



Provisioning Overlay Networks

This chapter has the following sections:



Note

Cisco VTS does not support out of band settings that have been done on devices. Any out of band configuration done on the devices will be lost while you provision overlays using Cisco VTS.

- [Provisioning Overlay Networks Using Cisco Virtual Topology System, on page 1](#)
- [Creating Overlays, on page 2](#)
- [Creating Network using VMware, on page 4](#)
- [Creating Subnetwork using VMware, on page 4](#)
- [Creating Routers using VMware, on page 4](#)
- [Attaching Network to Router, on page 5](#)
- [Attaching a Virtual Machine to Network, on page 5](#)
- [Creating a Network using Cisco VTS GUI, on page 5](#)
- [Creating Router using Cisco VTS GUI, on page 6](#)
- [Port Extensions Support , on page 8](#)
- [Assigning BVI Interface IP Address, on page 14](#)
- [Extending Layer 2 Network Across Data Centers, on page 15](#)
- [Enabling Global Route Leaking Service, on page 15](#)
- [Enabling L3VPN to EVPN Route Stitching, on page 17](#)
- [Adding Static Routes, on page 17](#)
- [OpenStack Allowed Address Pairs Support, on page 20](#)

Provisioning Overlay Networks Using Cisco Virtual Topology System

Virtual Topology System enables overlay connectivity orchestrated through an SDN-based control plane. This ensures instant availability of computing and application workloads in the virtualized data center, and removes network provisioning challenges.

Cisco VTS uses VXLAN to overcome scale limits in the data center and to segment the network better. VXLAN is designed to provide the same Ethernet Layer 2 network services as VLAN does, but with greater extensibility and flexibility. The dependence on a Layer 3 underlay network allows VXLAN to take complete

advantage of Layer 3 routing, equal-cost multipath (ECMP) routing, and link aggregation protocols. Virtual Topology System supports hardware and software VTEPs to segment the data center network.

Virtual Topology System supports both VXLAN overlays using the BGP EVPN control plane and VXLAN overlays using IP Multicast-based techniques.

Implementing VXLANs using MP-BGP EVPN based control plane to manage the VXLAN overlay provides a distributed network database, which enables federation and scaling. The BGP EVPN solution is the preferred option, and it can be flexibly implemented using the infrastructure policy constructs within the Virtual Topology System environment.

Virtual Topology System implements the highly scalable MP-BGP with the standards-based EVPN address family as the overlay control plane to:

- Distribute attached host MAC and IP addresses and avoid the need for unknown unicast, and multicast traffic
- Support multi-destination traffic by either using the multicast capabilities of the underlay or using unicast ingress replication over a unicast network core (without multicast) for forwarding Layer 2 multicast and broadcast packets
- Terminate Address Resolution Protocol (ARP) requests early

Control-plane separation is also maintained among the interconnected VXLAN networks. Capabilities such as route filtering and route reflection can be used to provide flexibility and scalability in deployment.

High-level Workflow for Establishing a VXLAN Overlay Network with Hardware and Software VTEPs using BGP EVPN

The following steps provide a high-level workflow for establishing a simple VXLAN overlay network with hardware and software VTEPs using a BGP EVPN control plane:

- Prepare the physical environment to be managed by Cisco VTS to build virtual overlays. See the *Prerequisites* section in the *Cisco VTS Installation Guide* for details.
- Discover the network topology in the data center. See the *Managing Inventory* chapter of the *Cisco VTS User Guide* for details.
- Define Admin Domains. See *Creating and Managing Admin Domains* chapter of the *Cisco VTS User Guide* for details.

After you commit the changes to the network group, Virtual Topology System automatically pushes all the relevant configuration information to the respective leafs, VTSR, and DCI gateways. At this point, the Admin Domain is ready to build overlay networks based on the intent defined by the service policy or through a Virtual Machine Manager (VMM) or orchestration environment.

Cisco VTS supports dual stack IPv4 and IPv6 addressing for overlay provisioning.

For a detailed, illustrated example, see *Cisco Virtual Topology System: Data Center Automation for Next-Generation Cloud Architectures White Paper*.

Creating Overlays

As part of overlay provisioning, you may need to:

- Create Tenant

- Create Network
- Create Subnet
- Create Router
- Create VM

This can be done using the VMM or Cisco VTS GUI.

**Note**

- If you create a Network in OpenStack, and then attach a Baremetal port from Cisco VTS, you must not delete the Network from OpenStack before all baremetal ports attached to this network are deleted from Cisco VTS.
- If you attach VTS subnets (Baremetal) to a router from Cisco VTS GUI, then attach the OpenStack subnets to the same router from the Cisco VTS GUI, all subsequent operations on these subnets need to be done from Cisco VTS.

Using OpenStack

**Note**

When you use a VMM such as OpenStack or VMware, the plugin will provide integration between the VMM and Cisco VTS. Once Tenant/ Network/ Subnets are created on the VMM, required overlay network(s) will automatically be created by Cisco VTS.

For information about performing these tasks via OpenStack Horizon dashboard, see OpenStack documentation.

Using VMware

For information about performing these tasks using VMWare, see the following sections:

- [Attaching Network to Router, on page 5](#)
- [Creating Network using VMware, on page 4](#)
- [Creating Subnetwork using VMware, on page 4](#)
- [Creating Routers using VMware, on page 4](#)
- [Attaching a Virtual Machine to Network, on page 5](#)

**Note**

The VTS tab appears under the Configure tab in vCenter 6.5. In vCenter 6.0, it appears under Manage tab.

Using Cisco VTS GUI

For information about creating Network and Router using Cisco VTS GUI, see the following sections:

- [Creating a Network using Cisco VTS GUI, on page 5](#)
- [Creating Router using Cisco VTS GUI, on page 6](#)

Creating Network using VMware

To create a network:

-
- Step 1** Select one of the vDS switches you created, then click **Manage** tab.
 - Step 2** Select the Cisco VTS Network tab and click **Add (+)** to add the network.
 - Step 3** Select create Tenant and enter Network Name field.
 - Step 4** Click **Create** to create the network.
 - Step 5** Click the **Refresh** icon to display the created network.
-

Creating Subnetwork using VMware

Before you create the subnetwork, you need to create the network in which the subnetwork has to be created.

To create subnetworks:

-
- Step 1** Select one of the vDS switches you had created, then click the **Manage** tab.
 - Step 2** Select Cisco VTS Network tab, and click the network name in which the subnetwork has to be created.
 - Step 3** Enter the subnet name, the network range in CIDR format, and the Gateway IP.
 - Step 4** Click **Create Subnet** button to create subnetwork.
 - Step 5** Click **Refresh** button to see the subnetwork.
-

Creating Routers using VMware

-
- Step 1** Select one of the vDS switches you had created, then click the **Manage** tab.
 - Step 2** Select Cisco VTS Router tab, and click **Add (+)** to add the Router.
 - Step 3** Select **Tenant Name** and enter the **Router Name**.
 - Step 4** Click **Create Router** button to create the router.
-

Attaching Network to Router

To attach a network and subnetwork to a router:

-
- Step 1** Select one of the vDS switches you had created, then click **Manage** tab.
- Step 2** Select VTS Router tab and click the **Router Name** where network has to be added.
The Router Details popup appears.
- Step 3** Select Network and subnet and click **Attach Subnet**.
-

Attaching a Virtual Machine to Network

To create VMs:

-
- Step 1** Create network and subnet using vCenter Cisco VTC plugin. This will create portgroup for the network.
- Step 2** Create the VM in vCenter and attach the created portgroup to the VM.
This will attach the VM to the network created via Cisco VTS.
-

Creating a Network using Cisco VTS GUI

To create a network from the Cisco VTS GUI:



Important

- You must verify that ARP Suppression is supported on the switches where the network will have ports attached. Cisco Nexus 9000 devices do not support ARP suppression for Fabric/Host networks when SVI is not created. ARP suppression must not be enabled in cases where ARP is used by applications for keep alive and monitoring.
 - If you create a Network in OpenStack, and then attach a Baremetal port from Cisco VTS, you must not delete the Network from OpenStack before all Baremetal ports attached to this network are deleted from Cisco VTS.
-

-
- Step 1** Go to **Overlay > Network**. The Overlay / Network window appears.
- Step 2** Click Fabric Host Networks or External Networks, based on your need.
- Step 3** Click **Add (+)** icon.

Note For External Networks you need to specify the Name, Tenant, and Zone.

- Step 4** Enter the network name. This is mandatory.
- Step 5** Select the Tenant for which you to create the network.
- Step 6** Select the Zone.
- Step 7** If the network is not external, enter the Static VNI. This can be an integer between 4096 and 65535.
- Step 8** Specify whether the network is a shared network.
- Step 9** Specify whether the network has to be Extended. If you select Yes, VPN Service becomes available for use.
- For the L2 Extended Network, click the **L2VPN** tab and enter the EVI number. This can be an integer between 1 and 65534. Select the Load Balance Per EVI check box to introduce the load balance CLI in the device. See [Extending Layer 2 Network Across Data Centers, on page 15](#).
- Step 10** Specify whether ARP Suppression needs to be enabled.
- Step 11** Click **Save**.
- You can also add a subnet, and add port.

Creating a Subnetwork

To create a subnetwork:

- Step 1** Click **Add (+)** in the Subnet pane of the Add Network page.
- Step 2** Enter the subnet name. Only IPv4/IPv6 addresses, alphabets, space, numbers, and special characters /, - and _ are allowed.
- Step 3** Enter the IP details. You can enter an IPv4 or IPv6 address. You must ensure that the network address and the gateway IP are in sync.

You can create subnets with /31 prefix. /31 subnet masks are used for point to point links. The gateway IP for a subnet with /31 prefix should be within the two allowed IP addresses. For Example:

- For subnet 10.20.30.0/31, the allowed IPs are 10.20.30.0 and 10.20.30.1
- For subnet 10.20.30.5/31, the allowed IPs are 10.20.30.5 and 10.20.30.4

Note When you have a network with /31 subnet, then you cannot make it as an external network and vice versa.

- Step 4** Click **OK**.
- The table displays the Subnet Name, Network Address, Gateway IP, and the IP Version (whether IPv4 or IPv6).

Creating Router using Cisco VTS GUI

To create a router using Cisco VTS GUI:



Important If you attach VTS subnets (Baremetal) to a router from Cisco VTS GUI, and then attach the OpenStack subnets to the same router from the Cisco VTS GUI, all subsequent operations on these subnets need to be done from Cisco VTS.

- Step 1** Go to **Overlay > Router**. The Overlay / Router window appears.
- Step 2** Click the **Add (+)** icon. The Add Router window appears.
- Step 3** Select the tenant from the **Select Tenant** drop-down list.
- Step 4** Select the Zone from the **Select Zone** drop-down list.
- Step 5** Enter the Static VNI. This can be an integer number between 4096 and 65535
- Step 6** Enter the **Router Name**.
- Step 7** Select a template that you might want to associate with the router, using the find icon in the Template field. See [Attaching Templates while Adding Routers](#) for details.
- Step 8** Enter a VRF name. This is optional. If this is left empty, when the **Save** button is clicked, a default VRF name gets automatically generated.
- The custom VRF name accepts up to 24 characters.
 - If there is no input for custom VRF name, a default VRF name gets generated in form of `<tenant-name>-<router-name>`. Both tenant-name and router-name accept up to 15 characters.
 - If Cisco ASR 9000 series router is configured as DCI in the domain, and you have not given a custom VRF name, then you must ensure that the default VRF name does not exceed 27 characters. Otherwise, the configuration will fail.
 - If VTSR is configured, and you have not given a custom VRF name, then you must ensure that the default VRF name does not exceed 24 characters. Otherwise, the configuration will fail.
 - If configuration fails because the default VRF name exceeds the limit, you can choose to use custom VRF name instead.
 - If the configuration fails because the default VRF name exceeds 27 characters, an error message appears on the Network > Port Attach screen, which indicates invalid input for “bridge-domain” configured on Cisco ASR 9000 series router.
 - For VTSR configuration, a similar error is displayed if the default VRF name exceeds 24 characters.
 - VRF name change from VTS GUI is not supported for VTSR. Cisco VTS does not allow changing the name of a router if it connects to a port on a V node. (A V node is compute node where there is a VTF present, and the workload is behind a VTF where the VXLAN Tunnel originates.)
 - If you modify the VRF name after saving the router, the Router Gateway IP address gets removed. You can reconfigure it back after saving the VRF name change.
- Step 9** Select the router gateway from the **Router Gateway** drop-down list. When you select External GW from drop-down, two additional fields for Router Gateway IPv4 and Router Gateway IPv6 get displayed. These are optional.

When you select Router Gateway, the Advertise Default Route toggle switch is displayed. It is enabled by default. When it is enabled, the default routes are pushed on the DC gateway device in VRF-Peering mode and on the DCI device in integrated mode. For example:

```
router static
vrf tl-rout
  address-family ipv4 unicast
    0.0.0.0/0 Null0 254
  exit
exit
exit
router bgp 65539
```

```
vrf t1-rohi-rout
rd 2.2.2.11:10009
address-family ipv4 unicast
  label mode per-vrf
  maximum-paths ebgp 2
  maximum-paths ibgp 2
  network 0.0.0.0/0
  aggregate-address 3.2.3.0/24 summary-only
  redistribute connected
exit
address-family ipv6 unicast
  label mode per-vrf
  redistribute connected
```

When set to No, the default routes are not pushed.

- Step 10** If the router is used to add shared networks from different tenants as interfaces, set the **Provider Router** toggle switch to **Yes**.
- Step 11** Click **Add (+)** icon. The Add Interface popup appears.
- Step 12** Select the subnet from the drop-down list, and click **OK**.
- Step 13** Click **Save** in the Add Router window to save the router and its interface.

Port Extensions Support

Port Extensions is a VTS construct that allows additional services to be configured on the TORs to which the associated overlay ports are connected. Port Extensions Type determines the nature and scope of the configuration that is pushed to the TORs.

Currently BGP service configuration—iBGP/eBGP, is supported.

For BGP Port Extension type, configuration pushed to relevant TOR devices is scoped to within the VRF to which the overlay port belongs. If the network to which the port belongs, is not associated with any VRF then no settings in this object take effect on the TOR. When port is associated with a VRF then Port Extension can be used to bring up BGP peering session on the TOR towards the VMs. This allows CE <--> PE L3 VPN style peering between the VMs that are playing the CE role, and respective TOR device that is playing the Port Extension role. BGP peering between VMs and TOR allows dynamic exchange of overlay routes between them. Upon VM migration from one TOR to the other TOR, any associated BGP configuration driven through Port Extensions also get automatically transferred to the new TOR. BGP peering sessions get automatically torn down from the old TOR and established on the new TOR.

Creating a Port Extension

You can create Port Extensions and attach them to Baremetal Ports and Virtual Machine Ports. To create Port Extension in Cisco VTS:

- Step 1** Go to **Overlay > Port Extensions**. The Overlay / Port Extensions window appears.
- Step 2** Select a tenant from the Tenant drop-down. The following details are displayed.
- Port Extension Name—The Port Extensions that have been created for the tenant.
 - Description

- Type
- Zone
- Attached Ports

Step 3 Click **Add (+)** icon to create a new Port Extension.

1. In the Add Port Extension pane, enter the following details:
 - Port Extension Name—This is a mandatory field.
 - Description
 - Type—Type it BGP.
 - Tenant—The tenant under which the Port Extension is being created.
 - Zone—Select any zone from the drop-down.
2. BGP Profile Information—BGP Profile Information (VRF Config) contains details that will influence the type and nature of the BGP peering sessions initiated by the TOR (connected to overlay ports) towards VM instances. Enter the following details:
 - Route Reflector Mode—Applicable in iBGP scenarios. Select **None** for eBGP. The following options are available:
 - None
 - Client
3. In the Neighbor List pane, you can either add or edit the neighbor list. Click **Add (+)**.

Note You can create many neighbors. You need to add at least one neighbor.

The Add Neighbor page appears with the following details:

1. BGP Neighbor Information:
 - Neighbor Id—Neighbor IP to which the BGP session needs to be established. This is a mandatory field, there is a format.
 - Description
 - Neighbor ASN—Applies only in eBGP case.
 - Local ASN—Applies only in eBGP case.
 - Note** If you enter a value for Neighbor ASN, then Local ASN value can not be the same as that of Neighbor ASN.
 - Local Source Loopback Number
 - Passphrase—Password to establish the BGP session with the neighbor.
 - Suppress 4-byte ASN—Suppress 4-byte AS Capability.
 - BFD—Bidirectional Fast Detection for the neighbor.

- eBGP-Multihop—Number of hops the eBGP peer is away. For directly connected peers, leave this field empty.

Note All eBGP fields need to be removed before moving the neighbor session from eBGP type to iBGP.

Due to platform limitation, to switch from eBGP to iBGP with an attach Port Extension, you need to follow the platform flow with the following steps:

- Edit the Port Extension by removing the values for all eBGP specific fields (except remote-as). Examples of eBGP fields are eBGP-Multihop, disable-peer-as-check, remove-private-as, and so on.
- Save the Port Extension.
- Remove **Local ASN** and then change Neighbor ASN to make it iBGP.
You can also convert from eBGP to iBGP by detaching the eBGP Port Extension and then attaching the iBGP Port Extension.
- Save the Port Extension.

- Remove-private-AS—Removes the private ASNs.
- Keep Alive—Time interval for transmission on keep alive messages between neighbors. Set this as 1/3 of Hold Timeout.
- Hold Timeout—Time interval in seconds until which the BGP session will be kept active in the absence of keep alive or other messages from the peer. Set this as 3x of Keep Alive.

2. Address Family List:

Note You can add more than one Address Family List. Make sure that at least one Address Family List exist all the time. You can make only four different entries to the Address Family List, that is, for IPv4 unicast, IPv4 multicast, IPv6 unicast, and IPv6 multicast.

- Address Family—Choose the address family type from the drop-down.
- AS Override—Override matching AS-number while sending update.
- Send Community—Choose from the drop-down. Selecting Both sends Extended and Standard community attributes.
- Soft-reconfig—Enable soft-reconfig if neighbor does not support dynamic soft reset.
- Default-originate—Advertise default route to this neighbor.
- Nexthop-type—Nexthop type for eBGP peering. Default value is next-hop-third-party.
- Disable Peer-AS—Disable checking of peer AS number while advertising.
- AllowAS-in—Have the radio button either enable or disable. If you click on **enable**, another field **occurrences** will show up.

3. Route Filter List—This is not a mandatory field. Specify the following:

- Type—Route Map or Prefix List.
- Name

- Direction—filter-in or filter-out.

You may add more Route Filter Lists using the Add (+) button.

Step 4 Click **Add** to add the neighbor details.

Step 5 Click **Save**. The Overlay / Port Extensions page appears where you can see that the Port Extension is created successfully.

Editing a Port Extension

You can modify a Port Extension that you have created.

Step 1 Go to **Overlay > Port Extension**. The Overlay / Port Extensions window appears.

Step 2 Select the check box corresponding to the Port Extension you need to edit, and click the **Edit** icon.

Step 3 Make the desired changes in the attributes.

Note You can edit the Port Extension name only for the ones that are not attached.
You can make only four entries to the Address Family List.

Step 4 Click **Add (+)** icon to add any number of Route Filter Lists based on your requirement. Click the **Remove (-)** icon to remove any Route Filter List.

Step 5 Click **Add** to add the neighbor details.

Step 6 Click **Update**.

Step 7 Click **Save**. The Overlay / Port Extensions page appears where you can see that the Port Extension is updated successfully.

Deleting a Port Extension

You can delete a Port Extension that you have created.

Step 1 Go to **Overlay > Port Extension**. The Overlay / Port Extensions window appears.

Step 2 Select the check box corresponding to the Port Extension you need to delete, and click the **Delete (x)** icon.

Note You will not be able to delete a Port Extension that is attached to any ports. You need to detach the Port Extensions from those ports and then delete the Port Extension.

Step 3 Click **Yes** to delete the selected Port Extension that does not have any ports attached.

Attaching Port Extension to Baremetal Ports

You can attach a single port extension to one or multiple Baremetal ports.

Before you begin

You need to have a Port Extension created already.

Step 1 Go to **Overlay > Baremetal Ports**. The Overlay / Baremetal Ports window appears.

Step 2 Select any tenant from the Tenant drop-down.

The table shows the following details:

- Baremetal Port ID
- Baremetal
- Device
- Device Port
- Network Name
- VLAN Number
- Attached Port Extension
- Attached Security Group

Step 3 Select the **Attach** icon. Attach Port Extension window appears.

Step 4 Select the **Zone**.

Step 5 Select a VRF from the **VRF** drop-down.

The list of available ports gets displayed for the selected VRF.

When you select the VRF, there is an option called **Ports with no VRF** for you to attach a Port Extension even if you do not have a router created. In case you create a router later, the corresponding VRF configuration will be pushed to the device.

Step 6 Select the ports that you want to attach from the displayed list.

Step 7 Select the Port Extension to attach to from the **Port Extension** drop-down.

Step 8 Click **Attach**. Baremetal ports page is displayed with the attached Port Extension displayed as a link.

Detaching Port Extension from Baremetal Ports

To detach a Port Extension from Baremetal Ports:

Step 1 Go to **Overlay > Baremetal Ports**. The Overlay / Baremetal Ports window appears.

Step 2 Select any tenant from the Tenant drop-down.

Step 3 Select the **Detach** icon. Detach Port Extension window appears.

Step 4 Select the **Zone**.

Step 5 Select a VRF from the **VRF** drop-down.

Step 6 Select the ports that you want to detach from the displayed list.

Step 7 Click **Detach**.

- Step 8** Click **Yes** to confirm.
If the Port Extension detach fails, you can see the tool tips for failure message.
-

Attaching Port Extension to Virtual Machine Ports

You can attach a single port extension to one or multiple Virtual Machine (VM) ports.

Before you begin

You need to have a Port Extension created already.

- Step 1** Go to **Overlay > Virtual Machine Ports**. The Overlay / Virtual Machine Ports window appears.

- Step 2** Select any tenant from the Tenant drop-down.

The table shows the following details:

- VM Port ID
- Binding Host
- Type
- Device
- Network Name
- SRIOV Enabled
- VLAN
- Attached Port Extension
- Attached Security Group

- Step 3** Select the **Attach** icon. Attach Port Extension window appears.

- Step 4** Select the **Zone**.

- Step 5** Select a VRF from the **VRF** drop-down.

The list of available ports gets displayed for the selected VRF.

When you select the VRF, there is an option called **Ports with no VRF** for you to attach a Port Extension even if you do not have a router created. In case you create a router later, the corresponding VRF configuration will be pushed to the device.

- Step 6** Select the ports that you want to attach from the displayed list.

- Step 7** Select the Port Extension to attach to from the **Port Extension** drop-down.

- Step 8** Click **Attach**. Virtual Machine ports page is displayed with the attached Port Extension displayed as a link.
-

Detaching Port Extension from Virtual Machine Ports

To detach a Port Extension from Virtual Machine Ports:

-
- Step 1** Go to **Overlay > Virtual Machine Ports**. The Overlay / Virtual Machine Ports window appears.
 - Step 2** Select any tenant from the Tenant drop-down.
 - Step 3** Select the **Detach** icon. Detach Port Extension window appears.
 - Step 4** Select the **Zone**.
 - Step 5** Select a VRF from the **VRF** drop-down.
 - Step 6** Select the ports that you want to detach from the displayed list.
 - Step 7** Click **Detach**.
 - Step 8** Click **Yes** to confirm.

If the Port Extension detach fails, you can see the tool tips for failure message.

Assigning BVI Interface IP Address

To assign a Bridge Group Virtual Interface (BVI) IP address:

-
- Step 1** Go to **Overlay > Network**. The Overlay / Network page appears.
 - Step 2** Click the **Add (+)** icon. The Add Network page appears.
 - Step 3** Enter the Network name.
 - Step 4** Check the External Network check box.
 - Step 5** Click the **Add (+)** icon to assign a **Subnet** to the network created.
 1. If a Subnet is assigned to this External Network, assign the Router Gateway IP address for BVI interface from this Subnet under Step 10.
 2. If Subnet is not assigned to this External Network, any IP address can be assigned to Router Gateway IP address tab for BVI interface under Step 10.
 - Step 6** Go to **Overlay > Router**. The Overlay / Router page appears.
 - Step 7** Click the **Add (+)** icon. The Add Router page appears.
 - Step 8** Click the **Add (+)** icon to assign an **Interface** to the Subnet created.
Note: This subnet belongs to the Internal network, and excludes the External network.
 - Step 9** Select an external network from the **Router Gateway** drop-down list. Router Gateway IP address field appears.
 - Step 10** Assign the **Router Gateway IP address** for the selected external network for BVI interface and click **Save**.
 - Step 11** Verify whether the configuration is pushed to DCI and the IP address is assigned to BVI interface.
-

Extending Layer 2 Network Across Data Centers

If there are multiple data center PODs managed separately, (one instance of Cisco VTS managing only one POD) and connected over the WAN/core using a BGP-EVPN MPLS cloud, the L2VNI routes can be distributed from within the BGP-EVPN VXLAN fabric by stitching them to BGP-EVPN MPLS routes over the WAN/core side. On the other side (POD) the BGP-EVPN MPLS routes can be stitched onto BGP-EVPN VXLAN routes.

To complete the L2VNI extension workflow:

**Note**

- VTS supports a redundant DCI pair per data center.
- Both the DCIs in the ICCP pair must be added before any configuration prior to an L2VNI extension. If a redundant DCI is added mid-way where some L2 networks are already extended, configuration for the extended networks may not be synced to the new DCI.
- If fabric side supports ESI, you can enter the ESI number when you create the admin domain.
- Day Zero configuration has to be done on the ASR 9000 Series DCI device. See Day Zero Configuration Examples on Cisco.com for details.

-
- Step 1** Complete the Day Zero configuration with route policies/filters and DCI redundancy group.
 - Step 2** Go to **Admin Domain > DCI Interconnect profile**, and create an MPLS L2 VPN profile.
 - Step 3** Create the admin domain, add the MPLS L2 VPN profile. Extend L2 GW to DC GW.
After you save, the neighbor details are pushed under BGP.
 - Step 4** Go to **Overlay > Network**, click **Add (+)** under Fabric Host Networks.
 - Step 5** Use the Extend Network toggle switch to extend the network.
 - Step 6** Under the L2VPN tab, enter the EVI number.
 - Step 7** Specify the subnets, then do a port attach.
-

Enabling Global Route Leaking Service

The global route leaking feature enables you to provide internet/external connectivity to the host inside the Data Center. This feature allows associating/dissociating of Global Route Leaking (also known as Global Routing Table [GRT]) Service to/from the Overlay Router. Once the Overlay Router gets realized (that is, when port attach happens on interface), VTS pushes the policies configured as part of GRT associated to a router. Route policies for core facing/external facing routes and route policies for fabric facing/internal routes gets pushed.

**Note**

Global Route Leaking feature is available only when an external router gateway is selected.

Router cannot get deleted if the GRT is still attached. Admin needs to disassociate the GRT profile before deleting the router.

You can add create and enable global route leaking service while you create a router, or at any other point in time.

Step 1 Configure the import and export route policy on DCI and perform a *sync from*. For example:

```
route-policy data-center-vrf-export-policy
  if destination in (101.1.1.0/24 eq 32, 102.1.1.0/24 eq 32, 103.1.1.0/24 eq 32, 104.1.1.0/24
  eq 32, 105.1.1.0/24 eq 32) then
    pass
  endif
end-policy
!
route-policy data-center-vrf-import-policy
  if destination in (60.0.0.0/24) then
    pass
  endif
end-policy
```

See [Synchronizing Configuration](#) for details about performing a sync from operation.

Step 2 Create Fabric and Core Facing Route Policy (underlay policy for Internet connectivity). This is not mandatory. For example:

```
route-policy vts-route-policy
  pass
end-policy
```

Step 3 Create Profile for Internet from **Admin Domains > DCI Interconnect Profiles**. See [Creating DCI Interconnect Profiles](#).

Step 4 Attach the internet profile to DCI in the admin domain. Configuration is pushed by VTS on saving the admin domain. For example, the below configuration, which has the neighbor details, will be pushed under router BGP on the DCI.

```
router bgp 65540
  bgp router-id 18.18.18.18
  .
  .
  .
  neighbor 5.1.1.1
    remote-as 65544
    ebgp-multihop 255
    update-source Loopback2
    address-family ipv4 unicast
      route-policy vts-route-policy in
      route-policy vts-route-policy out
```

Step 5 Go to **Overlay > Router**. The Overlay / Router window appears.

Step 6 Click **Add (+)**. The Add Router page is displayed.

Step 7 Click **Global Route Leaking tab**.

Note Ensure that you have chosen an external router gateway as the Router Gateway.

Step 8 Click **Add (+)**. The New Global Route Leaking popup window appears.

Step 9 Enter a name (this is mandatory), and description.

Step 10 In the Policies pane, enter at least one policy for the address family.

Note Ensure that this policy exists on the device. Policy names gets validated from the device. If policy names are wrong, VTS will throw an error.

- Import Policy Name—Route policy to control import of routes from Global Routing Table (GRT).
- Export Policy Name—Route policy to control export of routes to GRT.

- Step 11** Click **Add**. The Global Route Leaking service gets added. You can click on the name to get a summary of the global route leaking service you created.
- Step 12** Click **Save**. Once the service is attached to the router, all the networks for the router will be leaked outside. To disassociate the service you need to select the **Detach** button and save the edit.

Enabling L3VPN to EVPN Route Stitching

L3VPN to EVPN route stitching feature provides the capability to exchange the routes from core towards the data center and vice versa. EVPN is used inside the data center whereas L3VPN is used as an interconnect between two data centers.



Note As a prerequisite, you must create an external network and extend to L3. You must then attach the router interfaces to the external network. See [Creating a Network using Cisco VTS GUI, on page 5](#) and [Creating Router using Cisco VTS GUI, on page 6](#) sections for details.

- Step 1** Configure BGP VPNv4/v6 neighbor using Device Templates. A single template can be used for all the neighbors, or you can have a template each for each neighbor. Create the template at **Templates > Device Template Management**. Attach the template to the DCI. See [Managing Templates and Device Objects](#) chapter for details.
- Step 2** Create an External Route Stitching Template. Choose the routes which you want to leak between your core and EVPN, or vice versa. Create the template at **Templates > Overlay Template Management** (use the Fabric External RT option). Attach the template to the DCI.

Adding Static Routes

You can add static routes to a router while you add or edit a router.

Procedure

	Command or Action	Purpose
Step 1	Add the following types of static routes:	<ul style="list-style-type: none"> • Internal Static Route—See Adding Fabric Static Routes, on page 18 for details. • External Static Route—See Adding External Static Routes, on page 18 for details. • Port Static Route—See Adding Port Static Routes, on page 19 for details.

	Command or Action	Purpose
		<p>If the static route is marked external, then it gets applied to the VRF on Border Leaf/DCI.</p> <p>This release supports Static Route - BFD feature which enables association of static routes with a static Bidirectional Forwarding Detection (BFD) configuration.</p> <p>During upgrade, all the routes in the route templates would be converted to static route objects and pushed to all leaf devices. Post-upgrade, if you need to modify these, you can remove these static route objects.</p>

Adding Fabric Static Routes

Static Routes can be configured directly on the router, and it will be pushed to all the nodes that have vrf for that router.

-
- Step 1** Go to **Overlay > Router**. Select the router you need to edit, and click the **Edit** icon.
 - Step 2** Click **Static Route**, then select **Fabric Static Route** from the Static Route Scope pane.
 - Step 3** Click **Add (+)** in the Internal Static Route pane, to add an internal static route. The New Internal Static Route popup is displayed.
 - Step 4** Enter the **Destination Prefix**—Destination IP address and the Subnet Mask.
 - Step 5** Enter the **Static Route Tag**. This is the BGP tag which is used for redistribution. This is optional.
 - Step 6** Enter the **Next Hop Details**—IP Address and Subnet Mask.
 - Step 7** Specify the **VRF**, if a router need to send traffic to a specific destination, via a specific next hop, and if that next hop is reachable only via a certain VRF.
 - Step 8** Specify the **Weight/Preference**, when you add multiple next hops.
 - Note** If you want the packet to a destination to be dropped, you can configure a null interface, by selecting the **Drop traffic that matches Destination Prefix** radio button.
 - Step 9** Click **Save**. The static route will be pushed to all the nodes.
- You can also edit the static routes you have created using the Edit option. To edit an internal static route, select the internal static route and click **Edit**.

While you upgrade to Cisco VTS 2.6.1 from an earlier version, and you had route templates in that version for which static routes were defined, all static routes are migrated and are displayed in this page.

Adding External Static Routes

External Static Route can be used when a router needs to send traffic outside of the fabric and it needs to reach to a specific destination to forward the traffic to that destination. These static routes are only pushed on border leaf, and DCI in case of integrated mode.



Note External Static Routes are available only when an external router gateway is selected.

- Step 1** Go to **Overlay > Router**. Select the router you need to edit, and click the **Edit** icon.
- Step 2** Click **Static Route**, then select **External Static Route** from the Static Route Scope pane.
- Step 3** Click **Add (+)** in the Internal Static Route pane, to add an internal static route. The New Internal Static Route popup is displayed.
- Step 4** Enter the **Destination Prefix**—Destination IP address and the Subnet Mask.
- Step 5** Enter the **Tag**. This is the BGP tag which is used for redistribution. This is optional.
- Step 6** Enter the **Next Hop Details**—IP Address and Subnet Mask.
- Step 7** Specify the **VRF**, if a router need to send traffic to a specific destination, via a specific next hop, and if that next hop is reachable only via a certain VRF.
- Step 8** Specify the **Weight/Preference**, when you add multiple next hops.
- Step 9** Click **Save**.

You can also edit the static routes you have created using the Edit option. To edit an external static route, select the external static route and click **Edit**.

Adding Port Static Routes

If a router needs to forward traffic to a prefix which is behind a VM, say a VTSR, it needs to have VTSR as the next hop. You can then configure a port static route in that scenario. The benefit is that if the VM moves across the nodes, the static route will also move across the nodes the where the vrf is present.

Port Scoped Static Route and Fabric static Route will be supported for virtual VTEP only when VTSR is configured as VTEP. VTSR will not support multiple next hops static routes.



Note For Port Scoped Static Routes, the static routes will be configured only on ToR which has port connection to it. And the static routes will be distributed to other ToR's in fabric via BGP EVPN.

- Step 1** Go to **Overlay > Router**. Select the router you need to edit, and click the **Edit** icon.
- Step 2** Click **Static Route**, then select **Port Scoped Static Route** from the Static Route Scope pane.
- Step 3** Click **Add (+)** to add new Static Route. The New Port Static Route page appears.
- Step 4** Enter the: a Port Scope Static route with destination with multiple Port Id, Next Hop IP.
- Destination Prefix—The destination (subnet) that you want to reach to.
 - Static Route Tag—This is an optional parameter.
 - Next Hop Details

- Port ID—The ID of the port (VM) which you want to reach. You can select the subnet and choose from the ports on the interfaces.
- Next Hop Prefix
- Weight—Specify the preference here. This is optional.
- BFD—Enable Bidirectional Fault Detection. This is supported only on Cisco Nexus 9000 Series devices. Port Scope Static Route supports BFD on both VTF (V-side) and (P-side).

Use the **Add (+)** icon to add more details.

Step 5 Click **Save**. Static Route is saved successfully.

Note After save, configuration is pushed to the selected device where the corresponding port (V or P) is connected. After save, configuration is pushed to the selected device where the corresponding port (V or P) is connected. The following is the sample config that is pushed to device for port scope static router.

```
vtshr01# show running-config vtsr-config ip-routes
vtsr-config ip-routes ip-route admin-rtr-1 217.217.217.217/32
scope port
paths path 101.1.1.67
  next-hop-vrf admin-rtr-1
  bfd true
  interface vhost-f8d76e3b-8943-4cde-90f1-df8141274aad
!
```

OpenStack Allowed Address Pairs Support

Cisco VTS supports OpenStack allowed address pairs feature. Allowed address pairs feature allows one port to add additional IP/MAC address pairs on that port to allow traffic that matches those specified values. See [OpenStack documentation](#) and [RedHat documentation](#) for details about this feature.

Important Notes:

- VPP adds default allowed address pair:



Note This is specific to VTF.

- IPv4 DHCP address—0.0.0.0/32 for each MAC
- From IPv6 Link-local Multicast IPs for IPv6 ND—ff02::/16 – 33:33:00:00:00:00 to IPv6 Link-local Multicast IPs for IPv6 ND—ff00::/12 – 33:33:00:00:00:00



Note If DHCP is used for IPv6 with allowed address pair, you should configure Link-local IPv6 address as allowed address pair from OpenStack.