Managing Inventory

For Cisco VTS to manage the network entities, they have to be present in the Cisco VTS inventory. You need to discover the network entities in the network, and add these to the inventory.

You can discover these entities using the Auto Discovery option using a seed IP, and import the details into Cisco VTS inventory. You can also manually create a CSV file with the details, in a prescribed format, and import it into the Cisco VTS inventory.

Note For vCenter-based setups, Cisco VTS supports only discovery using the CSV option. Auto Discovery using seed IP is not supported for vCenter-based setups.

The discovery process discovers the new devices, fabric connections, and the host (including host interfaces).

The discovery framework displays the difference between the current inventory, and the discovered content. With this enhancement, after you discover the devices using the CSV file or Auto Discovery option, you can view the changes in the network, and compare it with the existing inventory, and accept the changes or make edits as required.

Cisco VTS supports secure device access and communicates with the device using a secure channel. This is the default behavior.

Note Cisco VTS device discovery is performed over secure ports and protocols. You must make sure that the Nexus OS devices are reachable through HTTPS (443). In Nexus 7000 series devices, https is disabled by default. You must make sure it is enabled on port 443.
To enable secure communication for IOS-XR devices over SSH, you need to have the SSH enabled on the devices. Day Zero Configuration for Cisco ASR 9000 has to be updated to support this. (See Day Zero Configuration Examples document for details).

This chapter has the following sections:

Creating Authorization Groups

Authorization Group is used by Cisco VTS to authenticate or to log in to the device.

You can create authorization groups and assign devices you import into Cisco VTS, to these groups. Authorization groups are used to group devices with the same credentials (that is, usernames and passphrases). Once the authorization groups are created, all the devices under these groups may be accessed without specifying the credentials every time they are accessed.

If the same credential are used for accessing all devices, one authorization group can be used. If the credentials are different for different devices, multiple authorization-groups (as many as username/passphrase pairs used by devices) need to be created.

When you do a manual import of devices, the CSV file that is used to import inventory details links the authorization group with a specific device. The applicable authorization group should be used for corresponding device entry in the CSV file.

Note Changing the VTS UI password on first time log in does not update the vts-default authgroup password. To sync vts-default password with VTS UI, change the password of vts-default authgroup after you change the password for VTS UI initially. You must do this before you import devices into the inventory, using the vts-default authgroup.

To create an authorization group:

Step 1 Go to **Inventory** > **Authorization Group**. The Inventory / Authorization Group window appears.

Step 2 Click **Add (+)** icon. The Add Auth Group popup window appears.
Enter the following details, and click **Save**:

- **Auth Group Name—**The authorization group name.
- **Controller User Name—**This is the VTC administrative user name.
- **Device User Name—**This is the login user name for the device.
- **Passphrase—**This is the login passphrase for the device.

The authorization group gets added to the Groups table.

To edit an authorization group, select the Auth Group Name check box and click the **Edit** icon.

To delete an authorization group, select the Auth Group Name check box and click the **delete (X)** icon.

Importing Inventory using CSV File

You can manually create a CSV file with device details, in a prescribed format, and import the CSV file into Cisco VTS.

The CSV file is used to define device mappings. If the format is incorrect, Cisco VTS displays an error and provides the details of the error. After a successful import, the topology gets displayed based on the mapping specified in the file.

Note You should be an admin user to download or upload the CSV file. Also, if you are uploading a CSV file for the first time and there are issues uploading the file, then only the partial information is uploaded. You may encounter problems due to the partial upload.

To download a sample inventory file, click **Download latest CSV Template** . You can use the **Export Inventory** option to export the current inventory details in CSV format, for reuse.

The CSV file has the following fields:

- device-name—The device host-name (leaf, spine, DCI)
- device-ip—IP address for the device (leaf, spine, DCI)
- device-platform—Can be Cisco Nexus 9000, Cisco Nexus 7000 etc based on the device that is part of the network.
- device-role—The role that a particular device plays in the data center.
    - leaf—If the device plays the role of a Leaf in the data center.
    - border leaf—If the device plays the role of a Border Leaf in the data center.
    - spine—If the device plays the role of a Spine in the data center.
    - spine-rr—If the Spine plays the role of a Route Reflector in the data center.
    - dci—If the device plays the role of a DCI in the data center.
- group-tag—Identifier for the group.
- port-name—Physical port connectivity (local interface)
- connection-type—server (if connected to compute host); fabric (if connected to another leaf, spine, DCI devices).
- server-id—Host-name or IP address of the connected device based upon what is configured on the actual host. If you enter hostname, ensure that it contains hostname in FQDN format, i.e <hostname>.<domain>.
- server-type—virtual-server for computes; baremetal for connections to spine, DCI.
- interface-name—Physical port connectivity (interface of the connected device)
- server-ip— IP address of the connected device.

- auth-group—Authorization group name, created as part of initialization, with correct credentials.
- sriov-enabled— If the interface (*interface-name*) is SR-IOV enabled, this has to be TRUE.
- physnet-name— Physnet name associated with the interface (*interface-name*) in OpenStack. If *sriov enabled* is TRUE, it is the Physnet to be used for SR-IOV. If it is FALSE, the other possibilities are that the port is associated to L2 switch or OVS. In case of OVS, you need to give Physnet intended to be used for OVS.
- bgp-asn— BGP ASN number.
- underlay-loopback-num— Underlay loopback number.
- overlay-loopback-num— Overlay loopback number.

Note In a VMware environment, each time you add a leaf, you must create a corresponding VMware vSphere Distributed Switch (vDS). See the Notes Regarding VMware vSphere Distributed Switch section for details.

Note While importing inventory with IPv6 addresses for compute hosts in vCenter, the host labels in vCenter (if they have IPv4 addresses) need to be changed. In order to change them, you need to disconnect the host in vCenter, add the host back to the Datastore with IPv6 address.

The CSV file should always have the columns for bgp-asn, underlay-loopback-num, and overlay-loopback-num, in that order from left to right, and adjacent to each other. If the bgn-asn column is not adjacent to the underlay-loopback-num column, all bgp-asn values provided in the CSV file will not show up after you import the file. Also, if this order is not followed in the CSV file, the values will be mixed up in the inventory upon CSV import. That is, if the order in the file is underlay-loopback-num, overlay-loopback-num, and bgp-asn, from left to right, then upon CSV import the bgp-asn value is taken as overlay loopback number, underlay loopback number is taken as bgp-asn, and overlay loopback number is taken as underlay loopback number.

Note These three fields are optional in inventory CSV file. Only when you decide to place them in inventory CSV, the order specific above has to be followed.

Step 1 Go to **Inventory** > **Import and Discovery**. The Inventory > Import and Discovery window appears.

Step 2 Select the **CSV** radio button.

Step 3 Click **Import CSV** to choose the CSV file. Browse for the CSV file, and click **Open**.
A summary of the data obtained from the CSV file is displayed as a matrix. If data already exists in the inventory, Cisco VTS compares it with the data you had provided in the CSV file and displays it in the summary. The Devices, Fabric Connections, and Hosts (including host interfaces) present in the CSV file are displayed in the following buckets in the matrix. If there is no data in the inventory, the summary displays everything as new.

- New—Shows the new devices, fabric connection, and hosts included in the CSV.
- Mis Matched—Shows the mismatch between the uploaded CSV and the existing inventory. You can see the new and existing values for each of the entities, in this view.
- Existing—Shows all existing devices, fabric connections, and hosts in the inventory, and also present in the CSV.

- Missing—Shows the devices that are in the inventory, but not present in the CSV. Missing devices will be removed from current inventory when you update the inventory.

The following details are displayed for Devices:

- Device Name—The green icon near the Device Name indicates that the device is accessed via a secure channel.
- Admin State
- IP Address
- Auth Group
- Device Platform
- Device Role
- Group Tag
- Templates Attached
- Sync
- Last Sync Operation

The following details are displayed for Fabric Connection:

- Target Device Name
- Device Type
- Target Device Interface
- Target Device IP
- Source Device Interface
- Connection ID

The following details are displayed for Hosts:

- Host Name
- Host Type
- Host IP Address
- Associated VMM
- Virtual Switch

Click the drop-down to view the Host Interface details pertaining to each bucket.

Only for new devices, you can use the **Bulk Edit** option to update BGP-ASN and Loopback Interface Number. You can also use the **Bulk Edit** option to disable secure communication. By default, this is enabled in Cisco VTS.

Step 4 Click **Update Inventory**, and confirm that you need to update the inventory. Based on what is uploaded from CSV, the entire Inventory get replaced.

Step 5 After the inventory is replaced successfully, you can choose the following options to add/update device.

- Network Inventory

- Host Inventory

Performing Auto Discovery

In the auto discovery option, Cisco VTS automatically discovers the network topology in the data center. You can modify the device details after discovery is complete and add details to the inventory.

After the VTS admin user provides the Seed device IP and credentials, upon completion of discovery, Cisco VTS displays the discovered data in a matrix that has the following buckets—New, Modified, Missing, Existing.

The auto discovery option has the following prerequisites:

- Link Layer Discovery Protocol (LLDP) has to be enabled on leafs, spine, DCI, and computes. See documentation for the respective devices for details about how to enable LLDP on these devices.
- Enable lldpd on computes. See Enabling lldpd on Computes, on page 7 for details.

Note As part of Topology discovery, once the compute hosts have been discovered using LLDP, you need to add the username and passphrase to each host entry. This update is required for installation of the host-agent (in case of OpenStack) and any subsequent passphrase change via VTS GUI to go through.

- A seed device has to be identified, and the IP should be provided. The seed IP is that of one of the leaf or spine devices.

Note You can provide an IPv6 or IPv4 address. If an IPv6 address is given, preference is given to IPv6 address in cases where the devices have both IPv4 and IPv6 addresses, and the IPv6 address will be displayed upon completion of discovery.

- All devices must have a common set of credentials. These credentials will be used during the discovery process. See Managing Inventory, on page 1 for more information. The credentials must be of the appropriate privilege level on the devices.

To perform auto discovery:

Step 1 Go to **Inventory** > **Import and Discovery**. The Inventory / Discovery window appears.
Step 2 Enter the **Seed Device IP**.
Step 3 Enter the **Seed Device User Name**.
Step 4 Enter the **Seed Device Passphrase**.
Step 5 Click **Discover**.

After the discovery is complete, the details are displayed in the matrix in the following buckets.

Step 6 Click the desired cell for respective details to be populated in the screen. You may review the details, make changes wherever applicable, and click the **Add to Inventory** button to add the details into the Cisco VTS inventory. See Working with Discovered Data, on page 8 for detailed information about the how to work with the discovered values.

Enabling lldpd on Computes

You can install and configure lldpd on computes using an Ansible script. You may also manually install and configure lldpd on the computes. The following sections give details.

Note This procedure is to be used in a non-OSPD OpenStack installation. However, for OSPD deployments where computes are already configured, the following procedures can be used to install and configure lldpd on the computes.

Enabling lldpd Using Anisble

To enable lldpd on computes:

Step 1 Set export ANSIBLE_HOST_KEY_CHECKING=False on the VM from which Ansible script should be run.

Step 2 Run Ansible script packaging/debian/vts-vtep/opt/vts/lib/ansible/playbooks/lldp_configure/lldpd_configure_port_desc.yaml.

```
ansible-playbook -i inventory_file lldpd_configure_port_desc.yaml
```

Inventory file should have host details on which lldpd needs to be installed. Multiple hostnames can be separated by a new line.

A sample inventory file is given below:

```
#SSH details of computes on which lldpd needs to be installed and configured
[all]
#<hostname1> ansible_ssh_host=<ip> ansible_connection=ssh ansible_ssh_user=<username>
ansible_ssh_pass=<password>
compute-abc ansible_ssh_host="1.1.1.1" ansible_connection=ssh ansible_ssh_user=root
ansible_ssh_pass=abc

#Details to get LLDPD and configure rpm
[all:vars]
LLDPD_URL="http://download.opensuse.org/repositories/home:/vbernat/RHEL_7/src/lldpd-0.9.8-1.1.src.rpm"
VTS_LLDPD_CONFIGURE_RPM="http://engci-maven-master.cisco.com/artifactory/vts-yum/vts-lldpd-configure/2.0/noarch/vts-lldpd-configure-2-0.noarch.rpm
```

Enabling lldpd Manually

When you enable lldpd manually, you must ensure that you do the following on each compute.

Step 1 Uninstall lldpad on hosts.

```
yum -y remove lldpad

killall lldpad
```

Step 2 wget http://download.opensuse.org/repositories/home:/vbernat/RHEL_7/src/lldpd-0.9.8-1.1.src.rpm --directory-prefix=/etc/yum.repos.d/

Step 3 yum -y install lldpd

Step 4 Start lldpd deamon process.

```
lldpd
```

Step 5 wget vts-lldpd configure rpm from artifactory to configure sriov port information wget http://engci-maven-master.cisco.com/artifactory/vts-yum/vts-lldpd-configure/2.0/noarch/vts-lldpd-configure-2-0.noarch.rp…

Step 6 Install the rpm.

```
rpm –ivh vts-lldpd-configure-2-0.noarch.rpm
```

Working with Discovered Data

Upon completion of discovery, the discovered details about the Devices, Fabric Connections, and Hosts are displayed as a matrix. It displays data in the following buckets. You can click each button, view the details that get displayed in the respective screens, and, wherever Cisco VTS allows edits, change the values. The tables below gives detailed information about the discovered values in each bucket for Devices, Fabric Connections, and Hosts, and specifies whether edit option (including Bulk Edit option) is available. Make sure you also review the Important Notes, on page 13 before you update the inventory.

- **New—**The new devices, fabric connections and host (including host interfaces) discovered.

  The following table gives details of the values that are discovered and editable for **New Devices**:

Values	Discovered	Notes
Device Name	Yes	
Device IP	Yes	Update this with a new value, or retain the discovered data.
Auth Group	No	Select the desired value from the drop-down. Can be edited using Bulk Edit option too.
Device Platform	Yes	Update this with a new value from the drop-down, or retain the discovered data. Can be edited using Bulk Edit option too.

Values	Discovered	Notes
Device Role	No	Select the desired value from the drop-down. Can be edited using Bulk Edit option too.
Group Tag	No	Enter the Group tag value in the text box. Can be edited using Bulk Edit option too.
BGP ASN	No	Enter the ASN value in the text box. Can be edited using Bulk Edit option too.
Underlay Loopback Interface Num	No	Enter the loopback int num in the text box. Can be edited using Bulk Edit option too
Overlay Loopback Interface Num	No	Enter the loopback int num in the text box. Can be edited using Bulk Edit option too.

The following table gives details of the values that are discovered for **New Fabric Connections**.

Note No edits allowed under these values. You can add to inventory, and then perform edits as required.

Values	Discovered	Notes
Source Device Name	Yes	
Source Device Interface	Yes	
Target Device Name	Yes	You can only choose the device discovered. Will be blank if Target Device Type is FEX.
Target Device Interface	Yes	You can only choose the interface that is discovered.
Target Device Type	Yes	Possible values are baremetal and fex.
Target Device IP Address	Yes	You cannot change this value. Also, not visible in the UI.

The following table gives details of the values that are discovered for **New Hosts**.

Note For Hosts and Host Interfaces you can use the *Unmanaged* checkbox to have Cisco VTS not manage that host or host interfaces.

Values	Discovered	Notes
Host Name	Yes	This should typically contain the hostname in FQDN format, that is, <hostname>.<domain>.
Host IP	Yes	Retain the discovered data or update it with a new value
Associated VMM	No	Select the desired VMM from drop-down list of registered VMMs. Can be edited using Bulk Edit option too.
Virtual Switch	No	Select from the drop-down list of supported virtual-switch types, based on VMM type. Can be edited using Bulk Edit option too.

.
The following table gives details of the values that are discovered for Host Interfaces for the new Hosts. You need to click the > for a host icon to see the Host Interface details.

Note If you do not want to add a host interface to the inventory, click Do not add to Inventory.

Values	Discovered	Notes
Host Interface	Yes	
SRIOV-Enabled	Yes	You can only choose the discovered data. Edit option is not available.
Physnet	Yes	You can only choose the discovered data. Edit option is not available.
Attached Device	Yes	You can only choose the discovered data. Edit option is not available.

Values	Discovered	Notes
Device Interface	Yes	You can only choose the discovered data. Edit option is not available.

- **Mis Matched—**The number of mismatched devices, fabric connections, and hosts between the ones that are discovered from the network and the ones that are existing in the inventory. For mismatch bucket, edit option is not available for values that are not discovered. You can only accept the value from existing inventory, for those entities. You can edit the discovered content. You have the option to accept what is discovered or what is existing in the inventory. Once the values are updated to inventory, you can proceed to modify all fields as necessary.

  The following table gives details about mis matches in values discovered for Devices:

Value	Discovered	Notes
Device Name	Yes	
Device IP	Yes	You can choose the existing value or update it with the discovered value.
Auth Group	No	Reconciled with existing in inventory. Can be edited after adding to inventory.
Device Platform	Yes	You can choose the existing value or update it with the discovered value.
Device Role	No	Reconciled with existing in inventory. Can be edited after adding to inventory.
BGP ASN	No	Reconciled with existing in inventory. Can be edited after adding to inventory.
Underlay Loopback Interface Num	No	Reconciled with existing in inventory. Can be edited after adding to inventory.
Overlay Loopback Interface Num	No	Reconciled with existing in inventory. Can be edited after adding to inventory.

The following table gives details about mis matches in values discovered for Fabric Connections.

Value	Discovered	Notes
Source Device Name	Yes	
Source Device Interface	Yes	You can choose the existing value or update it with the discovered value.
Target Device Name	Yes	You can choose the existing value or update it with the discovered value.
Target Device Interface	Yes	You can choose the existing value or update it with the discovered value.

The following table gives details about mis matches in values discovered for Hosts.

Value	Discovered	Notes
Host Name	Yes	This should typically contain the hostname in FQDN format, that is, <hostname>.<domain>.
Host IP	Yes	You can choose the existing value or update it with the discovered value.
Associated VMM	No	Reconciled with existing in inventory. Can be edited after adding to inventory.
Virtual Switch	No	Reconciled with existing in inventory. Can be edited after adding to inventory.

The following table gives details about mis matches in values discovered for Hosts Interfaces:

Value	Discovered	Notes
Host Interface	Yes	
SRIOV-Enabled	Yes	You can choose the existing value or update it with the discovered value.
Physnet	Yes	You can choose the existing value or update it with the discovered value.

Value	Discovered	Notes
Attached Device	Yes	You can choose the existing value or update it with the discovered value.
Device Interface	Yes	You can choose the existing value or update it with the discovered value.

- **Missing—**The number of devices, fabric connections, and hosts that are existing in the inventory, but not discovered in the current discovery. For missing bucket, you cannot edit any of the values. These are entities that are present in the current inventory but have not been discovered in the deployment. You have the following options:

  1. Remove the missing entries from inventory (You will be asked for confirmation whether the entities have ports or are attached to ports.)
  2. Keep the missing entries in inventory. (This means that you opt that the inventory continues to function as before.)

A missing device can be deleted from the inventory via import/discovery only if:

- None of its connected hosts have ports attached.
- It is not the last spine route reflector.

When you delete a missing device, Cisco VTS does the following before deleting the device:

- Detaches all the templates attached to the device and removes the configurations from the device.
- Removes the device from admin domain.
- Uninstalls host-agent or VTF from all the hosts solely connected to the device.

A missing host can only be deleted, if it does not have any ports attached. Before deleting the missing host, the host-agent or VTF is removed from the host/compute.

Important Notes

This section lists a few important notes related to the discovery framework.

- You must not add UCS 6200 Fabric Interconnects to the inventory even if these Fabric Interconnects are discovered during auto discovery.
- While adding new vCenter hosts into Cisco VTS, which has an existing inventory, you must:

  1. Export the current inventory.
  2. Update the exported inventory CSV file with the new vCenter Hosts.
  3. Reimport the CSV file into Cisco VTS, and update the inventory.

- If, in the CSV file you update existing devices authgroup and import again, in the GUI these will be shown under Mismatch devices. Clicking Update Inventory will update the authgroup of existing devices to the authgroup value you specified in the CSV. This change occurs even if you have a workload attached to the device. The same behavior occurs for BGP-ASN, and Loopback Interface Number also.
- After auto discovery is complete, for New devices, you must add the devices first, then add the fabric connection, and then the hosts.
- If you had changed the name of a TOR, which already exists in inventory, and then do a rediscovery, the TOR whose name is changes will be included in the New Device list, and a mismatch will be shown for Fabric Connection (Target Device Name). If you try to add the discovered fabric connection value to the inventory, it will throw an error. You must first add the new TOR to the inventory, and then add the newly discovered fabric connection.
- If Cisco Nexus 3000 device is used as a Leaf, then in Cisco VTS, the Device Platform needs to be set as Cisco Nexus 9000.
- For the New bucket, first add devices, fabric connections, and then hosts. For Missing bucket, first remove the hosts, then fabric connections, and then devices.
- The discovery process discovers only one connection for Cisco UCS B-Series hosts with multiple connections to the same interface. After discovery, you must manually add the details of the connections that are not discovered, via the Host Inventory page.
- When two ToRs are configured in vPC and no dual-homed host (connected to those ToRs) is in the VTS inventory, VTS does not correctly identify the vPC. You must add the dual-homed host connected to the ToRs in vPC to the VTS inventory, before provisioning a port on a host connected to the ToRs in vPC.
- Different ESI groups/domains must have different ES-id or system MAC. In other words, duplicate ES-id and system MAC are not allowed among ESI groups. This needs to be guaranteed by providing correct Day Zero configurations for ESI on Cisco Nexus 9000 switches.
- The Cisco VTS discovery log file is under /var/vts/log. Check for any errors/exceptions in this log file.

Viewing the Network Topology

Topology window provides a view of the data center fabric controlled by Cisco VTS. It displays the leafs, spines, border leafs, DCI, hosts, as well as the software VTEPs. You can get a tenant-based topology view using this feature.
To view the network topology:

Step 1 Go to **Inventory** > **Topology**. The Inventory / Topology window appears.

Step 2 Select the VMM from the VMMs drop-down.

Step 3 Select the tenant for which you need to view the topology, from the **Select Tenant** drop-down list.
The topology is displayed in the Topology window. You can use the following buttons to control the display:

- Select node mode
- Move mode
- Zoom in / Zoom out / Zoom Selection

- Fit Stage
- Full Screen mode

Hover the mouse cursor over the Topology Setting icon to view Topology Setting popup, where you can change the display icon appearance, and display color.

Note In case of FEX or vPC, if no host is connected, Cisco VTS will not show the vPC or FEX in the Topology. Also, you might encounter errors.

The legend provided at the left bottom of the screen help you identify the different types of links (Ethernet/vPC/Multi-Homing/ESI).

Hover the mouse cursor over the link to view the Info popup, which gives the information about the link.

Viewing Network Inventory

The network inventory table displays details about the devices which have been added to the inventory.

To view the network topology:

Go to **Inventory** > **Network Inventory**. The Inventory / Network Inventory window appears with the Network Inventory table displayed.
The following details are displayed:

- Device Name

  Note Click the info icon on the device name to view the detailed information about the device.

- Admin State
- IP Address
- Device Platform
- Device Role
- Group Tag
- Templates Attached
- Sync
- Last Sync Operation

For devices that have no Loopback Interface Numbers/Loopback IP/BGP-ASN Number, you can find a warning icon adjacent to device name. You must update these values if you need these devices to be a part of the admin domain.

Note If you are using VTSR, then the BGP ASN value should be between 0 and 655535.

You can add network devices via the Network Inventory table. To do this, click the **Add (+)** icon, and provide the details. You can use this option to add devices to the inventory.

To edit network device, select the device you want to edit and click the **Edit** icon.

Note For VTSR, Loopback Interface Number Underlay and Loopback Interface Number Overlay fields cannot be edited.

To delete network devices from the Network Inventory table, select the device you want to delete and click the **Delete (X)** icon.

If there is problem in deleting device, you need to make sure that fabric link is cleaned up manually. For example, when Device 1 is connected to Device 2, Inventory has two devices and two fabric links (this can be seen in Fabric Connection tab in Network Inventory)—one from Device 1 to Device 2, and the other from Device 2 to Device 1. While deleting Device 1 from network inventory, cleanup is done for Fabric link Device 1 to Device 2 and for the device from the inventory. The link Device 2 to Device 1 has to be cleaned up manually before you delete.

It is important that you remove the resource pool before deleting a device.

You need to discover the devices and add them to the inventory before you bring up the VTSR. If you do these tasks simultaneously, you might encounter errors.

To recalculate the inventory topology for a particular device, click the redeploy button. See Redeploying Device Inventory for more details.

Adding Fabric Connection

To add fabric connection:

Step 1 Go to **Inventory** > **Network Inventory**. The Inventory / Network Inventory window appears with the Network Inventory table displayed.

Step 2 Click Fabric Connection tab, then click **Add (+)** icon.
The Add Fabric Connection popup window appears.

Step 3 Enter the necessary details and click **Save**.

Synchronizing Configuration

You can check if the device configuration is in sync with Cisco VTS database, using the Check Sync option. Once Check Sync is complete, the sync status of the device along with the differences with the device is displayed. Options to Sync From, Sync To, and Reconcile Service are available. See Important Notes, on page 17section for important information related to Reconcile Service feature.

Note This operation can be done only on a device that has the Admin State as **Unlocked**. If Admin state is **Locked**, you must change the Admin State to **Unlocked**, and then do the check-sync operation. Also, the out-of-sync-commit behavior in System Settings must be set to **Reject** for this feature to be enabled.

Step 1 Go to **Inventory** > **Network Inventory**. The Inventory / Network Inventory page displays the Network Inventory table.

Step 2 Click the **Check Sync** link under the Sync column, for the device.
A popup window is displayed with Check Sync Results. The green + indicates additional configuration on the device and the red - indicates the additional configuration in the VTS database.

Step 3 To synchronize the configuration, you can use the following options:

- Sync From—Synchronize the configuration by pulling configuration from devices into VTS database. The configs marked as + will be added to the VTS database and the configs marked as - will be removed from the VTS database.
- Sync To—Synchronize the configuration by pushing configuration from VTS database to devices. The configs marked as + will be removed from the device and configs marked as - will be added to the device.
- Reconcile Service— Reconciles Out of Band (OOB) configuration from devices to VTS database. Reconcile service enables you to ensure that any out-of-band configuration on the device is absorbed into the VTS database and any subsequent VTS service or L2/L3 template update specific to that configuration will not overwrite the out-of-band configuration on the device.

  Note If switch name (switch hostname) is changed in the switch CLI, the sync to option will not work. The switch name has to be the same as the value in the VTS inventory.

You can choose to initiate these actions on multiple devices. The requests are placed in a queue and each will be initiated in the order initiated.
If the action succeeds, a success green check icon is displayed in the selected device row. If it fails, a red failure icon is displayed in the selected device row. The tooltip for the critical icon displays which action failed and the reason for failure.

Important Notes

This section lists a few important notes related to the out-of-band reconcile feature.

- We recommend that you use Out-of-Band reconciliation feature to reconcile configuration that is pushed to the device via ports created from VTS GUI only. Using this feature to reconcile configuration in a VMM integrated VTS setup, where ports are created from the VMM, might cause errors.
- You must ensure the day zero configuration on the device does not include configuration that will be pushed using Cisco VTS services or device templates. That is, device day zero configuration should not include configuration which would conflict with the configuration that VTS would be pushing into the device either via service configuration or device template configuration.
- In certain cases, if a port detach operation fails, you may need to remove any related out-of -band configurations from device, do an out-of-band reconcile operation from the Cisco VTS GUI, and then try the port detach operation again.

Viewing Host Inventory

You can view the details of the hosts connected to the switches.

To view host inventory details:

Step 1 Go to **Inventory** > **Host Inventory**. The Inventory / Host Inventory page appears. The Host Inventory page has two tabs—**Virtual Servers** and **Baremetals**. By default, the page displays Virtual Server details.

Step 2 To view host details on Virtual Servers, select the VMM from the Select VMM drop-down, and select the device from the Select Device drop-down list. The following details are displayed:

- Host Name
- IP Address
- Host Type
- Associated VMM
- Virtual Switch
- Interfaces
- Installation Status—Shows the installation status.
- VTF Mode—Displayed on the top right of the table shows the VTF mode you have chosen in the Administration > System Settings window.

Step 3 To view host details on Baremetals, select the **Baremetals** tab, then select the device from the Select Device drop-down.

Adding a new Host on Virtual Servers

To add a new host:

Step 1 Click the **Add (+)** icon. The Add New Host popup window appears. It has two tabs—Host Details and Host Interfaces. the Host Details tab is selected by default.

Step 2 Enter the following host details:

- Host Name—This is mandatory. Only letters numbers, underscore and dashes are allowed. Requires at lease one letter or number. The hostname entered here needs to be in FQDN format, that is. <hostname>.<domain>.
- Host IP Address—This is mandatory.
- User Name
- Passphrase— User Name and passphrase are mandatory if you choose Non-OSPD VMM name in the VMM Name drop-down of the 'Host Configuration' section in the current popup window.

- Host Configuration
  - VMM Name—The VMM to which you want to associate the host to. Depending on VMM chosen in the VMM Name section either the VTF Details information is pre-populated or you have to enter the details.
  - Virtual Switch—The following options exist:
    - not defined
    - ovs—If you want to install the VTS host agent on the compute, check the Install VTS agent on save check box.
    - vtf-l2—VTF is used as an L2 switch.
    - vtf-vtep

  Note The options displayed here depends on what you have specified in the VTF Mode field in Administration > System Settings and the VMM type.

  The same host cannot support OVS and L2 at the same time. However, in the same host OVS and L2 can reside together with SR-IOV. Some ports can be SR-IOV ports, and others can have L2 switch or OVS.

Step 3 If you choose vtf-vtep or vtf-l2, a new tab VTF Details is displayed. Go to VTF Details tab and enter the required information for the VTF-L2/VTEP.

- VTF Name—Only letters, numbers, underscores and dashes are allowed. Requires at least one letter or number.
- VTF IP—Enter Compute host underlay IPv4 address.
- Subnet Mask—Enter compute host underlay subnet mask.
- Max Huge Page Memory—Max huge page memory % that is being allocated on the host. This value is greater than 0 and less than or equal to 100. Default value is 40.
- Gateway—Enter the Compute host underlay gateway.
- PCI Driver—vfio-pci and uio-pco-generic are supported. Choose an option from the drop-down.
- Underlay Interfaces—Interface connected from compute host to the physical device (N9K/N7K/N5K). It has 2 options, Physical or Bond. Select Physical if you need to add only one interface that are connected from the compute host.
  Select Bond option if you need to add multiple interfaces that are connected from the compute host. i.e multiple entries in the Interfaces' tab.
- Bond Mode—Choose required Bond mode from the drop-down.
- Bond Interfaces—Add multiple Interfaces.
- Routes to Reach Via Gateway—Routes to reach other underlay networks from this VTF host.

Advanced Configurations Section:

- Multi-Threading—Set Enable Workers to true for Multithreading. By default it is set to true.
- Jumbo Frames Support—By default, it is true.
- Jumbo MTU Size—Enter Value Between Range of 1500 - 9000.

If you want to install VTF on the compute select the checkbox 'Install VTF on Save'. Depending on the type of VMM Name chosen in the Host Details tab, either you can 'Save' or 'Save and Validate'.
The VMM can be OSPD/Non-OSPD VMM based on the VMM registration. See Registering the Virtual Machine Manager using GUI. For OSPD, the Host will allow for validation of installed plugins (either OVS or VTF). For Non-OSPD, the Host will allow installation of plugins on Host Inventory UI.

Step 4 Enter the Host Interface details. At least one interface is mandatory.

- Host Interfaces—This is mandatory.
- SR-IOV Enabled—Choose Yes or No from the drop-down to specify whether the interface is SR-IOV enabled.
- Phys Net— Physnet name associated with the interface. If *SR-IOV Enabled* is Yes, it is the Physnet to be used for SR-IOV. If it is No, the other possibilities are that the port is associated to L2 switch or OVS. In case of OVS, you need to give Physnet intended to be used for OVS.
- Attached to Device—Choose the device from the drop-down.
- Device Port—This is mandatory. Choose the device port from the drop-down.
- Group

  To add more interfaces, use the **Add (+)** icon.

Step 5 Click Save. Host details and at least one interface have to be added for the Save button to be enabled.

Adding a new Host on Baremetal

To add a host:

Step 1 Click the **Add (+)** icon. The Add New Host popup window appears. It has two tabs—Host Details and Host Interfaces. the Host Details tab is selected by default.

Step 2 Enter the Host Name. Only letters numbers, underscore and dashes are allowed. Requires at lease one letter or number.

Step 3 Enter the Host IP Address. IPv4/IPv6 address of the host. This is mandatory.

Step 4 Enter the Host Interface details. At least one interface is mandatory.

- Host Interfaces—This is mandatory.
- SR-IOV Enabled—Choose Yes or No from the drop-down to specify whether the interface is SR-IOV enabled.
- Phys Net
- Attached to Device—Choose the device from the drop-down.
- Device Port—This is mandatory. Choose the device port from the drop-down.
- Group

  To add more interfaces, use the **Add (+)** icon.

To edit a host from the table, select the Host Name check box corresponding to the device and click the **Edit** icon. You can also click the port icon in the Interfaces column to open the Edit Host popup. You can also use the Bulk Edit option to make changes to more than one host.

You cannot edit hosts on which there are workloads associated.
To delete a host from the table, select the Host Name check box corresponding to the device and click the **Delete (X)** icon.

Note To convert a virtual server host to Baremetal, delete the host and add it as Baremetal.

Viewing the VTSR to VTF Mapping

Step 1 Go to **Inventory** > **Virtual Forwarding Groups**. The Inventory / Virtual Forwarding Groups window appears.
The window displays the number of VTFs that are attached to the VTSRs . The table on the right hand side shows the VTFs.

Step 2 To disassociate the VTF from Virtual Forwarding Group (VFG), select the VTF on the right pane, and click the detach icon.

Note When VTF is in L2 mode, this window is read only. You cannot detach the VTF in this mode. See also, the *Deleting VTF in a vCenter Environment* and *Deleting VTF in an OpenStack Environment* sections in the *Cisco VTS Installation Guide* for more details.

SR-IOV Support

Multiple NIC Cards are supported. The following combinations are supported:

- SR-IOV + OVS
- SR-IOV + VTF as L2 Switch
- SR-IOV + SR-IOV

SR-IOV is supported for OpenStack only. VXLAN, VLAN, and Flat network types are supported.

The following Provider network types are supported:

- VLAN and Flat provider network
- Static VLAN (segmentation ID) is honored for VLAN networks.

The default tenant network type is VXLAN.

Assigining VLAN Ranges

Based on the network_vlan_ranges in OpenStack, at /etc/neutron/plugins/ml2/ml2_conf.ini (Controller node), you need to configure:

- Device level and device interface level restricted vlan pool for Cisco Nexus 7000 devices in Cisco VTS.
- Device level restricted vlan pool for Cisco Nexus 7000 devices in Cisco VTS.

See the Managing Resources chapter for details about assigning VLAN ranges.

SR-IOV related fields SR-IOV Enabled and Phys Net can be edited on Host Interfaces tab in Host Inventory, when you add/modify the host.

Trunk Port Support

Cisco VTS supports OpenStack Trunk Port feature for SR-IOV. See OpenStack documentation for information about creating Trunks and Subports.

Migrating from vPC to ESI

This section provides details about the generic procedure to migrate from Virtual Port Channel (vPC) to Ethernet Segment Identifier (ESI).

Note Before you begin, ensure that the following TCAM regions are carved on Cisco Nexus 9000 series switch:

```
hardware access-list tcam region vpc-convergence 256
hardware access-list tcam region arp-ether 256
```

To migrate from vPC to ESI:

Step 1 In case of VTSR HA, bring down the VTSR.

Step 2 Upgrade VTS to a version which supports ESI.

Step 3 If the TCAM regions, as mentioned above, are not already carved on Cisco Nexus 9000 series switch, add the lines and save as running config.

```
hardware access-list tcam region vpc-convergence 256
hardware access-list tcam region arp-ether 256
```

Note Do not reboot device (as the TOR will be rebooted in the next step).

Step 4 Upgrade TORs to a new Cisco Nexus 9000 image, which has ESI feature. This will automatically cause device to reboot.

```
copy run start
install all nxos bootflash:/nxos.7.0.3.I4.1t.bin
```

Step 5 Upgrade Cisco ASR 9000 series DCIs to an ESI supporting image.

Step 6 Once the setup is up then remove feature vPC and configure ESI on the required TORs that you are planning to convert to ESI.

Remove vPC	`no feature vpc`
Remove other vPC related configuration under port channel and Ethernet Interfaces	

Remove secondary interface from loopback	`interface loopback0` `  no ip address 44.44.44.44/32 secondary`
Enable ESI	`evpn esi multihoming`

Create nve	`interface nve1`
	`  no shutdown`
	`  source-interface loopback0`
	`  host-reachability protocol bgp`

Enable core links	`interface Ethernet1/35`
	`  Description " Connected with Spine"`
	`  no switchport`
	`  evpn multihoming core-tracking  <<< Add here`
	`  ip address 16.1.1.2/24`
	`  ip router ospf 100 area 0.0.0.0`
	`  ip pim sparse-mode`
	`  no shutdown`

Add Ethernet-segment and system-mac address in the port-channel	`interface port-channel220`
	`  switchport mode trunk`
	`  switchport trunk allowed vlan none`
	`  ethernet-segment 220`
	`  system-mac eeee.1111.2222`

Apply the channel group to the TORs interface which are connected to compute.	`interface Ethernet1/5`
	`  switchport trunk allowed vlan none`
	`  channel-group 220 mode active`

Verify whether the ESI is up.	`tor1# show nve ethernet-segment`
	
	`ESI Database`
	`----------------------------------------`
	`ESI: 03aa.bbcc.ddee.ee00.002d,`
	`  Parent interface: port-channel30,`
	`  ES State: Up`
	`  Port-channel state: U`
	`  NVE Interface: nve1`
	`  NVE State: Up`
	`  Host Learning Mode: control-plane`
	`  Active Vlans: 1001`
	`  DF Vlans: 0-4095`
	`  Active VNIs: 30001`
	`  Number of ES  members: 1`
	`  My ordinal: 0`
	`  DF timer start time: 00:00:00`
	`  Config State: config-applied`
	`  DF List: 1.1.1.1`
	`  ES route added to L2RIB: True`
	`  EAD routes added to L2RIB: True`
	`----------------------------------------`

Step 7 On Cisco VTS, perform a sync-from operation for the TORs that have ESI enabled.

Step 8 Redeploy inventory from Cisco VTS only for devices that have new ESI configuration. This is to make sure that Cisco VTS recognizes ESI configuration on Cisco Nexus 9000 series devices. See Redeploying Device Inventory, on page 24 for details.

Step 9 Remove the peer links between previous vPC peer TORs (**Inventory > Network Inventory > Fabric Connection**).

Step 10 Add the ESI device group to appropriate functional groups in Admin Domain, and also disable ARP suppression at **(Overlay > Network)**.

Step 11 Upgrade VTSR to the latest image.

Step 12 Run the Migration script from the path /opt/cisco/package/vtc/bin/vpc-migration. For an HA setup, run this on the Active VM.
For example:

```
root@vtc1:/opt/cisco/package/vtc/bin/vpc-migration# ./VpcEsiMigration.py -u admin -p Cisco123! -s
-target esi -dev stb2-tor1 stb2-tor2
```

Where:

- – u is the VTS GUI username.
- – p is the VTS GUI password. Use a single quote (‘) before and after a password that contains special characters. Especially when the password contains an & character in it.
- target esi for the vPC to ESI Migration
- stb2-tor1 and stb2-tor2 are the hostname of a pair of TOR devices running ESI Day 0 configuration. Modify the name to fit your own hostnames. Also, run the script for one ESI TOR-pair at a time if there is more than one in your environment.

Redeploying Device Inventory

You can use the Redeploy feature to recalculate the inventory topology for a particular device. This is important in the context of vPC and ESI.

You need to Redeploy the inventory when device Day Zero configuration changes for:

- vPC or ESI. For example, vpc id for a port-channel is changed
- port-channel or ether-channel

Redeploy triggers the inventory for a device again. Since inventory reads the data from the device model in the database it is important to perform sync-from before doing a Redeploy.

Note Redeploy function is different from the sync-from function. Sync-from gets the configurations from the device and updates it in the device model in the database. However, it does not recalculate the topology. That is, the topology would still show old information/configuration. Redeploy recalculates the inventory topology. After you perform a Redeploy, the topology will be updated with the modified configuration.

To redeploy device inventory:

Step 1 Go to **Inventory > Network Inventory**, perform a sync-from for the device for which the configuration has changed. See Synchronizing Configuration, on page 16 for more details.

Step 2 Select the device, click Redeploy.

Note Redeploy just recalculates the inventory. Existing ports/VMs belonging to old device configurations, would not be updated or redeployed. You might need to delete and recreate the existing ports. We recommend that you use redeploy only if there are no existing ports/router/router interfaces.

Note If you delete devices from the inventory and also deleted VTSR with it, when you redeploy or reload the inventory, VTSR will not show up until it is reloaded or restarted. Power on the VTSR and wait for the registration with VTC to complete.

Enabling Static Multi Homing

Static multi homing can be enabled on Cisco Nexus 7000 series and Cisco Nexus 9000 series devices. You can enable static multi homing by connecting one compute to two ToRs.

When you perform a port attach on VMs attached this compute, the configuration is pushed on both the ToRs. Currently, static multi home feature is supported for two ToRs, that is, one compute can connect only to two ToRs. Static multi homing also builds in high availability where one of the interfaces is an active and the other is a standby.

Enabling Static Multi Homing on Cisco Nexus 7000

To enable static multi homing on Cisco Nexus 7000 devices:

Step 1 Group the interfaces using the **Resources > Devices > Interface Groups** UI.

Step 2 In Host Inventory, add the same tag for both the interfaces that are connected to the host for which you are enabling static multi homing.

Enabling Static Multi Homing on Cisco Nexus 9000

To enable static multi homing for Cisco Nexus 9000 devices:

Step 1 Group the devices using **Resources > Devices > Groups** UI.

Step 2 In Host Inventory, add the same tag for both the devices that are connected to the host for which you are enabling static multi homing.
If you have the devices already added to admin domain, you will need to update the admin domain to use the device group instead of individual devices.