | privilege | WRITE_ACCESS | | use_keystone_auth | True | |
</sensitive>
```