



Configuring Interfaces

- [Interface Configurations, on page 1](#)
- [Hardware Acceleration Support \(OpenStack Only\), on page 15](#)

Interface Configurations

The Interface configuration allows to choose various configuration for the interface including network, subnet, ip address, mac address, vim interface name, model, and so on.

This section describes these basic and advance interface configurations for Elastic Services Controller (ESC) and procedures to configure these.

Basic Interface Configurations

In ESC Datamodel, Interface refers to the VNIC attached to the VM. We can add one or more Interface under a VM Group. The interface section will have details to configure the VNIC.

This section describes basic interface configurations for Elastic Services Controller (ESC).

Configuring Basic Interface Settings

This section describes basic interface configurations, such as:

- Network
- Subnet
- IP address
- MAC address
- VIM interface name, and so on for Elastic Services Controller (ESC).

Configuring an Interface Name

To configure VIM interface name, specify attribute `<vim_interface_name>` for an interface in the Deployment XML file. Use `<vim_interface_name>` to use a specific name when generating an interface name. If these attribute is not specified, ESC will auto-generate an interface name, which is a combination of the

deployment_name, group_name, and a random UUID string. For example:
my-deployment-na_my-gro_0_8053d7gf-hyt33-4676-h9d4-9j4a5599472t.



Note This feature is currently supported only on OpenStack.

If the VM group is elastic and a `vim_interface_name` has been specified, a numeric index is added after the interface name for the second interface name onwards (the first one remains unchanged). For example, if the specified interface name is set as `<vim_interface_name>interface_1</vim_interface_name>` and scaling is set to 3, three VMs are created with three different interface name, `interface_1`, `interface_1_1`, and `interface_1_2`. If a VM group only has a single VM, then there is no "`<index>`" appended to the custom interface name. A single deployment can contain multiple VM groups, and each individual VM group can specify a different `vim_interface_name` value, if required. For example, a deployment could have two VM groups: the first group specifies a `vim_interface_name` and all VMs have their names generated as described above. The second VM group does not specify a `vim_interface_name`, therefore all VM names created from this group are auto generated. The same interface name can be used in separate interface sections within the same VM group, or in separate VM groups within a deployment, or in different deployments if required.

If attributes `<vim_interface_name>` or `<port>` are used for the same interface, the `vim_interface_name` value will be ignored and the value in the `port` attribute will be used.

```
<esc_datamodel xmlns="https://www.cisco.com/esc/esc"> <tenants><tenant>
<name>Admin</name>
<deployments>
<deployment>
<deployment_name>NwDepModel_nosvc</deployment_name>
<interface>
<nicid>0</nicid>
<vim_interface_name>interface_1</vim_interface_name>
<network>my-network</network>
</interface>
```



Note You can use a maximum of 61 characters for an interface name should not contain special characters and can only contain alphanumeric characters and "_" and "-". The following are some output samples with the custom port name. If the `vim_interface_name` was set during the deployment, the same value will be shown in the output. If this value was not set during the deployment, ESC will auto-generate the port name.

- Below is an example of the output operational data fetched using the `esc_nc_cli` script after adding a custom interface name. A new element called `vim_interface_name` will be shown under the interface element.

```
[admin@esc-3-1-xxx]$ esc_nc_cli --user <username> --password <password> get
esc_datamodel/opdata
<?xml version="1.0" encoding="UTF-8"?>
<rpc-reply xmlns="urn:ietf:params:xml:ns:netconf:base:1.0" message-id="1">
. . .
  <interface>
    <nicid>0</nicid>
    <type>virtual</type>
    <port_id>e4111069-5d00-493b-8ea9-1a2ca134b5c8</port_id>
    <vim_interface_name>interface_1</vim_interface_name>      <!-- NEW IN OUTPUT
-->
    <network>c7fafeca-aa53-4349-9b60-1f4b92605420</network>
```

```

    <subnet>255.255.255.0</subnet>
    <ip_address>192.168.2.1</ip_address>
    <mac_address>fa:16:3e:d7:5e:da</mac_address>
    <netmask>255.255.240.0</netmask>
    <gateway>192.168.2.255</gateway>
  </interface>

```

- Below is an example output operational data fetched using a REST API.

```

GET http://localhost:8080/ESCManager/v0/deployments/example-deployment-123
| xmllint --format -
<?xml version="1.0" encoding="UTF-8" standalone="yes"?>
<deployments>
  . . .
  <interface>
    <network_uuid>c7fafeca-aa53-4349-9b60-1f4b92605420</network_uuid>
    <gateway>172.16.0.1</gateway>
    <ip_address>172.16.12.251</ip_address>
    <mac_address>fa:16:3e:30:0c:99</mac_address>
    <netmask>255.255.240.0</netmask>
    <nic_id>0</nic_id>
    <port_forwarding/>
    <port_uuid>1773cdbf-fe5f-4af1-adff-3a9c1dd1c47d</port_uuid>
    <vim_interface_name>interface_1</vim_interface_name>      <!-- NEW IN OUTPUT
-->
    <security_groups/>
    <subnet_uuid>7b2ce63b-eb20-4ff8-8d49-e46ee8dde0f5</subnet_uuid>
    <type>virtual</type>
  </interface>

```

In all the above scenarios, if `vim_interface_name` is not specified in the `deployment.xml`, the output will still contain this element, however with an internally generated interface name. For example:

```

<vim_interface_name>vm-name-deployme_Grp1_1_0f24cd7e-cae7-402e-819a-5c84087103ba</vim_interface_name>

```

Assigning the MAC Address

ESC deployment on VMware vCenter supports assigning MAC address using the MAC address range, or MAC address list from the MAC address pool to deploy VMs to the network.

You can assign MAC address in the following ways:

Using the Interface

```

<interfaces>
  <interface>
    <nicid>1</nicid>
    <network>MANAGEMENT_NETWORK</network>
    <ip_address>172.16.0.11</ip_address>
    <mac_address>fa:16:3e:73:19:a0</mac_address>
  </interface>
</interfaces>

```

During scaling, you can assign the MAC address list or MAC address range from the MAC address pool.

```

<scaling>
  <min_active>2</min_active>
  <max_active>2</max_active>
  <elastic>true</elastic>
  <static_ip_address_pool>
    <network>MANAGEMENT_NETWORK</network>
    <ip_address>172.16.0.11</ip_address>
    <ip_address>172.16.0.12</ip_address>
  </static_ip_address_pool>

```

```

    <ip_address>172.16.0.13</ip_address>
  </static_ip_address_pool>
  <static_mac_address_pool>
    <network>MANAGEMENT_NETWORK</network>
    <mac_address>fa:16:3e:73:19:a0</mac_address>
    <mac_address>fa:16:3e:73:19:a1</mac_address>
    <mac_address>fa:16:3e:73:19:a2</mac_address>
  </static_mac_address_pool>
</scaling>

```

Assign MAC address using MAC address range.

```

<scaling>
  <min_active>2</min_active>
  <max_active>2</max_active>
  <elastic>true</elastic>
  <static_ip_address_pool>
    <network>MANAGEMENT_NETWORK</network>
    <ip_address_range>
      <start>172.16.0.25</start>
      <end>172.16.0.27</end>
    </ip_address_range>
  </static_ip_address_pool>
  <static_mac_address_pool>
    <network>MANAGEMENT_NETWORK</network>
    <mac_address_range>
      <start>fa:16:3e:73:19:b0</start>
      <end>fa:16:3e:73:19:b2</end>
    </mac_address_range>
  </static_mac_address_pool>
</scaling>

```



Note You cannot change the MAC or IP pool in an existing deployment, or during scaling (when min and max value are greater than 1) of VM instances in a service update.

In VMware vCenter, while assigning the MAC address, the server might override the specified value for "Generated" or "Assigned" if it does not fall in the right ranges or is determined to be a duplicate. Because of this, if ESC is unable to assign the MAC address the deployment fails.

Configuring Subnet for an Interface

Subnets can be passed through the datamodel. Subnet within interfaces can be specified in the Interface section of the Deployment XML file. If there is no subnet specified in the datamodel, ESC will let OpenStack select the subnet for interface creation and will use the subnet from the port created by OpenStack.

```

<interface>
  <nicid>0</nicid>
  <network>my-network</network>
  <subnet>my-subnet</subnet>
</interface>

```

The `no_gateway` attribute allows ESC to create a subnet with the gateway disabled. In the example below, the `no_gateway` attribute is set to true to create a subnet without gateway.

```

<networks>
  <network>
    <name>mgmt-net</name>
    <subnet>

```

```

<name>mgmt-net-subnet</name>
<ipversion>ipv4</ipversion>
<dhcp>false</dhcp>
<address>172.16.0.0</address>
<no_gateway>true</no_gateway><!-- DISABLE GATEWAY -->
<gateway>172.16.0.1</gateway>
<netmask>255.255.255.0</netmask>
</subnet>
</network>
</networks>

```

Configuring an Out-of-Band Port

ESC also allows you to attach an out-of-band port to a VNF. To do this, pass the UUID or the name of the port in the deployment request file while initiating a service request. For more information, see, Out-of-band Volumes section in the [Cisco Elastic Services Controller User Guide](#).



Note While undeploying or restoring a VNF, the ports attached to that VNF will only be detached and not deleted. ESC does not allow scaling while using out-of-band port for a VM group. You can configure only one instance of VM for the VM group. Updating the scaling value for a VM group, while using the out-of-band port is not allowed during a deployment update.

```

<esc_datamodel xmlns="https://www.cisco.com/esc/esc">
  <name>tenant</name>
  <deployments>
    <deployment>
      <name>depz</name>
      <vm_group>
        <name>g1</name>
        <image>Automation-Cirros-Image</image>
        <flavor>Automation-Cirros-Flavor</flavor>
        <bootup_time>100</bootup_time>
        <reboot_time>30</reboot_time>
        <recovery_wait_time>10</recovery_wait_time>
        <interfaces>
          <interface>
            <nicid>0</nicid>
            <port>057a1c22-722e-44da-845b-a193e02807f7</port>
            <network>my-network</network>
          </interface>
        </interfaces>
      </vm_group>
    </deployment>
  </deployments>
</esc_datamodel>

```

Dual Stack Support

A dual stack network allows you to assign multiple IP addresses. These multiple IP addresses can be assigned on different subnets to a given interface within a VNF deployment using ESC.

ESC supports the following for dual stack:

- Configuring the network and list of subnet
- Configuring the network and list of subnet and ip address
- Configuring the network and list of ip address (no subnet)

- Specifying the network and list of subnet/ip (same subnet but different ip)



Note Currently, ESC supports dual stack only on OpenStack. ESC supports end-to-end IPv6 for OpenStack deployments.

A new container element named `addresses` is added to the `Interface`. This container holds a list of address elements. An address element must have an `address_id` (key). The `subnet` and `fixed-ip` address fields are optional, but you must specify either one.

The container address is as follows:

```
container addresses {
  list address {
    key "address_id";
    leaf address_id {
      description "Id for the address in address list.";
      type uint16;
      mandatory true;
    }
    leaf subnet {
      description "Subnet name or uuid for allocating IP to this port";
      type types:escnetname;
    }
  }
  leaf ip_address {
    description "Static IP address for this specific subnet";
    type types:escipaddr;
    must ".../.../.../.../scaling/max_active = 1 or
count(.../.../.../.../scaling/static_ip_address_pool) > 0"
    {
      error-message "Static ip address pools must be configured when static ip addresses are
configured.";
    }
  }
}
```

Dual stack now supports KPI monitoring. A new child element `address_id` has been added to the `metric_collector` element. This accepts a value which points to an address within the specified `nicid` to be used for KPI monitoring. That is, it allows one of the addresses defined beneath an interface to be used for KPI monitoring.

```
...
<interface>
  <nicid>1</nicid>
  <network>demo-net</network>
  <addresses>
    <address>
      <address_id>0</address_id>
      <subnet>demo-subnet</subnet>
    </address>
  </addresses>
</interface>

  <kpi_data>
    <kpi>
      <event_name>VM_ALIVE</event_name>
      <metric_value>1</metric_value>
      <metric_cond>GT</metric_cond>
      <metric_type>UINT32</metric_type>
      <metric_occurrences_true>5</metric_occurrences_true>
      <metric_occurrences_false>5</metric_occurrences_false>
    </kpi>
  </kpi_data>
</interface>
```

```

    <metric_collector>
      <type>ICMPping</type>
      <nicid>1</nicid>
      <address_id>0</address_id>
      <poll_frequency>10</poll_frequency>
      <polling_unit>seconds</polling_unit>
      <continuous_alarm>>false</continuous_alarm>
    </metric_collector>
  </kpi>
</kpi_data>

```

...



Note The address_id under the metric_collector element must be the same as one of the address_id beneath the interface.

Dual stack interfaces can now be used in day-0 variable substitution. This means the ability to substitute the values from the multiple addresses defined under a single interface. Day 0 configuration is defined in the datamodel under the config_data tag.

In case of dual stack with multiple IP addresses, the variables are in the form NICID_<n>_<a>_<PROPERTY> where:

- <n> is the nicid for the interface.
- <a> is the address_id of an address within that interface.

The list of possible day-0 substitution variables from dual stack is:

NICID_n_a_IP_ALLOCATION_TYPE	string containing FIXED DHCP	ipv4 or ipv6
NICID_n_a_IP_ADDRESS	IP address	ipv4 or ipv6
NICID_n_a_GATEWAY	Gateway address	ipv4 or ipv6
NICID_n_a_CIDR_ADDRESS	CIDR prefix address	ipv4 or ipv6
NICID_n_a_CIDR_PREFIX	Integer with CIDR prefix-length	ipv4 or ipv6
NICID_n_a_NETMASK	If an ipv4 CIDR address and prefix are present, ESC will automatically calculate and populate the netmask variable. This is not substituted in the case of an IPv6 address and should not be used.	ipv4 only

For information on day-0 configuration for single IP address, see Day Zero Configuration chapter in the [Cisco Elastic Services Controller User Guide](#).

The template file defined in the config_data with day-0 configurations is as follows:

```

NICID_0_NETWORK_ID=${NICID_0_NETWORK_ID}
NICID_0_MAC_ADDRESS=${NICID_0_MAC_ADDRESS}

NICID_0_0_IP_ALLOCATION_TYPE=${NICID_0_0_IP_ALLOCATION_TYPE}
NICID_0_0_IP_ADDRESS=${NICID_0_0_IP_ADDRESS}
NICID_0_0_GATEWAY=${NICID_0_0_GATEWAY}
NICID_0_0_CIDR_ADDRESS=${NICID_0_0_CIDR_ADDRESS}

```

```

NICID_0_0_CIDR_PREFIX=${NICID_0_0_CIDR_PREFIX}
NICID_0_0_NETMASK=${NICID_0_0_NETMASK}

NICID_0_1_IP_ALLOCATION_TYPE=${NICID_0_1_IP_ALLOCATION_TYPE}
NICID_0_1_IP_ADDRESS=${NICID_0_1_IP_ADDRESS}
NICID_0_1_GATEWAY=${NICID_0_1_GATEWAY}
NICID_0_1_CIDR_ADDRESS=${NICID_0_1_CIDR_ADDRESS}
NICID_0_1_CIDR_PREFIX=${NICID_0_1_CIDR_PREFIX}

```

The datamodel is as follows:

```

<?xml version="1.0" encoding="ASCII"?>
<esc_datamodel xmlns="https://www.cisco.com/esc/esc">
  <tenants>
    <tenant>
      <name>dep-tenant</name>
      <deployments>
        <deployment>
          <name>cirros-dep</name>
          <vm_group>
            <name>Grp1</name>
            <bootup_time>600</bootup_time>
            <recovery_wait_time>30</recovery_wait_time>
            <flavor>Automation-Cirros-Flavor</flavor>
            <image>Automation-Cirros-Image</image>
          <interfaces>
            <interface>
              <!-- No dual stack support on mgmt interface in ESC 4.1 -->
              <nicid>0</nicid>
              <network>my-network</network>
            </interface>
            <interface>
              <nicid>1</nicid>
              <network>ent-network1</network>
              <addresses>
                <address>
                  <!-- IPv4 Dynamic -->
                  <address_id>0</address_id>
                  <subnet>v4-subnet_A</subnet>
                </address>
                <address>
                  <!-- IPv6 Dynamic -->
                  <address_id>1</address_id>
                  <subnet>v6-subnet_B</subnet>
                </address>
              </addresses>
            </interface>
            <interface>
              <nicid>2</nicid>
              <network>ent-network2</network>
              <addresses>
                <address>
                  <!-- IPv4 Static -->
                  <address_id>0</address_id>
                  <subnet>v4-subnet_C</subnet>
                  <ip_address>172.16.87.8</ip_address>
                </address>
                <address>
                  <!-- IPv6 Static -->
                  <address_id>1</address_id>
                  <subnet>v6-subnet_D</subnet>
                  <ip_address>fd07::110</ip_address>
                </address>
              </addresses>
            </interface>
          </interfaces>
        </deployment>
      </deployments>
    </tenant>
  </tenants>
</esc_datamodel>

```



```

</interface>
<interface>
  <nicid>3</nicid>
  <network>ent-network3</network>
  <addresses>
    <address>
      <!-- Only ip config - ipv6 but no subnet -->
      <address_id>0</address_id>
      <ip_address>fd07::110</ip_address>
    </address>
    <address>
      <!-- Only ip config - ipv4 but no subnet -->
      <address_id>1</address_id>
      <ip_address>172.16.88.9</ip_address>
    </address>
  </addresses>
</interface>
<interface>
  <nicid>4</nicid>
  <network>ent-network4</network>
  <addresses>
    <address>
      <!-- ipv4 same subnet as address_id 6 -->
      <address_id>0</address_id>
      <subnet>v4-subnet_F</subnet>
      <ip_address>172.16.86.10</ip_address>
    </address>
    <address>
      <!-- ipv4 same subnet as id 5 -->
      <address_id>1</address_id>
      <subnet>v4-subnet_F</subnet>
      <ip_address>172.16.86.11</ip_address>
    </address>
  </addresses>
</interface>
</interfaces>
<kpi_data>
...

```

After successful deployment using multiple IPs, ESC provides a list of addresses as notification, or opdata.

A list of multiple <address> elements under the parent <interface> element containing the following:

- **address_id**—the address id specified in the input XML
- **subnet element**—subnet name or uuid
- **ip_address element**—the port's assigned IP on that subnet
- **prefix**—the subnet CIDR prefix
- **gateway**—the subnet gateway address
- ESC Static IP support

Notification:

```

<vm_id>1834124d-b70b-41b9-9e53-fb55d7c901f0</vm_id>
<name>jenkins-gr_g1_0_e8bc9a81-4b9a-437a-807a-f1a9bbc2ea3e</name>
<generated_name>custom_vim_name
<host_id>dc380f1721255e2a7ea15932c1a7abc681816642f75276c166b4fe50</host_id>

<hostname>my-server</hostname>

```

```

<interfaces>
  <interface>
    <nicid>0</nicid>
    <type>virtual</type>
    <vim_interface_name>custom_vim_name
    <port_id>4d57d4a5-3150-455a-ad39-c32fffb10b1</port_id>
    <mac_address>fa:16:3e:d2:50:a5</mac_address>
    <network>45638651-2e92-45fb-96ce-9efdd9ea343e</network>
    <address>
      <address_id>0</address_id>
      <subnet>6ac36430-4f58-454b-9dc1-82f7a796e2ff</subnet>
      <ip_address>172.16.0.22</ip_address>
      <prefix>24</prefix>
      <gateway>172.16.0.1</gateway>
    </address>
    <address>
      <address_id>1</address_id>
      <subnet>8dd9f501-19d4-4782-8335-9aa9fbd4dab9</subnet>
      <ip_address>2002:dc7::4</ip_address>
      <prefix>48</prefix>
      <gateway>2002:dc7::1</gateway>
    </address>
    <address>
      <address_id>2</address_id>
      <subnet>a234501-19d4-4782-8335-9aa9fbd4caf6</subnet>
      <ip_address>172.16.87.8</ip_address>
      <prefix>20</prefix>
      <gateway>172.16.87.1</gateway>
    </address>
  </interface>
</interfaces>

```

Sample opdata:

```

<interfaces>
  <interface>
    <nicid>0</nicid>
    <type>virtual</type>
    <vim_interface_name>custom_vim_name
    <port_id>4d57d4a5-3150-455a-ad39-c32fffb10b1</port_id>
    <mac_address>fa:16:3e:d2:50:a5</mac_address>
    <network>45638651-2e92-45fb-96ce-9efdd9ea343e</network>
    <address>
      <address_id>0</address_id>
      <subnet>6ac36430-4f58-454b-9dc1-82f7a796e2ff</subnet>
      <ip_address>172.16.0.22</ip_address>
      <prefix>24</prefix>
      <gateway>172.16.0.1</gateway>
    </address>
    <address>
      <address_id>1</address_id>
      <subnet>8dd9f501-19d4-4782-8335-9aa9fbd4dab9</subnet>
      <ip_address>2002:dc7::4</ip_address>
      <prefix>48</prefix>
      <gateway>2002:dc7::1</gateway>
    </address>
  </interface>
</interfaces>

```

You can also see that the day-0 substitution values are replaced in the output data. Sample output data with the values populated in the day-0 configuration is as follows:

```

NICID_0_NETWORK_ID=45638651-2e92-45fb-96ce-9efdd9ea343e
NICID_0_MAC_ADDRESS=fa:16:3e:d2:50:a5

```

```

NICID_0_0_IP_ALLOCATION_TYPE=DHCP
NICID_0_0_IP_ADDRESS=172.16.0.22
NICID_0_0_GATEWAY=172.16.0.1
NICID_0_0_CIDR_ADDRESS=172.16.0.0
NICID_0_0_CIDR_PREFIX=24
NICID_0_0_NETMASK=255.255.255.0

NICID_0_1_IP_ALLOCATION_TYPE=DHCP
NICID_0_1_IP_ADDRESS=2002:dc7::4
NICID_0_1_GATEWAY=2002:dc7::1
NICID_0_1_CIDR_ADDRESS=2002:dc7::/48
NICID_0_1_CIDR_PREFIX=48

```

Dual Stack with Static IP Support

ESC supports dual stack with static IP support. As part of the initial configuration the user can provide the subnet and IP to be configured.



Note ESC supports static IP only when the scaling is false or minimum /maximum =1.

When you create a VM with out-of-band network, and specify a list of subnets with static IP (the network has multiple subnets), then ESC applies both subnet and the corresponding static IP.

In the example below, two subnets (ipv4 and ipv6) are added to a single interface.

```

<?xml version="1.0" encoding="ASCII"?>
<esc_datamodel xmlns="https://www.cisco.com/esc/esc">
  <tenants>
    <tenant>
      <name>dep-tenant</name>
      <deployments>
        <deployment>
          <name>cirros-dep</name>
          <vm_group>
            <name>Grp1</name>
            <bootup_time>600</bootup_time>
            <recovery_wait_time>30</recovery_wait_time>
            <flavor>Automation-Cirros-Flavor</flavor>
            <image>Automation-Cirros-Image</image>
            <interfaces>
              <interface>
                <nicid>0</nicid>
                <network>ent-network2</network>
                <addresses>
                  <address>
                    <!-- IPv4 Static -->
                    <address_id>0</address_id>
                    <subnet>v4-subnet_C</subnet>
                    <ip_address>172.16.87.8</ip_address>
                  </address>
                  <address>
                    <!-- IPv6 Static -->
                    <address_id>1</address_id>
                    <subnet>v6-subnet_D</subnet>
                    <ip_address>fd07::110</ip_address>
                  </address>
                </addresses>
              </interface>
            </interfaces>
          </vm_group>
        </deployment>
      </deployments>
    </tenant>
  </tenants>
</esc_datamodel>

```

```

    </interfaces>
    <kpi_data>

```

For information on deploying VNFs, see [Deploying Virtual Network Functions on OpenStack](#).

Advanced Interface Configurations

This section describes several interface configurations for Elastic Services Controller (ESC) and the procedure to configure the hardware interfaces.

For information on basic interface settings, see [Basic Interface Configurations](#).

Configuring Advance Interface Settings

Configuring SR-IOV in ESC

Single Root I/O Virtualization (SR-IOV) allows multiple VMs running a variety of guest operating systems to share a single PCIe network adapter within a host server. It also allows a VM to move data directly to and from the network adapter, bypassing the hypervisor for increased network throughput and lower server CPU burden.

Configuring SR-IOV in ESC for OpenStack

Before you configure SR-IOV in ESC for OpenStack, configure the hardware and OpenStack with the correct parameters.

To enable SR-IOV in ESC for OpenStack, specify the interface `type` as `direct`. The following snippet shows a sample datamodel:

```

<interfaces>
  <interface>
    <nicid>0</nicid>
    <network>my-network</network>
    <type>direct</type>
  </interface>
</interfaces>
...

```

Configuring SR-IOV in ESC for VMware

Before you configure SR-IOV in ESC for VMware, consider the following:

- Enable SR-IOV Physical Functions on desired ESXi hosts. For more information, see [VMware documentation](#).
- Consider the following important points before enabling SR-IOV:
 - Review the list of physical network adaptors that VMware supports for SR-IOV. See [VMware documentation](#).
 - Review the list of VM features that are not supported on a VM with SR-IOV configured. See [VMware documentation](#).
 - In a cluster deployment (defined by "zone" in the datamodel) with SR-IOV, make sure that each ESXi host has identical Physical Functions enabled for SR-IOV selection. For example, if a VM is going to use `vmnic7` as the Physical Function, make sure that each host has `vmnic7` and SR-IOV status for each `vmnic7` is enabled.

To enable SR-IOV in ESC for VMware, specify `interface<type>` as `direct` and also extension `<name>` as `sriov_pf_selection` in the deployment datamodel. Interface Type `direct` indicates an SR-IOV device and extension name `sriov_pf_selection` indicates the physical function. The following snippet shows a sample datamodel:

```
<vm_group>
...
<interface>
  <nicid>2</nicid>
  <network>MgtNetwork</network>
  <type>direct</type>
</interface>
<interface>
  <nicid>3</nicid>
  <network>MgtNetwork</network>
  <type>direct</type>
</interface>
...
<extensions>
  <extension>
    <name>sriov_pf_selection</name>
    <properties>
      <property>
        <name>nicid-2</name>
        <value>vmnic1,vmnic2</value>
      </property>
      <property>
        <name>nicid-3</name>
        <value>vmnic3,vmnic4</value>
      </property>
    </properties>
  </extension>
</extensions>
</vm_group>
```

Configuring Allowed Address Pair

Cisco Elastic Services Controller allows you to specify the address pairs in the deployment datamodel to pass through a specified port regardless of the subnet associated with the network.

The address pair is configured in the following ways:

- **List of Network**—When a list of network is provided on a particular interface, ESC will get the subnet details from the OpenStack for these networks and add them to the corresponding port or interface. The following example explains how to configure address pairs as a list of network:

```
<interface>
  <nicid>1</nicid>
  <network>network1</network>
  <allowed_address_pairs>
    <network>
      <name>bb8c5cfb-921c-46ea-a95d-59feda61cac1</name>
    </network>
    <network>
      <name>6ae017d0-50c3-4225-be10-30e4e5c5e8e3</name>
    </network>
  </allowed_address_pairs>
</interface>
</interfaces>
```

- **List of Address**— When a list of address is provided, ESC will add these addresses to the corresponding interface. The following example explains how to configure address pairs as a list of address:

```
<interface>
    <nicid>0</nicid>
    <network>esc-net</network>
    <allowed_address_pairs>
        <address>
            <ip_address>10.10.10.10</ip_address>
            <netmask>255.255.255.0</netmask>
        </address>
        <address>
            <ip_address>10.10.20.10</ip_address>
            <netmask>255.255.255.0</netmask>
        </address>
    </allowed_address_pairs>
</interface>
```

Configuring Security Group Rules

Cisco Elastic Services Controller (ESC) allows you to associate security group rules to the deployed instances on OpenStack. These security group rules are configured by specifying the necessary parameters in the deployment datamodel. In addition to configuring security group rules, if any VNF instance fails, ESC recovers the instance and applies the security group rules for the redeployed VNF.

To configure security group rules, do the following:

Before you begin

- Make sure you have created a tenant through ESC.
- Make sure you have security groups created.
- Make sure you have the security group name or UUID.

Step 1 Log in to the ESC VM as a root user.

Step 2 Run the following command to check the UUIDs of a given security group:

```
nova --os-tenant-name <NameOfTheTenant> secgroup-list
```

Step 3 Pass the following arguments in the deployment data model:

```
<interfaces>
  <interface>
    <nicid>0</nicid>
    <network>my-network</network><!-- depends on network name -->
    <security_groups>
      <security_group>0c703474-2692-4e84-94b9-c29e439848b8</security_group>
      <security_group>bbcd62-a0de-4475-b258-740bfd33861b</security_group>
    </security_groups>
  </interface>
  <interface>
    <nicid>1</nicid>
    <network>sample_VmGrpNet</network><!--depends on network name -->
    <security_groups>
      <security_group>sample_test_SQL</security_group>
```

```
</security_groups>
</interface>
```

Step 4 Run the following command to verify whether the security groups are associated with the VM instance:

```
nova --os-tenant-name <NameOfTenant> show <NameOfVMInstance>
```

Hardware Acceleration Support (OpenStack Only)

You can configure hardware acceleration features on OpenStack using the *flavor data model*. The following hardware acceleration features can be configured:

- **vCPU Pinning**—enables binding and unbinding of a process to a vCPU (Virtual Central Processing Unit) or a range of CPUs, so that the process executes only on the designated CPU or CPUs rather than any CPU.
- **OpenStack performance optimization for large pages and non-uniform memory access (NUMA)**—enables improvement of system performance for large pages and NUMA i.e., system's ability to accept higher load and modify the system to handle a higher load.
- **OpenStack support for PCIe Passthrough interface**—enables assigning a PCI device to an instance on OpenStack.

The following example explains how to configure hardware acceleration features using *flavor data model*:

```
$ cat example.xml
<?xml version='1.0' encoding='ASCII'?>
<esc_datamodel xmlns="http://www.cisco.com/esc/esc">
  <flavors>
    <flavor>
      <name>testfl6</name>
      <vcpus>1</vcpus>
      <memory_mb>2048</memory_mb>
      <root_disk_mb>10240</root_disk_mb>
      <ephemeral_disk_mb>0</ephemeral_disk_mb>
      <swap_disk_mb>0</swap_disk_mb>
      <properties>
        <property>
          <name>pci_passthrough:alias</name>
          <value>nic1g:1</value>
        </property>
      </properties>
    </flavor>
  </flavors>
</esc_datamodel>
$ /opt/cisco/esc/esc-confd/esc-cli/esc_nc_cli --user <username> --password <password>
edit-config ./example.xml
```

Creating Additional Parameters for VMware vSphere NUMA Attributes

ESC enhances NUMA for VMware vSphere by adding additional configuration parameters.

This enhancement adds the additional or advanced configuration for VMware vSphere as a prefix to pass configuration parameters instead of passing these values through the day-0 configuration files.

Prefix: extConfigParam

Example:

```
<configuration>
  <dst>extConfigParam:mgmt-ipv4-addr</dst>
  <data>$NICID_1_IP_ADDRESS/16</data>
</configuration>
```

The additional configuration helps to minimize the data model changes, and restrict configuration changes to the VIM layer.

Configuring PCI or PCIe Device Passthrough on VMware vCenter

ESC supports VMware vCenter PCI or PCIe device passthrough (VMDirectPath I/O). This enables VM access to physical PCI functions on platforms with an I/O memory management unit.

Before You Begin

For the PCI / PCIe devices of a host VM to enable passthrough, the vSphere administrator must mark these devices in the vCenter.



Note You must reboot the host after PCI settings. Put the host to maintenance mode, power off or migrate all VMs to other hosts.

To specify PCI device passthrough request in ESC deployments, include the `<type>` attribute with value set to *passthru*. To specify the PCI device to be selected for a particular `vm_group` or network, include the *pci_id*. The data model is as follows:

```
<tenants>
  <tenant>
    <name>admin</name>
    <deployments>
      <deployment>
        <name>test</name>

        <vm_group>
          <name>test-g1</name>
          <image>uLinux</image>
          <bootup_time>300</bootup_time>
          <recovery_wait_time>10</recovery_wait_time>
          <interfaces>
            <interface>
              <nicid>1</nicid>
              <network>MgtNetwork</network>
              <ip_address>192.168.0.102</ip_address>
            </interface>
            <interface>
              <nicid>2</nicid>
              <network>VM Network</network>
              <type>passthru</type>
              <ip_address>172.16.0.0</ip_address>
            </interface>
            <interface>
              <nicid>3</nicid>
              <network>VM Network</network>
              <type>passthru</type>
              <ip_address>192.168.46.117</ip_address>
          </interfaces>
        </vm_group>
      </deployment>
    </deployments>
  </tenant>
</tenants>
```



```
</interface>
  <interface>
    <nicid>3</nicid>
    <type>passthru</type>
    <network>MgtNetwork</network>
    <pci_id>0000:07:10.3</pci_id>
  </interface>
</interfaces>
```

After successful deployment, the *passthru* value is set in the interface section of the notification as well as in the operational data.

Auto Selecting PCI or PCIe PassThrough Device

ESC needs one or more PCI or PCIe passthrough devices to be attached to each deployment without a particular PCI ID. ESC first selects a host. ESC selects the next available PCI or PCIe passthrough enabled device and attaches it during the deployment. If there is no PCI or PCIe passthrough enabled device available, ESC fails the deployment. The vSphere administrator has to ensure all computing-host within the target computing-cluster have enough number of PCI or PCIe passthrough enabled devices.



Note

- PCI or PCIe passthrough is not considered by ESC placement algorithm. For example, ESC does not select a host because it has available resources to complete the PCI or PCIe passthrough requests.
 - ESC selects the PCI or PCIe passthrough device randomly. ESC does not consider the type or specification of the device. It selects the next available PCI or PCIe device from the list.
 - Recovery fails if the VNF is recovered to a computing-host that ESC has selected based on the ESC placement algorithm, and if that computing-host does not have any PCI or PCIe passthrough enabled devices available.
 - DRS must be turned off for the passthrough to work.
-

