Cisco Virtual Topology System (VTS) 2.5.1 User Guide

First Published: 2017-06-23

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
http://www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 527-0883

© 2017 Cisco Systems, Inc. All rights reserved.

CONTENTS

CHAPTER 1 **Introduction** 1

Understanding Cisco VTS 1

Cisco VTS Architecture Overview 3

Virtual Topology Forwarder 3

Virtual Topology System High Availability 4

CHAPTER 2 **Getting Started with Cisco Virtual Topology System** 5

Logging in 5

Using the Setup Wizard 6

System Settings 6

Virtual Machine Manager 7

Authorization Group 8

Discovery 9

Inventory 9

Using the Expert Mode—Quick Guide 11

Initial Configuration Tasks 11

Important Notes Regarding VMware vSphere Distributed Switch (vDS) 15

For Non-vPC Specific Configuration 15

For vPC Specific Configuration 16

CHAPTER 3 **Managing Local Users and Groups** 17

Creating Users 17

CHAPTER 4 **Managing Inventory** 19

Creating Authorization Groups 20

Performing Auto Discovery 20

Importing Inventory using CSV File 22

Day Zero Configuration for IPv6 Management on vCenter 24

Viewing the Network Topology 24
Viewing Network Inventory 25
Adding Fabric Connection 26
Synchronizing Configuration 26
Viewing Host Inventory 27
Adding a new Host on Virtual Servers 28
Adding a new Host on Baremetal 29
Viewing the VTSR to VTF Mapping 30
SRIOV Support 30
Migrating from vPC to ESI 30
Redeploying Device Inventory 33
Enabling Static Multi Homing 34
Enabling Static Multi Homing on Cisco Nexus 7000 34
Enabling Static Multi Homing on Cisco Nexus 9000 34

CHAPTER 5 Managing Resources 37
Specifying Global VNI Range 38
Specifying VLAN Range 38
Specifying Device VLAN Range 39
Specifying Group VLAN Range 40
Specifying Multicast IP Pool 40
Specifying Interface VLAN Range 41
Creating Interface Groups 41
Important Note 43

CHAPTER 6 Creating and Managing Admin Domains 45
Admin Domain Overview 45
Viewing Admin Domain 46
Creating an Admin Domain 46
Creating DCI Interconnect Profiles 50

CHAPTER 7 Managing Templates 53
Creating Route Templates 55
Adding Static Routes 55
Adding Route Targets 55

Adding Fabric Internal Route Targets 56
Adding Fabric External Route Targets 56
Disabling Auto Route Target Configuration 57
Creating L3 Extension Templates 58
Supported System Variables 58
Editing Templates 59
Copying a Template 60
Deleting Templates 60
Attaching Templates to Tenants 60
Attaching Template while Adding Tenants 61
Attaching Template while Editing a Tenant 61
Detaching Template from Tenants 61
Attaching Templates to Routers 62
Attaching Templates while Adding Routers 62
Attaching Template while Editing a Router 63
Creating L2 Extension Templates 63
Important Notes—L2 QoS Template 64
Important Notes—VPLS Template 64
Associating L2 Extension Templates 65
Creating Underlay Templates 65
Associating Devices to Underlay Template 66

CHAPTER 8 **Managing Tenants** 67
Viewing Tenant Details 67
Adding Tenants 68
Editing Tenants 68

CHAPTER 9 **Provisioning Overlay Networks** 69
Provisioning Overlay Networks Using Cisco Virtual Topology System 69
Creating Overlays 71
Using OpenStack 71
Using VMware 71
Using Cisco VTS GUI 71
Creating Network using VMware 72
Creating Subnetwork using VMware 72

Creating Routers using VMware 72
Attaching Network to Router 73
Attaching a Virtual Machine to Network 73
Creating a Network using Cisco VTS GUI 73
Creating a Subnetwork 74
Attaching a Port 74
Creating Router using Cisco VTS GUI 75
Assigning BVI Interface IP Address 77
Extending Layer 2 Network Across Data Centers 77
Enabling Global Route Leaking Service 78
Enabling L3VPN to EVPN Route Stitching 80

CHAPTER 10 Viewing Overlay Details 81
Viewing Device Details 81
Viewing Virtual Machine Details 82
Viewing Baremetal Details 82
Viewing Network Details 83
Viewing Router Details 84

CHAPTER 11 Administering Cisco VTS 85
Setting up System 85
Modifying Login Banners 86
Setting Global Route Reflector 87
Registering the Virtual Machine Manager using GUI 87
Integrating Cisco VTS with Multiple Virtual Machine Managers 89
Performing Merge Operation from VMM 92
Performing Merge Operation from Tenant 93
Performing Merge Operation from Network 94
Performing Publish Operation from VMM 94
Performing Publish Operation from Tenant 95
Performing Publish Operation from Network 95
Performing Publish Operation from VTS 95
Backing up the Database 96
Restoring the Database 97
Viewing HA Status 97

Enabling External Authentication and Authorization 98
Setting up Remote Authentication Server 99
Enabling Accounting and Logging 99
Setting up Accounting 100

APPENDIX A Service Extension Templates-Supported Configuration 101
Supported Service Extension Template Configuration Examples for Cisco Nexus 7000 Series Switches 101
Supported Service Extension Template Configuration for Cisco Nexus 9000 Series Switches 104
Supported Service Extension Template Configuration for Cisco ASR 9000 Series Routers 105

APPENDIX B Supported Underlay Configuration 109
Supported Underlay Configuration 109

APPENDIX C OpenStack Configuration for SR-IOV Support 131

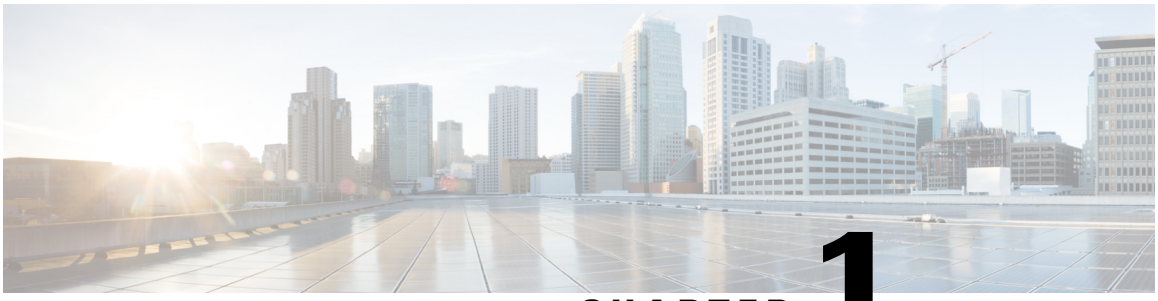

CHAPTER 1

Introduction

This chapter provides an overview of Cisco Virtual Topology System (VTS). It has the following sections:

- Understanding Cisco VTS, page 1
- Cisco VTS Architecture Overview, page 3
- Virtual Topology Forwarder, page 3
- Virtual Topology System High Availability, page 4

Understanding Cisco VTS

The Cisco Virtual Topology System (VTS) is a standards-based, open, overlay management and provisioning system for data center networks. It automates DC overlay fabric provisioning for both physical and virtual workloads.

Cisco VTS provides a network virtualization architecture and software-defined networking (SDN) framework that meets the requirements of multitenant data centers for cloud services. It enables a policy-based approach for overlay provisioning.

Cisco VTS automates complex network overlay provisioning and management tasks through integration with cloud orchestration systems such as OpenStack and VMware vCenter and abstracts out the complexity involved in managing heterogeneous network environments. The solution can be managed from the embedded Cisco VTS GUI or entirely by a set of northbound Representational State Transfer (REST) APIs that can be consumed by orchestration and cloud management systems.

Cisco VTS provides:

- Fabric automation
- Programmability
- Open, scalable, standards based solution
- Cisco Nexus 2000, 3000, 5000, 7000, and 9000 Series Switches. For more information, see Supported Platforms in *Cisco VTS 2.5 Installation Guide*.
- Software forwarder (Virtual Topology Forwarder [VTF])

VTS performs the role of an overlay orchestrator in data-center networks. In this role, it would manage configuration on the data center leaf and spine devices. The configuration of the devices is dependent on the type of overlay service that the Cisco VTS user intends to create. The Cisco VTS user in this context could either be manual users interfacing via GUI or APIs, or could be virtual machine managers like OpenStack or vSphere. Since the device configuration is derived from overlay service instances, Cisco VTS holds the 'desired' device configuration in its database.

Whenever, there is a change to the overlay service instances, it generates desired device configuration and applies them to the relevant set of devices. This is the prime functionality of an orchestrator. Changing any of the device configuration outside of Cisco VTS (For example, using CLI or other programmatic interfaces to the device), can result in service disruption. Hence Cisco VTS always reconciles its view of the device configuration and pushes that to the devices. Cisco VTS holds the master database of all device configuration in the fabric.

However, there are some practical use cases where Cisco VTS accommodates out-of-band device configuration.

- Day0 underlay configuration—Cisco VTS is an overlay manager, but overlays cannot be established without an underlay. Underlay configuration on each leaf/spine device is unique. Typically, such underlay configuration is laid out even before Cisco VTS can manage the overlays. Assuming all the devices in the fabric are physically connected, the data center administrator establishes the underlay configuration by manually connecting to the devices and configuring them OR using an underlay manager to perform this. When Cisco VTS performs a scan of the fabric inventory and discovers the topology, it is expected that all the underlay configuration has been fully established. At this point, Cisco VTS treats all the pre-existing device configuration to be Day0 configuration. Day0 configuration is synced up from the devices and stored in VTS database as a baseline. All overlay service configuration is built on top of this day-0 device configuration.
- Day*N* underlay configuration—While Cisco VTS manages overlay specific device configuration, there is always the need for the fabric operators to customize device underlay configuration. Typical operations include physical link management, applying link specific features, managing underlay routing protocols, security settings and the like. Recognizing this need, Cisco VTS supports the concept of 'device' templates. These are essentially device configuration parameters exposed to the VTS user using GUI/APIs. VTS users can customize device configuration using device templates and use that to create the consolidated device configuration.
- Day*N* overlay configuration—While the overlay specific configuration pushed by Cisco VTS is sufficient to establish overlays, every deployment may require some customization around this configuration. Since VTS holds the master device configuration, it is essential that any customization flows through Cisco VTS. To address this, Cisco VTS supports the concept of a 'service' template. Service templates allow the Cisco VTS user to extend the service specific device configuration via GUI/API.

Note Service templates always 'augment' the configuration. They cannot modify or remove configuration that is constructed by the VTS service layer.

We recommend that you do not modify device configuration using outside of VTS. Doing so, can result in misconfiguration of devices and will result in service outage. If there is a real need to do so, you may follow one of the three models of device configuration to achieve the desired customization.

Cisco VTS Architecture Overview

Cisco VTS architecture has two main components: the Policy Plane and the Control Plane. These perform core functions such as SDN control, resource allocation, and core management function.

- **Policy Plane**: The policy plane enables Cisco VTS to implement a declarative policy model designed to capture user intent and render it into specific device-level constructs. The solution exposes a set of modular policy constructs that can be flexibly organized into user-defined services for use cases across service provider and cloud environments. These policy constructs are exposed through a set of REST APIs that can be consumed by orchestrators and applications to express user intent, or instantiated through the Cisco VTS GUI. Policy models are exposed as system policies or service policies.

  System policies allow administrators to logically group devices into pods within or across data centers to define Admin Domains with common system parameters (for example, BGP-EVPN control plane with distributed Layer 2 and 3 gateways).

  The inventory module maintains a database of the available physical entities (for example, data center interconnect [DCI] routers and top-of-rack leaf, spine, and border-leaf switches) and virtual entities (for example, VTFs) in the Virtual Topology System domain. The database also includes interconnections between these entities and details about all services instantiated within a Virtual Topology System domain.

  The resource management module manages all available resource pools in the Virtual Topology System domain, including VLANs, VXLAN Network Identifiers (VNIs), IP addresses, and multicast groups.

- **Control Plane**: The control plane module serves as the SDN control subsystem that programs the various data planes including the VTFs residing on the x86 servers, hardware leafs, DCI gateways. The Control plane hosts Service Routing (SR) module, which provides routing services to Cisco VTS. The Service Routing (SR) module is responsible for calculating L2 and L3tables and routes to provide connectivity between the different VMs for a given tenant and service chaining. The main components of this module are the VTSR and VTF. VTSR is the controller and Virtual topology forwarder (VTF) runs on each compute server hosting the tenant VMs.

Virtual Topology Forwarder

Virtual Topology Forwarder (VTF) runs on each compute server in the DC and provides connectivity to all tenant VMs hosted on the compute server. VTF supports both intra and inter DC/WAN connectivity. VTF allows Cisco VTS to terminate VXLAN tunnels on host servers by using the VTF as a Software VXLAN Tunnel Endpoint (VTEP). Cisco VTS also supports hybrid overlays by stitching together physical and virtual endpoints into a single VXLAN segment.

VTF has 2 major components—Cisco's VPP (Vector Packet Processing)and VPFa. VPFA is a Cisco agent running on each VMM compute resource. VPFA is FIB agent which receives L2/L3 table forwarding information from VTSR needed to provide the connectivity to local tenant VMs hosted on its compute, and programs them in the VPP.

VTF is deployed as a virtual machine or in vhost mode, to deliver a high-performance software data plane on a host server.

Virtual Topology System High Availability

The Virtual Topology System solution is designed to support redundancy, with two solution instances running on separate hosts in an active-standby configuration.

During initial setup, each instance is configured with both an underlay IP address and a virtual IP address. Virtual Router Redundancy Protocol (VRRP) is used between the instances to determine which instance is active.

The active-instance data is synchronized with the standby instance after each transaction to help ensure consistency of the control-plane information to accelerate failover after a failure. BGP peering is established from both Virtual Topology System instances for the distribution of tenant-specific routes. During the switchover, nonstop forwarding (NSF) and graceful restart help ensure that services are not disrupted.

See the *Installing VTS in High Availability Mode* section of the *Cisco VTS 2.4.2 Installation Guide* for the detailed procedure about setting up high availability.

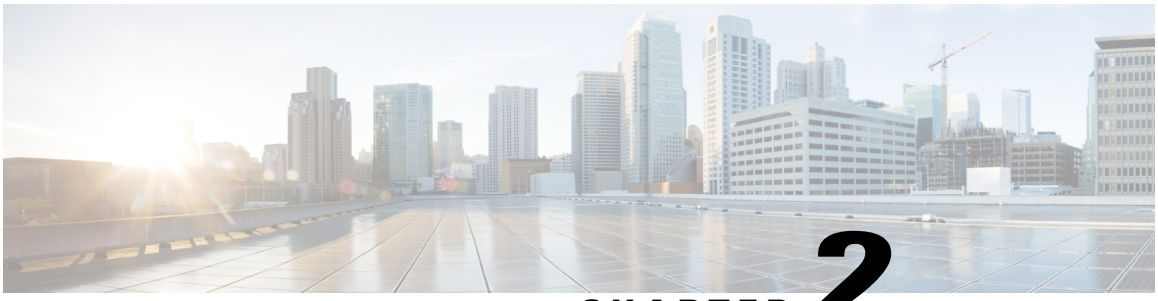

CHAPTER 2

Getting Started with Cisco Virtual Topology System

This chapter provides an overview of Cisco Virtual Topology System (VTS). It also provides a high level workflow of the tasks that you need to perform after you install Cisco VTS.

- Logging in, page 5
- Using the Setup Wizard, page 6
- Using the Expert Mode—Quick Guide, page 11
- Initial Configuration Tasks, page 11
- Important Notes Regarding VMware vSphere Distributed Switch (vDS), page 15

Logging in

To log in to the Cisco VTS GUI:

Step 1 Open a supported browser, and enter the URL of the server. For example: http://<*IP Address*>.
Cisco VTS supports Google Chrome and Mozilla Firefox browsers. Cisco VTS has been tested on Google Chrome version 49.0.2623.112 m and Mozilla Firefox version 45.0.1.

Step 2 Enter the username and passphrase, and click **Login**. The default username and passphrase is admin/admin.
The Change Password window appears.

You will be required to change the passphrase for the admin account the first time you are logging in. Click the Passphrase guidelines link for details about the passphrase guidelines.

Step 3 Enter the **New Passphrase**, and reenter the new passphrase in the **Confirm New Passphrase** field.

Step 4 Click **Change Passphrase**.

Note
- To change your passphrase subsequently, click **Change Passphrase** on the top right settings button.

The Cisco VTS welcome screen is displayed. The screen provides two options for you to continue with the set up tasks:

- Guide me to set up—This option takes you to the Cisco VTS Setup Wizard. This wizard guides you through the initial configuration tasks to get your system ready for further operations.
- I will set up myself—This option takes you to the Quick Guide, which displays the tasks you need to complete in order to get started with the system.

  If you are familiar with the Cisco VTS setup tasks, you may opt to close the Quick Guide and proceed with the tasks. You can access the Quick Guide anytime from the Settings menu on the top right corner of the Cisco VTS GUI.

Using the Setup Wizard

The Cisco VTS Setup wizard guides you through a series of steps that helps you set up the system and get started. We recommend that you use the Setup wizard if you are not familiar with the VTS UI and the sequence of the initial configuration tasks.

Note Not all steps in the wizard are mandatory. Green check mark indicates that the task is not mandatory. These tasks can be done outside of the wizard flow or have default values.

The following sections provide information about the series of screens that are part of the wizard.

System Settings

This is the initial screen in the Setup wizard. This helps you to set up the following system parameters:

- DHCP Server IPv4
- DHCP Server IPv6
- AnyCast Gateway Mac
- Out of Sync Commit Behavior
- Default Range for VLAN pools
- Multiple VNIs to Multicast Address Mapping

Enter the values, and click Next.
The Virtual Machine Manager screen is displayed.

Virtual Machine Manager

The Virtual Machine Manager screen in the Setup wizard helps you to register the VMM. You can also specify whether the VMM you register is a trusted or an untrusted VMM.

Note For cluster-based deployments, you must install the plugin on each node.

Step 1 Click the Add (+) button.
The Add Virtual Machine Manager popup is displayed.

Step 2 Enter the following details:

- VMM Name—Name of the VMM.
- VMM Type—Specify the VMM type. Choose vcenter or openstack from the drop-down list.
- Version Name—Specify the version.
- Datacenter—The name of the datacenter for which Cisco VTS acts as the controller. This is applicable only for vCenter.
- IP Address/Port—The IPv4/IPv6 address of the VMM end point. The default port is 443. This field is applicable only for vCenter.
- Identity IP Address—The IPv4/IPv6 address of the VMM end point. This field is applicable only for OpenStack.
- Protocol—Specify whether the protocol is http or https.
- HTTPS port—In case of vCenter, the vShpere webclient https port. For OpenStack, the OpenStack Neutron service port.
- Keystone port—Keystone port for OpenStack.
- VTF Deployment Mode—If you select openstack-liberty-rhel as the OpenStack version, this option appears. To install VTS in vhost mode, choose vtf-vhost. To install VTF as a VM, choose vts-as-vm.
- User Name—Enter the VMM username.
- Passphrase—Enter the VMM passphrase.
- Mode—Specify whether the VMM is a trusted VMM or an untrusted VMM. This field cannot be edited.
- Openstack Admin Project —The admin project in OpenStack.
- Admin User Name— admin user for the admin project in OpenStack.
- Admin Password—Password of the admin user.

Step 3 Click Add
The VMM you added is listed in the Virtual Machine Manager screen.

You can check the status of VMM registration in the Status column.

Step 4 Click Next if the VMM registration is successful.
To delete a VMM, select the VMM and click **X** (delete).

Note For the unregistration to be complete, you need to run the manual clean up procedures. See <section/awaiting inputs> for details

The Administration > Virtual Machine Manager page displays the following details:

- VMM Name
- VMM Type
- Version
- Openstack Admin Project
- IP Address:Port
- Mode—Whether the VMM has been registered as Trusted or Untrusted.
- Username
- Status—The VMM registration status.
- MultiVMM Operations—Whether Merge, Publish, and Publish from VTS are allowed. Click the respective icon to initiate a Merge/Publish operation.

Authorization Group

The Authorization Group screen in the Setup wizard helps you to create authorization groups. The authorization group will have credentials for logging into your devices. You can create as many authorization groups as necessary.

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

- **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 authorization group 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.

Discovery

The Discovery screen helps you to discover the devices in your network, and add them to the inventory. You can perform auto discovery, or import devices manually using a .CSV file. Before you perform this step, ensure that:

- 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.
- 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.
- All devices must have the same set of credentials. These credentials will be used during the discovery process. See Managing Inventory, on page 19 for more information. The credentials must be of the appropriate privilege level on the devices.

To perform discovery:

Step 1 Click **Discovery**.

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 **Start Discovery**.

Note After the discovery is complete, you may edit the editable fields in the table to modify the device details, before you add the details into the Cisco VTS inventory.

To import device details into inventory via CSV file, manually, click the **Import CSV**, then select the CSV file and import it.

If you want to edit details for more than one devices, you may select the devices, and then click **Bulk Edit** option. Click Apply after you make the changes.

Step 6 To add the devices to Cisco VTS inventory, select the devices and click **Add to Inventory** button.

Note When you add the devices to inventory using the Add to Inventory button, all the devices in the table are added to the inventory.

Step 7 Click **Next**.

The Inventory screen is displayed.

Inventory

The Inventory screen of the Setup wizard displays view inventory details: It has three tabs:

- Network Inventory—Details about the switches in the inventory. It shows the following details.
  - Device Name
  - Admin State
  - IP Address

- Auth Group
- Device Platform
- Device Role
- Group Tag
- bgp-asn
- Loopback Interface Number
- Loopback Interface IP

Note 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.

- Fabric Inventory—Details about the fabric connection between the switches. It shows the following details:
  - Device Name
  - Device Type
  - Device Interface
  - Device IP Address
  - Device Port Name
  - Connection ID
- Host Inventory—Details about the servers connected to the switches It shows the following details:
  - Host Name
  - Host Type
  - Host Interface
  - Host IP Address
  - Device Port Name
  - Connection ID
  - Capability—The host capability - Physical or Virtual.
  - Status—Shows the installation status.

  Click on the info icon adjacent to the **Host Name**. Host Summary table popup appears.

  If Virtual Capability is '**virtual-switch**', both **Additional Host Details** and **VTF Details** can be viewed in the Host Summary table.

  If Virtual Capability is '**no-virtual-switch**' only **Additional Host Details** can be viewed in the Host Summary table.

Click **Finish** to exit the Setup wizard.

Using the Expert Mode—Quick Guide

On logging in for the first time, if you choose the **I will set up myself** option, the Quick Guide appears.

Note You may opt to close the Quick Guide and proceed to the set up tasks on your own, via the Cisco VTS GUI. To get a list of tasks that need to be performed to set up and get started with Cisco VTS, see Initial Configuration Tasks, on page 11 section. At any time, you can access the Quick Guide from the settings menu on the top right corner of the Cisco VTS GUI.

The Quick Guide has the following tasks listed:

Task	Subtasks	Doc Section
Administration	• System Settings • Virtual Machine Manager	• Setting up System, on page 85 • Registering the Virtual Machine Manager using GUI, on page 87
Set up Inventory	• Discover Devices	• Managing Inventory, on page 19
Admin Domains	Create Admin Domains	Creating an Admin Domain, on page 46
Set up Tenants	Add Tenants	Adding Tenants, on page 68
Set up Overlay	• Add Network • Define Baremetal • Add Virtual Machines • Add Routers	Provisioning Overlay Networks, on page 69

Initial Configuration Tasks

After bringing up the Virtual Topology Controller (VTC) Virtual Machine (VM), do the following:

Note Before you perform the tasks below, ensure that installation is complete, day zero configuration on leafs is done, and all underlay configurations are working.

Sequence	Task	Navigation in VTS GUI / User Guide Section	Additional Notes
1	Create an Authorization Group	Inventory > Authorization Group For more information about creating authorization group, see Managing Local Users and Groups, on page 17	This will have credentials for logging into your devices. You can create as many auth groups as necessary.
2	Discover the Topology of all the leafs, spine, border-leafs, and DCI	Discovery > Topology Discovery For more information about adding devices and host information, see Performing Auto Discovery, on page 20	VTFs are not detected in topology discovery.
3	Import the devices after adding the auth group	Inventory > Import Inventory For more information about adding devices and host information, see Importing Inventory using CSV File, on page 22	
4	Add the Domain ID, DHCP Server IP, and Anycast Gateway MAC	Administration > System Settings For more information about adding devices and host information, see Administering Cisco VTS, on page 85.	

Sequence	Task	Navigation in VTS GUI / User Guide Section	Additional Notes
5	Perform VTSR and VTF Registration. **Note** This step is required only if you have a VTF-based deployment.	See the *Installing the Virtual Topology Forwarder* section in the *Cisco VTS Installation Guide*.	Register the VTF only after you see the VFG group on the GUI. To verify that the VFG group is created, go to Inventory > Virtual Forwarding Groups.
6	Login into the VTSR and verify that the network-controller config is pushed **Note** This step is required only if you have a VTF-based deployment.		
7	Update the BGP ASN information for the devices	Inventory > Network Inventory For more information, see Viewing Network Inventory, on page 25.	
	Create resource pools	• Resource Pools > Global VNI Pool • Resource Pools > Device Specific VLAN Pools • Resource Pools > Multicast IP Pool For more information about creating an admin domain, see Managing Resources, on page 37	

Sequence	Task	Navigation in VTS GUI / User Guide Section	Additional Notes
8	Create an Admin Domain	Admin Domains > Domains For more information about creating an admin domain, see Creating an Admin Domain, on page 46	Properties for the L2/L3 Gateway Group are as follows: • Control Protocol: BGP-EVPN • Replication Modes : Multicast and Ingress • Distribution Mode: Decentralized
9	Add the devices to the Gateway Group	Admin Domains > Domains For more information about creating an admin domain, see Creating an Admin Domain, on page 46	See the *Supported Platforms* section in the *Cisco VTS Installation Guide* for details about devices support for different roles.
10	Add the ToR and VTSR to the L2 and L3 Gateway Group	Admin Domains > Domains For more information about creating an admin domain, see Creating an Admin Domain, on page 46	

Sequence	Task	Navigation in VTS GUI / User Guide Section	Additional Notes
11	Save the Admin Domain you created	Admin Domains > Domains For more information about creating an admin domain, see Creating an Admin Domain, on page 46	

Important Notes Regarding VMware vSphere Distributed Switch (vDS)

The following points need to be taken care of while you create a vDS.

Note

- All the ToRs in the inventory should be part of the vDSs.
- One vDS can represent one or more ToRs.
- All the hosts that are connected to a particular ToR should be part of the same vDS.

For Non-vPC Specific Configuration

If you are not using vPC on the leaves:

- Associate one or more leafs per vDS.
- Attach the hosts' data interface to the vDS uplinks.

Note See VMware documentation for the detailed procedure.

For vPC Specific Configuration

If you are using vPC on the leaves:

Step 1 Create one vDS switch for one or more vPC pairs.

Step 2 Enable enhanced LACP.
See VMware documentation for the detailed procedure.

Step 3 Create a Link Aggregation Group for each vDS.
See VMware documentation for the detailed procedure.

Step 4 You may remove the default port group that gets created as it will not be used.

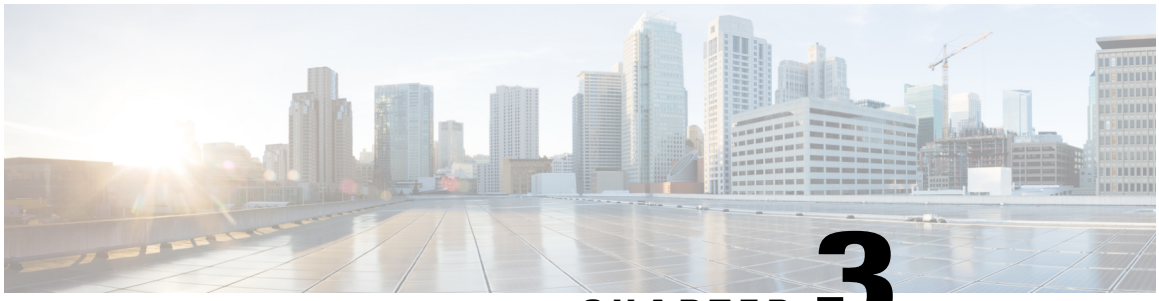

CHAPTER 3

Managing Local Users and Groups

You can create users and define the roles they have to control the access to Cisco VTS GUI operations.

See the Managing Inventory chapter for details about importing inventory.

This chapter has the following sections:

- Creating Users, page 17

Creating Users

You can create users to define the role that the users have when they log in to Cisco VTS. There are two default roles available:

- Administrator
- Operator
- ncsadmin—Has the same permissions as Administrator.
- ncsoper—Has the same permissions as Operator.

To create users:

Step 1 Click **Administration** > **User Management**. The Administration / User Management window appears.

Step 2 Click **Add (+)** icon. The Add New User popup window appears.

Step 3 Enter the **User Name** and **Passphrase**, and then select the desired role from the Role drop down list.

Step 4 Click **Save**.
The user details get added to the Users table.

Note To edit the user name, check the User Name check box, click **Edit** icon.

To delete the user name, check the User Name check box, click **Delete (X)** icon.

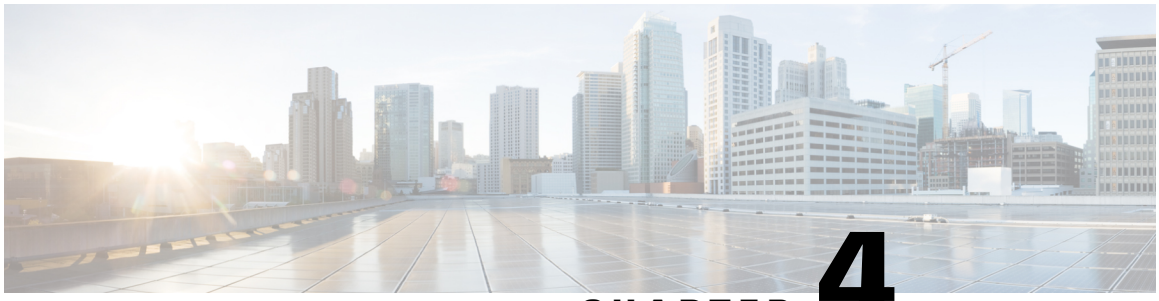

CHAPTER 4

Managing Inventory

You can discover the devices using the auto discovery option, and import the details into Cisco VTS. You can also manually create a CSV file with device details, in a prescribed format, and import it into Cisco VTS.

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 (i.e. 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.

This chapter has the following sections:

- Creating Authorization Groups, page 20
- Performing Auto Discovery, page 20
- Importing Inventory using CSV File, page 22
- Viewing the Network Topology, page 24
- Viewing Network Inventory, page 25
- Viewing Host Inventory, page 27
- Viewing the VTSR to VTF Mapping, page 30
- SRIOV Support, page 30
- Migrating from vPC to ESI, page 30
- Redeploying Device Inventory, page 33
- Enabling Static Multi Homing, page 34

Creating Authorization Groups

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

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.

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.

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.

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 19 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 Click **Discovery** . The Discover Network popup window appears.

Step 3 Enter the **Seed Device IP**.

Step 4 Enter the **Seed Device User Name**.

Step 5 Enter the **Seed Device Passphrase**.

Step 6 Click **Start Discovery**.

Step 7 After the auto discovery is complete, the details are displayed in the table.
You can use the **Bulk Edit** option to modify the device details for more than one devices, before you add the devices to inventory, using the **Add to Inventory** button. To avoid errors, make sure you review the details in the table before you add the devices to Inventory. The Cisco VTS discovery log file is under /var/vts/log. Check for any errors/exceptions in this log file.
If, during auto discovery, Cisco VTS fails to log into any of the nodes, then the information of the nodes along with the error details can be found in /var/vts/log.

When you add devices to the inventory using the Add to Inventory button, all devices in the table are added to inventory.

Note 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.

Note 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.

Complete LLDP settings for OpenStack compute host:

If 'lldpad' application is used on compute for LLDP then the compute may not send interface-name and management-ip-address TLVs by default. These need to be explicitly enabled using the below script. On the other hand, if 'lldpd' were to be used to run LLDP then these TLVs are sent by default and should not require any additional settings prescribed below.

For 'lldpad', follow below steps to facilitate compute host exchange interface name as Port-ID and management ip address with connected ToR: (Below instructions should go into compute host's recommended Day0 configuration.)

1 Execute the script:

```
for i in `ls /sys/class/net/ | grep eth` ;

        do echo "enabling lldp for interface: $i" ;

        lldptool set-lldp -i $i adminStatus=rxtx  ;
```

```
lldptool -T -i $i -V  sysName enableTx=yes;

lldptool -T -i $i -V  portDesc enableTx=yes ;

lldptool -T -i $i -V  sysDesc enableTx=yes;

lldptool -T -i $i -V sysCap enableTx=yes;

lldptool -T -i $i -V mngAddr enableTx=yes;
```

2 Execute below command:

```
# lldptool -T -i <mgmt-if-name> -V mngAddr ipv4=<mgmt-ip>
```

For example:

```
# lldptool -T -i eth2 -V mngAddr ipv4=192.168.24.185
```

3 Restart lldp daemon.

```
service lldpad restart
```

If the above steps have not been performed through a startup script during compute startup, then lldp discovery might have already happened and the corresponding ToR might not reflect the mgmt-ip of the compute. To make this happen disable and then enable lldp on the ToR port using below commands:

```
1. #config t
2. (config)#interface ethernet slot/port
3. (config)#[no] lldp transmit
4. (config)#[no] lldp receive
5. (config)#(Optional) show lldp interface ethernet slot/port
```

Importing Inventory using CSV File

The device inventory details need to be uploaded using a CSV file, in the prescribed format. The inventory 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 need to delete all the devices from network inventory and re-upload the file after fixing the issues. Or else, you may encounter problems due to the partial upload.

To download a sample inventory file, click **Download Template**.

Note In a VMware environment, each time you add a leaf, you must create a corresponding VMware vSphere Distributed Switch (vDS). See the Important Notes Regarding VMware vSphere Distributed Switch (vDS), on page 15 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.

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

Step 2 Click **Import CSV**
The Import Devices popup window appears.

Step 3 Click **Browse** to choose the CSV file.
Click **Download Template** to download a sample CSV file.

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.
- server-type—virtual-server for computes; baremetal for connections to spine, DCI.
- interface-name—Physical port connectivity (interface 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.

Step 4 Browse for the CSV file, and click **Import**.

Step 5 After the import operation complete, the details are displayed in the table.

You can use the **Bulk Edit** option to modify the device details for more than one devices, before you add the devices to inventory, using the **Add to Inventory** button. To avoid errors, make sure you review the details in the table before you add the devices to Inventory.

Click the device name to open the connected devices popup window, which gives details about connected devices for a device.

Day Zero Configuration for IPv6 Management on vCenter

Step 1 Make sure you configure IPv6 and default IPv6 route on all Cisco Nexus 9000 and Cisco Nexus 7000 series devices. For example

```
Leaf1# show run interface mgmt 0

!Command: show running-config interface mgmt0
!Time: Tue Apr 11 06:32:35 2018

version 7.3(1)D1(1)

interface mgmt0
vrf member management
ip address 172.23.92.16/24
ipv6 address 2001:422:10e:2005:c00::88/70
#########  default ipv6 route  #####
vrf context management
ip route 0.0.0.0/0 172.23.92.1
ipv6 route 0::/0 2001:420:10e:2015:c00::1
```

Step 2 Log in to the vCenter Server Appliance on the Controller (vCenter 5.5 or 6.0), and configure IPv6 (Network > Address), and save the settings.

Step 3 For ESXI Host, configure IPV6 through CIMC.

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
- Admin State
- IP Address
- Auth Group
- Device Platform
- Device Role
- Group Tag
- BGP-ASN
- Templates Attached
- Sync

You need to add the BGP-ASN information, loopback interface number, and loopback interface id of every device prior to starting any provisioning.

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.

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 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 state is in sync with VTS configuration database.

To check whether the configuration is in sync:

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 **Config Sync** link under the Sync column, for the device.
If the configuration is in sync with the VTS database, the green In Sync icon is displayed.

If the configuration is not in sync with the device, you will get an option to compare the configuration. You can compare the actual device configuration with the VTS configuration, and can view the diff in compare-config popup. You can also use the copy to clipboard option to copy and paste the configuration in a text editor. To synchronize the configuration, you can use the following options:

- sync-from—Synchronize the VTS configuration by pulling configuration from the device.
- sync-to—Synchronize the VTS configuration by pushing configuration to the device.

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

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 left 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.
- Host Interface—IPv4/IPv6 address of the host. This is mandatory.
- Host IP Address
- Device Port Name
- User Name
- Passphrase
- Host Configuration
  - VMM ID—The VMM ID of the VMM to which you want to associate the host to.
  - Virtual Switch—The following options exist:
    - not defined
    - ovs—If you want to install the VTF host agent on the compute, check the Install VTS agent on save check box.
    - vtf-l2—VTF is used as an L2 switch. If you choose this option, you need to install VTF in vhost mode. See *Cisco VTS Installation Guide* for details about installing VTS in vhost mode. After you install VTF in vhost mode, you must refresh the Host Inventory page for the details to get updated here.
    - vtf-vtep—On OpenStack, you need to install VTF as a vhost, if you choose 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 Enter the Host Interface details. At least one interface is mandatory.

- Host Interfaces—This is mandatory.
- SRIOV Enabled—Choose Yes or No from the drop down to specify whether the interface is SRIOV enabled.
- Phys Net— Physnet name associated with the interface. If *Sriov 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 4 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.
- SRIOV Enabled—Choose Yes or No from the drop down to specify whether the interface is SRIOV 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 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

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

SRIOV Support

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

- SRIOV + OVS
- SRIOV + VTF as L2 Switch
- SRIOV + SRIOV

SRIOV 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, on page 37 chapter for details about assigning VLAN ranges.

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

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 33 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 in Admin Domain.

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 26 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.

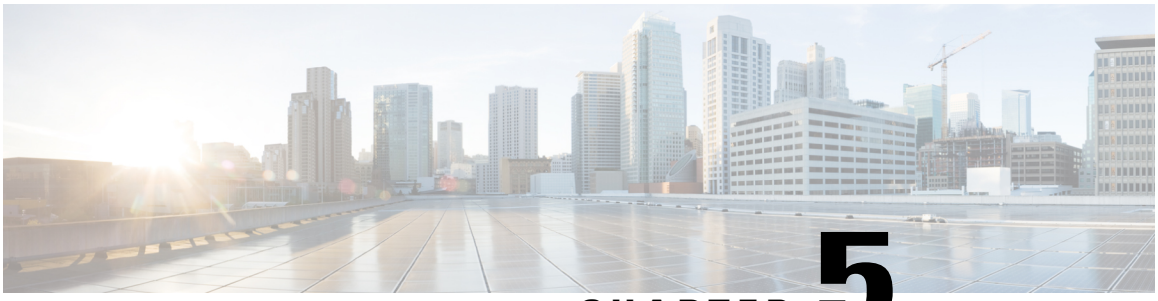

CHAPTER 5

Managing Resources

Cisco VTS enables you to define global and device local resource pools so that the resources can be allocated when VMs or tenants are added.

You can define the multicast addresses that can be allocated sequentially for VXLAN tunnel creation so that you do not have to manage them separately. VNI and multicast pools are global.

You can define the multicast addresses as global resource pool. These multicast addresses will be allocated sequentially when new VXLAN tunnels are created. Network administrator can define multiple multicast address ranges.

VLAN range can be assigned or each device. You can also group devices and assign VLAN range to the device group.

Additionally, for Nexus 7000 devices, you can assign VLAN resources at physical or FEX interface level. You can also group the interfaces from different devices and assign VLAN ranges to the interface group.

Note Default resource pools are device-specific VLAN pools that are also created automatically when leafs are added to the inventory. The default VLAN range is from 1001 to 2000. You can modify the range as per your requirement.

You can edit the range and also delete any unused ranges.

This chapter has the following sections:

- Specifying Global VNI Range, page 38
- Specifying VLAN Range, page 38
- Specifying Multicast IP Pool, page 40
- Specifying Interface VLAN Range, page 41
- Creating Interface Groups, page 41

Specifying Global VNI Range

You can specify the global VNI range. To do this:

Step 1 Go to **Resources** > **Global VNI Pool**. The Resources / Global VNI Pool window appears.
In Global VNI Pool window, the range table lists the following details:

- Range From
- Range To
- Restricted Range - data is Boolean (Yes or No)
- Used
- Available
- Total

Step 2 Click the **Add (+)** icon. Add VNI Pool popup appears.

Step 3 Specify the ranges, select the **Restricted Range** radio button to enable or disable the range, and click **Save**.

Step 4 To edit the range, select the Range From check box, and click the **Edit** icon as required .
All ranges are editable. Overlapping of range is allowed if Restricted Range field is Yes.

Step 5 To delete the range, select the Range From check box, and click the **Delete (X)** icon.

Specifying VLAN Range

VLAN ranges needs to be created for all the leafs and DCIs. You can create device specific VLAN range. You can also group devices together, and create a VLAN range for the group.

For Cisco Nexus 7000 devices, you can specify VLAN range per interface. You can also group interfaces together and specify ranges.

When you add Nexus 7000 devices to inventory, Cisco VTS checks whether these devices are in vPC and the compute links attached to these vPC devices are dual homed. If the devices are in vPC, VTC automatically creates a device group containing these two devices. This device group would have BD range associated with it, which VTS uses to provision overly networks. The default BD range is 1000 to 2000. you can configure this value.

- For computes attached to Cisco Nexus 7000 switch as single homed, Cisco VTS automatically creates a default device interface pool per interface attached to the compute. This default interface pool is of name vts-Device-InterfaceName with a default VLAN range of 2-4094.
- If a compute attached to the Cisco Nexus 7000 is dual homed, Cisco VTS automatically creates a default device interface group pool containing the dual homed interfaces of two switches. This default interface group pool is of name vts-group-<number> with a default VLAN range of 0-4096.
- For single homed and non vPC interfaces, Cisco VTS creates a default per interface level VLAN range of 0-4096.

- If computes are attached to FEX, the interface pool with default range is created for the FEX device. This FEX VLAN pool is of name vts-device-<fexId> with default range of 0-4096.
- For computes attached in vPC to two different FEX, Cisco VTS automatically creates a group for the two interfaces going to two different FEXs. Overlay network provisioning on this vPC compute uses a common VLAN from the two FEXs ranges.

  Cisco VTS only allows grouping of two FEXs to form a logical group. It does not allow a FEX from one logical group to form a grouping with a FEX from a different logical group. For example, if a host compute1 is connected in vPC to two FEXs 101 and 102, these two FEXs will form a logical group. Cisco VTS does not allow having a host compute2 connected in vPC to FEX 102 and FEX 103. This is because the same VLAN across multiple FEXs for a given network would be difficult to maintain.
- The default VLAN pool of a device gets deleted when the device is added to a Device Group. This is because it will be using the default VLAN pool of the Device Group once it is part of the group. However, when the device is removed from the Device Group, Cisco VTS does not add back the original VLAN pool to it. The Cisco VTS admin has to add the VLAN pool back to the device, manually.

Note We recommend that you check the supported VLAN range for the device that is created automatically, and also take a note of the reserved VLAN range. Every device has its own limitation. You need to ensure that you are not using a reserved VLAN range for your particular device.

See the following sections for details:

- Specifying Device VLAN Range, on page 39
- Specifying Group VLAN Range, on page 40
- Specifying Interface VLAN Range, on page 41
- Creating Interface Groups, on page 41

Specifying Device VLAN Range

To specify device VLAN pool:

Step 1 Go to **Resources** > **Devices**. The Resource / Devices window appears. It lists all the device VLAN ranges.

Step 2 Click **Devices**.

Step 3 Click the **Add (+)** icon. The Add Range pop up window appears.

Step 4 Enter the Device details, and specify the From and To values.
The device name should match the leaf name in the inventory. From is VLAN start number and To is VLAN end number to be used for the leaf.

Step 5 Select the **Restricted Range** radio button to enable or disable the range, and click **Save**.
To edit a device specific VLAN pool, select the Device check box, and click the **Edit** icon.

To delete a device specific VLAN pool, select the Device check box, and click the **Delete (X)** icon.

Specifying Group VLAN Range

You can group devices and assign VLAN range for the device group. To specify VLAN range for a device group:

Step 1 Go to **Resources** > **Devices**. The Resource / Devices window appears.
Step 2 Click **Groups**.
Step 3 Click the **Add (+)** icon. The Set Range pop up window appears.
Step 4 Enter the Group Name and select the devices that need to be part of the group. Click the help icon for guidelines about the group name.
Step 5 Select the **Restricted Range** radio button to enable or disable the range, and Click **Save**.
Step 6 To view the devices associated with a group, select the group and click **Associated Devices**.
Step 7 Click **Save**. The group gets created and is listed in the table.
Step 8 To add range to the group, select the group and click the **Add (+)** icon.
Step 9 Specify the From and To values.
Step 10 Click **Save**.
Step 11 To edit a device specific VLAN pool, select the Device check box, and click the **Edit** icon.
All ranges are editable.
Step 12 To delete a group specific VLAN pool, select the Device check box, and click the **Delete (X)** icon.

Specifying Multicast IP Pool

You can specify the number of overlay networks that can be mapped to a single multicast address. Choose Enter VNI(NetworkCount) from the drop down, and enter the number of networks you want to map to a single multicast address. You can also opt to have all networks to map to a single multicast IP. To do this choose All Network from the drop down.

You can specify the IP range. The range must be within the multicast IP address range configured on leaf devices via day zero configuration file. The valid range is from 239.0.0.0 to 239.255.255.255.

Step 1 Go to **Resources** > **Multicast IP Pool**. The Resources / Multicast IP Pool window appears.
Step 2 Click the **Add (+)** icon, and enter the Start and End values.
Click the **Question Mark (?)** icon to view the Multicast IP address range.

Use the Restricted Range toggle button to restrict or disallow allocations from this range.

Step 3 Click **Save**.
To delete Multicast IP Pool, select the desired check box, and click the **Delete (X)** icon.

Specifying Interface VLAN Range

To specify VLAN range for an interface:

Step 1 Go to **Resource** > **Devices**. The Resource / Devices window appears.

Step 2 Click **Interfaces**. It lists all the Cisco Nexus 7000 devices.

Step 3 Click on the corresponding chassis icon.
The Interfaces pop up window appears. You can view the Physical Interfaces and the FEX interfaces.

a) Click Physical Interfaces tab to view the physical interface.
The interfaces are displayed based on odd and even numbered interfaces, with the odd numbered interfaces on top and the even numbered interfaces at the bottom.

You can control the display using the filter options.

- Choose the desired option from the Module drop down to filter ports for a specific module.
- Enter port details in the search field to display a desired port.

.

b) Click Physical Interfaces tab to view the physical interface.
Reserved ports are grayed out. Editing is enabled for ports that are connected to computes.

Clicking on an available port shows the ranges.

c) To edit the VLAN range, click the Edit button.
You can use the Restricted range toggle button to restrict the range.

d) Click **Save**.

Step 4 Click the FEX Interfaces tab to view the details about the FEX modules. By default, the range for the first FEX module is shown.
You can control the display using the filter options.

- Choose the desired option from the FEX drop down to filter ports for a specific module.
- Enter port details in the search field to display a desired port.

.

a) To edit the VLAN range associated with the FEX module, click the Edit button.
You can use the Restricted range toggle button to restrict the range.

b) Click **Save**.

Creating Interface Groups

You can group interfaces and assign VLAN ranges for the group. You can create groups for physical interfaces or FEX interfaces and assign ranges.

Note You cannot group physical and FEX interfaces together.

Step 1 Go to **Resource** > **Devices**. The Resource / Devices window appears.

Step 2 Click **Interface Groups**. It lists all the interface groups for the Cisco Nexus 7000 devices.

Step 3 To add an interface group, click the Add **(+)** icon.
The Create Interface Group pop up window appears.

Step 4 Enter a group name.

Step 5 Choose a group type—Physical Interface Group or FEX Interface Group.
To create a physical interface group:

a) Click **Physical Interface Group** .
b) Click **Select Devices** to select the devices. The Select Devices popup window appears.
c) Click **Select Interfaces** to select the interfaces. The Select Interfaces popup window appears.
The interfaces display sorted based on odd and even numbered interfaces, with the odd numbered interfaces on top and the even numbered interfaces at the bottom.

You can control the display using the filter options.

- Choose the desired option from the Module drop down to filter ports for a specific module.
- Enter port details in the search field to display a desired port

.

Reserved ports are grayed out. Editing is enabled for ports that are connected to computes.

d) Click **Define Ranges** to define VLAN ranges. The Define Ranges and Group Details popup window appears.
You can use the Restricted range toggle button to restrict a range.[Also, addinfo about the Add button]

e) Click **Review and Save**. The Summary popup window displays the interface group range details you have modified. Click **Edit** if you need to modify any details. You can edit the interface ranges, the devices in the group, or edit the interfaces you have chosen for the device.
f) Click **Save**.
To delete an interface group, select the group and click Delete **(X)**.

To create a FEX interface group:

a) Click **FEX Interface Group** .
b) Click **Select Devices** to select the devices. The Select Devices popup window appears.
c) Click **Select Interfaces** to select the interfaces. The Select Interfaces popup window appears.
The interfaces display sorted based on odd and even numbered interfaces, with the odd numbered interfaces on top and the even numbered interfaces at the bottom.

You can control the display using the filter options.

- Choose the desired option from the Module drop down to filter ports for a specific module.
- Enter port details in the search field to display a desired port

.

Reserved ports are grayed out. Editing is enabled for ports that are connected to computes.

d) Click **Define Ranges** to define VLAN ranges. The Define Ranges and Group Details popup window appears. You can use the Restricted range toggle button to restrict a range. You can add ranges using the Add (+) button.

e) Click **Review and Save**. The Summary popup window displays the interface group range details you have modified. Click **Edit** if you need to modify any details. You can edit the interface ranges, the devices in the group, or edit the FEX interfaces you have chosen for the device.

f) Click **Save**.
To delete an interface group, select the group and click Delete **(X)**.

Auto Select/Auto Delete functionality
The auto select/auto delete functionality gets triggered on the devices that have port channel config on them from devices Day 0 config.

Only auto delete functionality gets triggered on the devices that have Static Multi Homed (SMH) group attached to them from VTC UI (Host Inventory page).

Note Auto select functionality is not applicable for devices with only SMH group.

When you select one of the peer ports/devices that is part of a system defined group (port channel in this case), then the corresponding peer port/device also gets auto selected and gets added to the interface group (both Physical and FEX interface group).

If this is the only device in the interface group then the group cannot be saved. If there are other devices in this group, user should be able to save the group even after the system defined/SMH tagged group devices are deleted

When you deselect one of the peer ports/devices that is part of a system defined group (port channel in this case) or SMH group, then the corresponding peer port/device also gets auto de-selected/deleted and gets deleted from the interface group (both Physical and FEX interface group).
If this is the only device in the interface group then the group cannot be saved. If there are other devices in this group, user should be able to save the group even after the system defined/SMH tagged group devices are deleted.

Important Note

For versions earlier than 2.3.1, Cisco VTS restricted the use of static allocation to within a range, for all resources, and all attempts to allocate outside the range returned an exception. Currently, by default, ranges are not required for static allocation. A static allocation may be done both inside a range and outside it.

If you wish to enable the restriction, this can be done using REST API. See the *Cisco VTS 2.5 Developer Guide* for details

Important Note

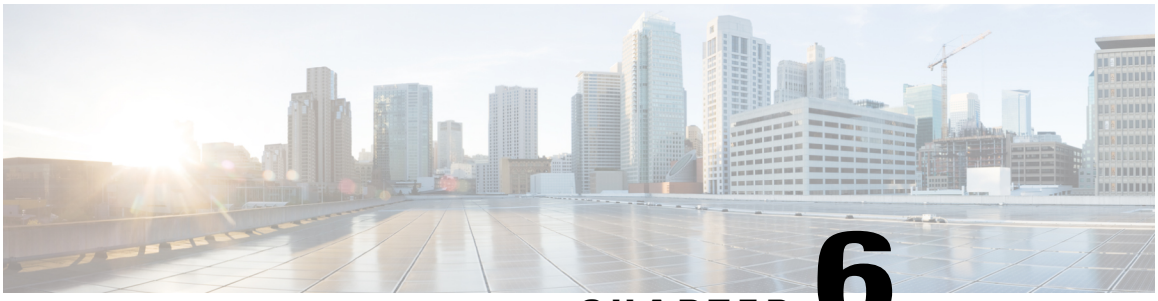

CHAPTER 6

Creating and Managing Admin Domains

This chapter has the following sections:

- Admin Domain Overview, page 45
- Viewing Admin Domain, page 46
- Creating an Admin Domain, page 46
- Creating DCI Interconnect Profiles, page 50

Admin Domain Overview

The Admin Domain feature enables you to partition the data center and define data center pods to group hardware and software VTEPs, Layer 3 gateways, and DCI gateways into administrative domains with similar properties. Admin Domains are independent of each other. You can create an admin domain, and specify certain functional roles within the admin domain. Admin domains are logical groups you create, based on the functional roles, which makes centralized L3 or Distributed L2/L3 deployments flexible and extendable.

Cisco VTS provides the functional roles, which you can use as desired to create the admin domains. You can set the system mode, control protocols, other parameters like replication mode (multicast/ingress), for each admin domain , and also assign devices to each of the functional roles. For example, you can pick certain leafs and put it in one group, and associate certain functional parameters to that group. The following functional roles are available:

- L2 Gateway
- L3 Gateway
- DC Gateway
- DCI

For the L2 Gateway group you can pick the desired leafs and associate certain functional parameters to that group. Similarly, you can define another L3 gateway group, and you can link between these two groups. All L2 configuration can be pushed into the L2 gateway group; and all L3 configuration can be pushed into L3 gateway group.

You can create an L3 gateway group and can link from the L3 group to the DC gateway. You can have the DCI at the top, and this can be linked to the DC gateway.

The DC gateway can be outside the Admin Domain, and more than one Admin Domains may connect to this. You can have the DC gateway inside an Admin Domain, and connect it to an external DCI.

See for detailed information about creating Admin Domains.

The design validated in this release has:

- L2/L3 gateway groups in all Admin Domain-Each Admin Domain can have its own L2 / L3 gateway.
- DC Gateway outside the Admin Domain
- DCI outside the Admin Domain.

Viewing Admin Domain

The **Admin Domains** home page lists all the Admin Domains that you have created. It provides the option to create a new Admin Domain.
It also displays the status of the Admin Domains. You can also edit an Admin Domain.

To view admin domains:

Step 1 Go to **Admin Domains** > **Domains**.
The Admin Domains / Domains window appears.
You can see two types of views on the Admin Domain page. The two types of views are as follows:

- List view
- Tree view

Step 2 To view the details of an Admin Domain, click the desired admin domain.
You can create an Admin Domain from the table. To do this, click the **Add (+)** icon in the table, and provide the required details. You can also edit or delete an Admin Domain.

Creating an Admin Domain

To create an admin domain:

Before You Begin

Ensure that you have:

- Created authorization groups populated with the correct credentials.
- Discovered the topology and imported the CSV file (after assigning / reviewing device roles). See Performing Auto Discovery, on page 20 and Managing Inventory, on page 19 sections for details.

- Reviewed the Supported Platforms section in the *Cisco Virtual Topology System Installation Guide,* which provide information about the platforms that Cisco VTS support, and their roles.

Step 1 Go to **Admin Domains** > **Domains**.
The Admin Domains / Admins window appears.

Step 2 Click **+ Create**.
The Create New Admin Domain popup window appears.

Step 3 Enter the name and description in the **Create New Admin Domain** popup window.

Step 4 Click **Create**.
The Admin Domain canvas appears.
You can see the following functional groups on the left-hand side of the canvas:

	Functional Group	Description
1	DCI	DCI is an external gateway.
2	DC GW	DC GW is a border leaf. **Note** If it is a DCI mode, then you need to add DCI device to both the DC GW and DCI. In an integrated mode, we need to add DCI to both DC GW functional group and DCI functional group.
3	L3 GW	A group of all L3 devices that can be within an admin domain and that particular device share a particular property or same functionalities. **Note** An admin can create a logical L3 groups and map devices that will exhibit a similar policy behavior under this group.
4	L2 GW	A group of all L2 devices that can be within an admin domain and that particular device share a particular property or same functionalities. **Note** An admin can create a logical L2 groups and map devices that will exhibit a similar policy behavior under this group.

Step 5 Click the functional group. The functional group icon appears on the canvas. You need to drag and drop the functional group and assign properties to them.

Note We recommend that you do not enable ARP suppression for Cisco Nexus 9000 devices. This is due to a limitation in Cisco Nexus 9000. Further, Cisco Nexus 9000 devices, currently, do not support ARP suppression in L2.

Functional Group	Property

DCI	Specify:
	• Whether it is a New or Shared DCI.
	• The Redundancy / Availability settings:
	  • Enable/Disable Redundancy using the toggle switch.
	  • ICCP—VXLAN/fabric ICCP group number. Valid range is 1 to 4294967295. MPLS/core ICCP group number. Valid range is 1 to 4294967295
	  • ESI—Ethernet Segment ID for NVE overlay. Valid entry is a nine octet string. Each octet can contain one or two numbers in the range 0 to F.
	
	Click Stitching Profile and choose the required profile.

DC GW	Specify: • Whether it is a New or Shared DC GW. • The Control Protocol—BGP EVPN.

Functional Group	Property
L3 GW	Specify: • Whether it is a New and Shared L3 GW. • The Control Protocol—BGP EVPN. • The Replication Mode—Multicast or Ingress. This is the data plane replication mode that will be used for VXLAN data plane traffic. The admin domain can contain devices that support common replication mode. **Note** • Cisco Nexus 5600 and Cisco Nexus 7000 supports Multicast replication mode only. • VTF supports Ingress mode only. • Cisco Nexus 9000 supports both modes. • Distribution Mode—Decentralized. **Note** L3 GW group is created as Decentralized when the L2/L3 VXLAN are terminated on the same leaf. Therefore, if you have multiple L2 VXLAN and you want to connect them together using an L3 VXLAN, you need to create a decentralized L3 GW group and add all the L2GW group devices to this L3GW group, and connect the L2 GW and L3 GW group together. An L3 GW group can be created as a Decentralized Gateway group when the L3 GW groups are distributed between multiple L2 GW group within an Admin Domain. • ARP Suppression—On.
L2 GW	Specify: • Whether it is a New and Shared L2 GW. • The Control Protocol—BGP EVPN. • The Replication Mode—Multicast or Ingress. • The Distribution Mode—Decentralized. • ARP Suppression—On. If you have ESI configured on the leaf, we recommend that you disable ARP suppression.

Step 6 Assign Devices for each functional group.

Note If you had created a device group (under **Resource Pools > VLAN Pool**), the device group information does not get displayed in the device list for DCI and DC GW functional groups, while you create an admin domain. However, the device group gets displayed in the device list for L3 GW and L2 GW functional group.

Click the drop down icon on the right hand side to see how many devices are placed in this group or how many devices are available to be placed in this group. The **All** option shows both placed devices and available devices.

For more information about supported devices, see the Supported Platforms section in the *Cisco Virtual Topology System Installation Guide*.

Step 7 Link the functional groups based on your requirement. You can click a functional group and drag the mouse pointer to the functional group you want to connect to, to form a link.

Note For L2VNI, you can extend the connection from L2 gateway to DC gateway by connecting them. To remove the link, click on the link that needs to be removed and click on Remove Link in the pop up box. Click Yes in the confirmation box to remove the link. See Extending Layer 2 Network Across Data Centers, on page 77 for details.

Step 8 Click **Save** to save the new Admin Domain with all the nodes, properties, and links.
Click **Cancel** icon if you want to go back to the main menu.

Creating DCI Interconnect Profiles

The DCI Interconnect Profiles page lets you create DCI interconnect profiles. These profiles enables services like route leaking to internet, and L2 VNI extension.

To create a DCI Interconnect Profile:

Step 1 Go to **Admin Domains** > **DCI Interconnect Profiles**. The DCI Interconnect Profiles page appears.

Step 2 Click the **Add (+)** button. In the Create Profile page, enter the DCI Interconnect Profile properties:

- Name—The profile name. This is mandatory.
- Description
- Control Plane Protocol—Specify the control plane protocol. It is BGP by default.

Step 3 Choose the interconnect type. You may choose one or both of the following interconnect types:

- Internet—IPv4 unicast and IPv6 unicast address families are added.
- MPLS L2 VPN—L2VPN EVPN address family are added.

Enter the following for the interconnect type you choose. This is optional:

- Fabric Facing Route Policy Route Map—Route filter to apply for fabric facing routes. Maximum length is 64 characters.

- Core Facing Route Policy Route Map—Route filter to apply for core facing routes. Maximum length is 64 characters.

Step 4 Click **Remote Neighbors Settings**, and enter the following:

- AS Number—Enter a natural number between 1 and 65000.
- Loopback Interface Number—Loopback interface which connect to the remote neighbor. Enter an integer. Range is 0 to 2147483647.

Step 5 Click the **Add (+)** button to add remote neighbors.
You may add one or more remote neighbors. Use the **Add (+)** button to add more remote neighbors. The IP address can be IPv4 or IPv6.

Step 6 Click **Save Profile**.

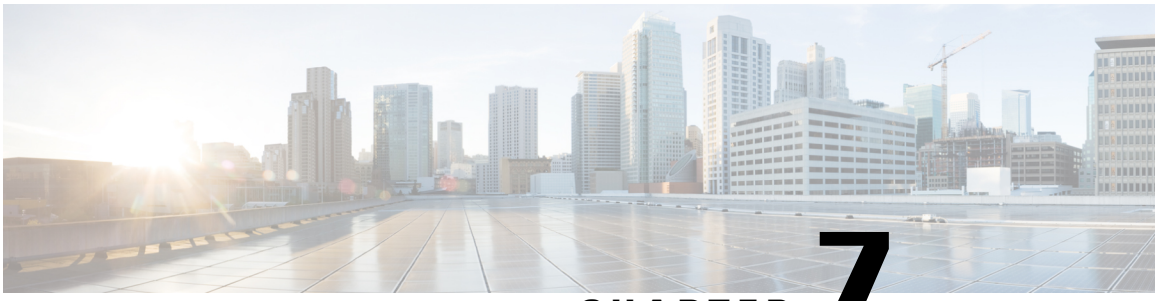

CHAPTER 7

Managing Templates

A template is a container of configurations, which can be applied to a target such as a device or a router.

Cisco VTS supports the following template types:

- **Overlay Templates—**The following types of overlay templates are supported.
  - Route templates—A route template is a template that lets you configure static routes and route targets. This template can be applied to a tenant or a router. When it is applied to a tenant, all the routers within that tenant inherit the static route and route target configuration. When it is applied to a router, it overrides the route template that was inherited from the tenant. It is supported in a set up that has only Cisco ASR 9000 Series Aggregation Services Routers as DCI.

    Only integrated DCI mode is supported. VRF-peering mode is not supported.
  - L3 Service Extension templates—An L3 Service Extension template allows you to extend Cisco VTS Layer 3 service configuration on routers or tenants.

    The configuration you define in the service extension template, along with the out-of-the-box Cisco VTS L3 configuration, will be pushed to the device to get the combined configuration on the device. Service extension templates do not allow you to configure any parameter that Cisco VTS configures out-of-the-box.

    One device can have multiple templates. One template can be attached to multiple devices. The admin has to ensure that the templates do not have conflicting configuration.

    Currently, L3 Service Extension templates are supported for the following platforms:

    - Cisco Nexus 5000 series
    - Cisco Nexus 9000 series
    - Cisco Nexus 7000 series
    - Cisco ASR 9000 series

    You can modify any L3 service related configuration that is pushed on Cisco ASR 9000 series devices in integrated mode (DC gateway and DCI is the same physical box) or VRF peering mode. In VRF peering mode, Cisco Nexus 9000 series device has to be configured as te border leaf. You must have an external network, and the External network should be set as the Router gateway. If you do not have the external gateway set, the template will be attached to the router, but configuration will not be pushed. After you have your external network as the router gateway, it will push the configuration.

- L2 Service Extension templates—An L2 Service Extension template allows you to extend Cisco VTS Layer 2 service configuration.
  The creation of L2 SE templates is done by authoring configuration that are specific to a device type. You can modify the L2 configuration that is attached to the network. The configuration is for the L2 construct which can be applied to the virtual interface. Following L2 service configuration templates are supported:
  - L2 QoS—Supported on Cisco Nexus 9000 series. See Important Notes—L2 QoS Template, on page 64.
  - VPLS—Supported on Cisco ASR 9000 series. The L2 Extension templates for Cisco ASR 9000 series devices include the VPLS configuration under the L2VPN and EVPN containers. See Important Notes—VPLS Template, on page 64.

- **Underlay template**—Underlay templates enables you to configure the Day zero configuration on underlay devices via the VTS UI.

Note If devices templates overwrite any VTS service configurations, these configurations would stay even after the device templates are detached/removed. To reinstate the service configurations, you will need to redeploy the services.

The overlay template can be used with multiple routers or tenants You can either associate the template while you create the router or tenant, or associate a template to a tenant or router you have already created, while you edit the tenant or router.

When you detach the overlay template from the tenant, it cleans up the configuration on the device.

Note For route templates, you can have only one instance of a template type per tenant/router. For example, let the template types be Temp A and Temp B, and the instance of Temp A be Ins A and that of Temp B be Ins B. Now, Ins A and Ins B can be applied to a tenant/router. However, Ins A and Ins A', where Ins A' is a second instance of Temp A, cannot be applied at the same time on the tenant/router.

The following sections provide more details about working with templates:

- Creating Route Templates, page 55
- Creating L3 Extension Templates, page 58
- Editing Templates, page 59
- Copying a Template, page 60
- Deleting Templates, page 60
- Attaching Templates to Tenants, page 60
- Attaching Templates to Routers, page 62
- Creating L2 Extension Templates, page 63
- Creating Underlay Templates, page 65

Creating Route Templates

To create templates:

Step 1 Go to **Templates > Overlay Template Management**. The Templates / Overlay Template Management page appears.

Step 2 Click +. The Create Template page appears.

Step 3 Enter a name for the template in the Template Name field. Only alphabets, numbers, and special characters **.**, _ and **-** are allowed. The maximum character limit is 128.

Step 4 Enter a description for the template, in the Description field. This is optional.

Step 5 Select the template type from the Template Type. For route templates, select Route.

Step 6 Click **Add Configuration**. The New Template page appears.

Step 7 Enter a route target seed. This can be an integer value in the range of 1-16777215. Route Target with seed is pushed to DCI, and Leaf if eBGP is enabled.

Step 8 Enable or Disable the Auto Route Target option. By default, it is enabled. See Disabling Auto Route Target Configuration, on page 57 for details.

Step 9 Add Static Routes. See Adding Static Routes, on page 55 for details.

Step 10 Add Route Targets. See Adding Route Targets, on page 55 for details.

Step 11 Click **Save**. The template is saved and listed in the Template Management page.

Adding Static Routes

You can add static routes to a template while you create the template.

Step 1 In the Template Management page, while you create or modify a template, click the **Static Route** tab.

Step 2 Click +. The Add Static Route(s) pop-up appears.

Step 3 Add a Destination IP address and netmask.

Step 4 Add the Next Hop IP address and netmask. Netmask is optional.

Step 5 Click **Add**. The static routes get added to the template.
You can add five static routes at a time.

Adding Route Targets

You can add route targets to be imported / exported to the leafs or DCI.

Adding Fabric Internal Route Targets

To add Fabric Internal route targets:

Step 1 Click **Fabric Internal RT** tab.

Step 2 If you want to use system defined RT, choose Auto RT- System Defined from the dropdown. Else, choose Auto RT - Custom. You must add at least four custom route targets.

Step 3 Click +. The Add Route Target(s) pop-up appears.

Note You can add five route targets at a time.

Enter the route targets to be shared across the different VRFs. The valid route target formats are:

- ASN2:NN4
- ASN4:NN2
- IPv4:NN2

Where:

- NN2 and ASN2 has a range of 1-65535
- NN4 and ASN4 has a range of 1-4294967295
- IPv4 is an IPv4 address in the dotted decimal format

Step 4 Specify whether route targets are to be imported or exported. To do this, select the desired value from the Direction drop-down.

Step 5 Specify following from the **Type** drop-down:

- Internal— To import / export on leafs.
- Stitching— To import / export on DCI.

Step 6 Click Add. You can add five route target at once.

Adding Fabric External Route Targets

To add Fabric Eternal route targets to import /export route targets to the DCI:

Step 1 Click **Fabric External RT** tab.

Step 2 If you want to use system defined RT, choose Auto RT- System Defined from the dropdown. Else, choose Auto RT - Custom. You must add at least two custom route targets.

Step 3 Click +. The Add Route Target(s) pop-up appears.

Note You can add five route targets at a time.

Enter the route targets to be shared across the different VRFs. The valid route target formats are:

- ASN2:NN4
- ASN4:NN2
- IPv4:NN2

Where:

- NN2 and ASN2 has a range of 1-65535
- NN4 and ASN4 has a range of 1-4294967295
- IPv4 is an IPv4 address in the dotted decimal format

Step 4 Specify whether route targets are to be imported or exported. To do this, select the desired value from the Direction drop-down.
The type will be Fabric External, by default

Step 5 Click **Add**.

Disabling Auto Route Target Configuration

To enable or disable automatic route target configuration, use the Auto Route Target toggle switch in the New Template screen while you create route templates. By default, Auto Route Target (RT) is enabled. When this is enabled, Cisco VTS adds route target configurations automatically, in addition to any static / manual route targets you have defined in route template, while the template configuration is pushed to the VTEPS.
If you choose to disable Auto Route Target, ensure that:

1. At least one import route target and one export route target are defined for internal devices (that is, for leaf switches controlled by Cisco VTS) in the same route template where auto RT gets disabled.
2. When a DCI is present in Admin Domain, and you choose to disable Auto Route Target, ensure that:
   - At least one import route target and one export route target are defined for external devices.
   - At least one import route target and one export route target are defined for "both".

This is to make sure that when auto RT is disabled, the static route targets defined in the template (which will then be pushed to the DCI) are sufficient to enable the DCI to communicate with the TORs properly.

Whenever you enable Auto Route Target again the route targets created using the route target seed (if provided) or ASN# will be pushed to devices. In addition, the RT seed textbox will be enabled again.

Creating L3 Extension Templates

To create L3 Extension templates:

Step 1 Go to **Templates > Overlay Template Management**. The Templates / Overlay Template Management page appears.

Step 2 Click +. The Create Template page appears.

Step 3 Enter a name for the template in the Template Name field. Only alphabets, numbers, and special characters **.**, **_** and **-** are allowed. The maximum character limit is 128.

Step 4 Enter a description for the template, in the Description field. This is optional.

Step 5 Choose L3 Extension as the template type.

Step 6 Choose the Device Platform. Currently, the following platforms are supported:

- Cisco Nexus 9000 Series
- Cisco Nexus 7000 Series
- Cisco ASR 9000 Series

Step 7 Click Add Configuration to add configuration to the template. The Author Template window appears.

Step 8 Click Configuration icon to get the Add Configuration menu. The flyout menu displays all the configuration options that are available at the root level. You can search for the desired configuration in the Search field.

Note Configuration options available are limited to configuration that Cisco VTS does not provide out of the box. The User Interface(UI) is schema driven and shows the configuration tree based on the device platform selected and the service extension template type, for example, L3 Service Extension Template.

Step 9 Choose the desired configuration. The configuration you chose gets added as a child node in the Config tree, on the left pane.
You may add further configuration to the node that you have added by clicking the Configuration icon. If you want to delete the configuration, click the delete (X) icon.

For configuration items which can take multiple instances, the Add Instance button appears in the authoring pane. You can add an instance by clicking **Add Instance**. Click **Add** after you add the configuration for the instance.

Note For certain configurations, some of the options that are available for selection have the %v suffix. These are system variables. See Supported System Variables, on page 58 for details.

Step 10 Click **Save Template**.
The template gets added to the Template Management screen. You can click on the template to get a summary of the template, in the Template Summary page. You can expand the Config node to view the template configuration. You can edit the template from the Summary screen, by clicking **Edit Template Config**.

Supported System Variables

For certain configurations, some of the options that are available for selection have the %v suffix. The %v suffix denotes that it is a system variable. When the Cisco VTS comes across such a variable in the template, it translates that into the value that it had configured for that device.

Supported variables in Cisco Nexus 5000, Cisco Nexus 7000, and Cisco Nexus 9000 Series for L3 Service Extension

- BGP AS number
- VRF name
- NVE Interface
- Host side SVI
- Fabric SVI

Supported variables in Cisco Nexus 9000 Series for L2 Service Extension

- VLAN ID

Supported variables in Cisco ASR 9000 Series for L3 Service Extension

- VRF Name
- BGP AS Number
- Neighbor address
- Bridge Group
- Bridge Domain Name
- BVI
- NVE interface

Supported variables in Cisco ASR 9000 Series for L2 Service Extension

- Bridge Group
- Bridge Domain Name
- EVI
- NVE interface
- VNI

Editing Templates

You can modify a template that you have created.

Step 1 Go to **Templates > Template Management**. The Templates / Template Management page lists all the templates you have created.

Step 2 Select the check box corresponding to the template you need to edit, and click the edit icon.

Note You can also edit a template from the Template Summary screen. Click **Edit Template Config** button to edit a template.

Based on the template type, the edit template page is displayed. For a Route template, the edit page for that route template is displayed. For a L3 Service Extension templates, the Edit Template page is displayed. You may change the Template Description, if required. Then click **Edit Configuration** to modify the service extension configuration.

Step 3 Make the desired changes, then click Save.

Copying a Template

You can copy a template and save it with a different name. You can also modify the parameters while you copy.

Step 1 Go to **Templates > Template Management**. The Templates / Template Management page lists all the templates you have created.

Step 2 Select the check box corresponding to the template you need to copy, and click the copy icon.

Step 3 Modify the details if required, then click **Save**.

Deleting Templates

You can delete a template that you have created.

Note You can delete a template only if it is not attached to either a tenant or router. If it is attached to a tenant or router, an error is displayed when you try to delete. You need to detach the template from tenant or router before deleting the template.

Step 1 Go to **Templates > Template Management**. The Templates / Template Management page lists all the templates you have created.

Step 2 Select the check box corresponding to the template you need to delete, and click the **X**.
The Delete pop-up appears.

Step 3 Click **Delete** to delete the template.

Attaching Templates to Tenants

You can attach template to tenants when you create or edit tenants.

• Attaching Template while Adding Tenants, on page 61
• Attaching Template while Editing a Tenant, on page 61

Attaching Template while Adding Tenants

You can attach templates to a Tenant while you add a tenant.

Step 1 When you are adding a new tenant, click the find icon in the Template field.
The Select Template pop-up appears.

Step 2 Select the template type, then select the desired template. If you want to create a new template and attach that, click Create New Template.

Step 3 Click **Select Devices(s)** to select the devices.

Step 4 Click **Apply Template(s)**.

Step 5 Click **Save**.

Attaching Template while Editing a Tenant

You can also modify an existing tenant to attach a template.

Step 1 Go to **Tenants > Tenant Management**. The Tenants / Tenant Management page lists all the tenants you have created.

Step 2 Select the check box corresponding to the tenant you need to modify, then click the edit icon.

Step 3 Click the find icon in the Template field. The Select Template pop-up appears.

Step 4 Select the template type from the Template Type drop down list.

Note If you want to create a new template and attach that, click **Create New Template**. See Creating L3 Extension Templates, on page 58 for details about how to create a new template.

Step 5 Select the desired template. Then click **Select Device(s)**, to select the devices.

Step 6 From the Select Device(s) popup window, select the devices, then click **Apply Template(s).**.

Step 7 Click **Save** in the Edit Tenant page.

Detaching Template from Tenants

You can detach a template that you had attached to a Tenant. When you detach a template from a tenant, all the template configuration is removed from the routers that inherited the template.

To detach the template:

Step 1 Go to **Tenants > Tenant Management**. The Tenants / Tenant Management page lists all tenants.
Step 2 Select the tenant from which you want to detach the template, then click the Edit icon.
Step 3 Click the Find icon in the Templates field. The Select Template popup is displayed.
Step 4 Check the check box.
Step 5 Click **Save**.

Attaching Templates to Routers

You can attach templates to routers while adding a router or while modifying an existing router. By default, router inherits the template from the tenant. You can override an inherited template while you create or modify a router, by selecting a different template, or creating a new template and attaching to the router.

- Attaching Templates while Adding Routers, on page 62
- Attaching Template while Editing a Router , on page 63

Note The configuration is pushed to the device only when Port Attach is done. If a router is created and template is applied to the tenant or router, but port is not attached, then the template configuration is not pushed. See Creating a Network using Cisco VTS GUI, on page 73 for details about attaching port.

Attaching Templates while Adding Routers

While you add the router, by default, the router inherits the template from the tenant. The template field displays <template name> (Inherited). You may inherit the template from the tenant or may chose to select a template of your choice.

Step 1 Click the Find icon in the Template selection box. The Select Template popup is displayed.
Step 2 Uncheck the Inherit from Tenant check box. The popup lists all the available templates.
Step 3 Select the desired template, then click **Apply**. The template gets applied to the router.
If you want to create a new template, click **Create New Template**. You will be redirected to the **Templates > Template Management > New Template** page. After you create the new template, you must select the new template from the Select Template pop up, as described in Step 2 and 3.
Step 4 Click **Save** after you complete the router creation procedure.

Attaching Template while Editing a Router

You can attach a template to a router while you edit a router.

Step 1 Go to **Overlay > Router**. The Overlay / Router page lists all the routers.

Step 2 Select the Router that you need to modify, then click the edit icon. The template field lists the template that the router has inherited from the tenant.

Step 3 Click the Find icon in the Template selection box. The Select Template popup is displayed.

Step 4 Uncheck the Inherit from Tenant check box. The popup lists all the available templates.

Step 5 Select the template type from the Template Type drop down list.

Step 6 Select the desired template, then click **Select Device(s)**
If you want to create a new template, click **Create New Template**. You will be redirected to the **Templates > Template Management > New Template** page. After you create the new template, you must select the new template from the Select Template pop up, as described in Step 2 and 3.

Step 7 From the Select Devices popup window, select the device(s) on which you want to apply the template, click **Apply Template(s).**.

Note To detach a template from a device, in the same flow, you may uncheck the checkbox for the device from which you want to detach the template.

Step 8 Click **Save**.

Creating L2 Extension Templates

To create L2 Extension templates:

Step 1 Go to **Templates > Overlay Template Management**. The Templates / Overlay Template Management page appears.

Step 2 Click +. The Create Template page appears.

Step 3 Enter a name for the template in the Template Name field. Only alphabets, numbers, and special characters**.**, **_** and **-** are allowed. The maximum character limit is 128.

Step 4 Enter a description for the template, in the Description field. This is optional.

Step 5 Choose L2 Extension as the template type.

Step 6 Choose the Device Platform.

Step 7 Click Add Configuration to add configuration to the template. The Author Template window appears.

Step 8 Click Configuration icon to get the Add Configuration menu. The flyout menu displays all the configuration options that are available at the root level. You can search for the desired configuration in the Search field or configuration items which can take multiple instances, the Add Configurations pane appears in the authoring pane.

Note Configuration options available are limited to configuration that Cisco VTS does not provide out of the box. The User Interface(UI) is schema driven and shows the configuration tree based on the device platform selected and the service extension template type, for example, L2 Extension Template. For this release, only configurations under vlan > configuration and vlan > vlan-list are qualified.

Step 9 Choose the desired configuration. The configuration you chose gets added as a child node in the Config tree, on the left pane.

Step 10 Click **Add Instance** to add an instance. New Instance gets added on the left pane.
For configuration items which can take multiple instances, the Add Instance button appears in the authoring pane. Click **Add** after you add the configuration for the instance.

Note For certain configurations, some of the options that are available for selection have the %v suffix (for example, %vVLAN_ID). The %v suffix denotes that it is a system variable. Here, the vlan-id gets set-up once you associate this particular template at a network level.

Step 11 Click **Save Template**, after you finish adding the desired configuration.

Important Notes—L2 QoS Template

Review the section below before you create QoS template:

- Policy map has to be created on the devices as part of Day Zero configuration or configured using underlay templates. Only those devices with the service policy map configured will be displayed in the template association UI.
- Under the vlan configuration Cisco Nexus 9000 supports only service policy with type qos. Both input and output service policies may be configured in the same template.
- While applying L2 QoS templates, only those devices with the service policy map configured will be shown. The policy map has to be configured on the device for the device to be displayed here.

Important Notes—VPLS Template

Review the section below before you create a VPLS template:

- Ensure that the necessary Day Zero configuration is complete on the DCI. See *Cisco VTS Day Zero Configuration Examples* document for details.
- Create the Admin Domain. Make sure you link the L2GW and DCGW, which is required for L2VNI. See Creating an Admin Domain, on page 46 for details.
- Create the DCI Interconnect profile with the specific details. See Creating DCI Interconnect Profiles, on page 50 for details.
- Associate the DCI Interconnect profile to the DCI in the Admin Domain under DCI Interconnect profile tab.
- Create the L2 Extended Network with an EVI ID assigned to the Network under L2VPN tab. See Creating a Network using Cisco VTS GUI, on page 73 for details.
- Make sure the subnets and port attach is done for the Network.
- You can create a template with just PW (pseudo wire) configuration, or just Access VFI configuration, or can have both configurations in one template. Make sure that you enter correct PW and Acess VFI under l2vpn, in the DC that is not managed by VTS.

- The following are supported as system variables in VPLS templates:
  - Bridge Domain Name
  - Bridge Group Name
  - Interface NVE ID
  - EVI ID

Associating L2 Extension Templates

You can attach an L2 Extension template while you create a network, or to an existing network. To attach an L2 Extension template:

Step 1 Click **Overlay > Network**. The Overlay > Network page is displayed.

Step 2 Select the network for which you want to attach the L2 extension template, then click **Edit** icon.

Step 3 Click the **Template** selection search box. The Select Templates pop up displays a list of all L2 extension templates you have created.

Step 4 Select the desired template(s). Click **Next**. You may also create a new template, using the **Create New Template** button. The Select Devices popup appear. You may choose to apply the selected templates to all of the devices, or choose specific devices on which you may want to apply the template.

Note For an L2 QoS template, if Service Policy is defined in the template, then only the devices that have the same policy map configured will be displayed for selection. If no Service Policy is provided in the L2 QoS template, then all the devices that have the matching platform type will be available for template application.

Step 5 Select the device(s), then click **Apply Template(s)**.

Step 6 Perform a port attach to get the template pushed to the device.

Creating Underlay Templates

To create Underlay templates:

Step 1 Go to **Templates > Device Template Management**.
The Templates / Device Template Management page appears.

Step 2 Click +. The Create Template page appears.

Step 3 Enter a name for the template in the Template Name field. Only alphabets, numbers, and special characters**.**, **_** and **-** are allowed. The template name requires at least one alphabet or number. The maximum character limit is 128.

Step 4 Enter a description for the template, in the Description field. This is optional.

Step 5 Choose the **Device Platform**. Currently, the following platforms are supported:

- Cisco ASR 9000 Series

- Cisco Nexus 7000 Series
- Cisco Nexus 9000 Series

Step 6 Click **Add Configuration** to add configuration to the template.
The Author Template window appears.

Step 7 Click Configuration icon to get the Add Configuration menu. The flyout menu displays all the configuration options that are available at the root level. You can search for the desired configuration in the Search field.

Step 8 Choose the desired configuration. The configuration you chose gets added as a child node in the Config tree, on the left pane.
You may add further configuration to the node that you have added by clicking the Configuration icon. If you want to delete the configuration, click the delete (X) icon.

For configuration items which can take multiple instances, the Add Instance button appears in the authoring pane. You can add an instance by clicking **Add Instance**. Click **Add** after you add the configuration for the instance.

Note

Step 9 Click **Save Template**.
The template gets added to the Template Management screen. You can click on the template to get a summary of the template, in the Template Summary page. You can expand the Config node to view the template configuration. You can edit the template from the Summary screen, by clicking **Edit** icon.

Associating Devices to Underlay Template

To attach or detach an Underlay template from a device:

Step 1 Go to **Inventory > Network Inventory**.
The Inventory / Network Inventory page appears.

Step 2 Click **Device Template**.

Note You can filter templates by Device Platforms.

Step 3 Click the **Edit** adjacent to template name you want to edit.
The Associate Devices for <*template*> page appears where you can see the devices attached or detached from the underlay template.

Step 4 Check or uncheck **Devices(s)** to attach or detach the devices as per your requirement.

Step 5 Click **Done**.
The devices are successfully added or removed from the underlay template.

Note You can filter devices by associated and non-associated devices.

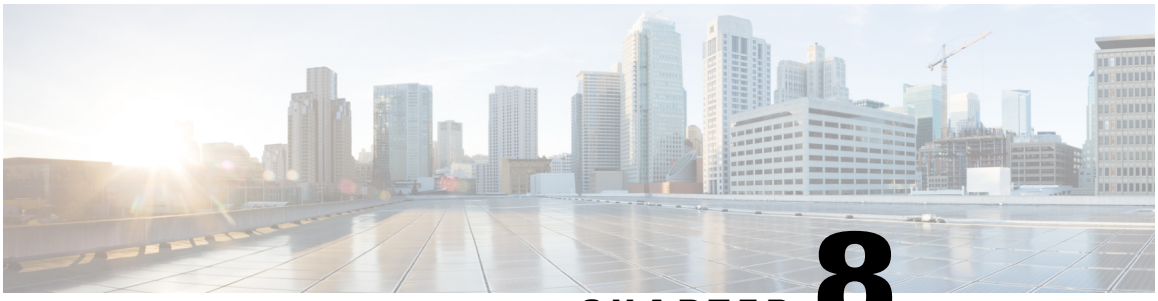

CHAPTER 8

Managing Tenants

The Tenant Management page displays a list of all tenants you have created. You can add, modify or delete a tenant. You can also attach templates to tenants.

This chapter has the following sections:

- Viewing Tenant Details, page 67
- Adding Tenants, page 68
- Editing Tenants, page 68

Viewing Tenant Details

Go to **Tenants**> **Tenant Management**. The Tenant Management page lists all the available tenants.
By default, the tenants under VTS are displayed. You can choose individual VMMs from the drop down to display the tenants under these.

The page displays the following:

- Name
- Description
- Zones
- Attached Templates
- Multi VMM Operations

Adding Tenants

To add tenants:

Step 1 Go to **Tenants** > **Tenant Management**. The Tenant Management page appears.
Step 2 Click **Add (+)** icon.
Step 3 Enter the following:

- Tenant Name
- Description

Step 4 If you want to attach template to the tenant, click the find icon in the Template field. See Attaching Templates to Tenants, on page 60 for details about attaching templates.
Step 5 Click **Save**.
Step 6 To add a Zone, click the **Add (+)** icon. The Add Zone popup appears.
Step 7 Enter the zone name, and click **OK**.
Step 8 Click **Save**.

Editing Tenants

To edit a tenant:

Step 1 Go to **Tenants** > **Tenant Management** . The Tenant Management page appears.
Step 2 Select the tenant, then click the edit icon.
Step 3 Modify the following:

- Tenant Name
- Description

Step 4 If you want to associate a new template or change the existing template, click the find icon in the Template field. See Attaching Templates to Tenants, on page 60 for details about associating templates.
Step 5 Click **Save**.
Step 6 To enable or disable network extension, use the **Extend all networks** toggle switch. By default, Extend all networks is **Yes**.
Step 7 Modify the zone details.
Step 8 Click **Save**.

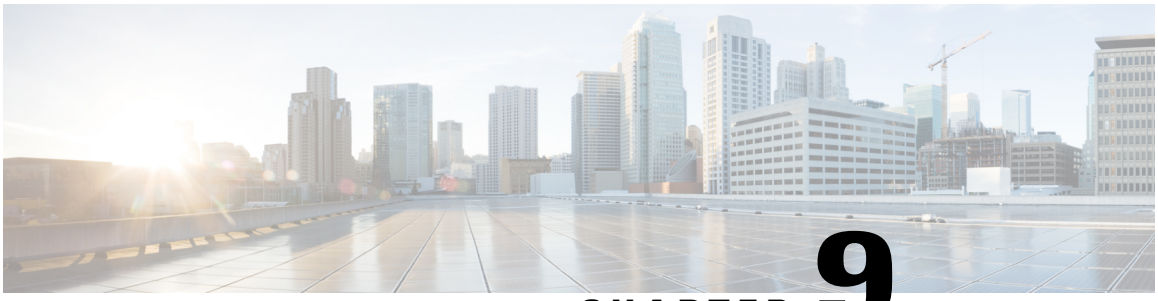

CHAPTER 9

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, page 69
- Creating Overlays, page 71
- Creating Network using VMware, page 72
- Creating Subnetwork using VMware, page 72
- Creating Routers using VMware, page 72
- Attaching Network to Router, page 73
- Attaching a Virtual Machine to Network, page 73
- Creating a Network using Cisco VTS GUI, page 73
- Creating Router using Cisco VTS GUI, page 75
- Assigning BVI Interface IP Address, page 77
- Extending Layer 2 Network Across Data Centers, page 77
- Enabling Global Route Leaking Service, page 78
- Enabling L3VPN to EVPN Route Stitching, page 80

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 multidestination 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.

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 73
- Creating Network using VMware, on page 72
- Creating Subnetwork using VMware, on page 72
- Creating Routers using VMware, on page 72
- Attaching a Virtual Machine to Network, on page 73

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 73
- Creating Router using Cisco VTS GUI, on page 75

Creating Network using VMware

To create a network:

Step 1 Select one of the vDS switches you created, then select 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 pop up 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:

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

Step 2 Select the Tenant for which you to create the network.

Step 3 Click **Add (+)** icon.

Step 4 Enter the network name. This is mandatory.

Step 5 Specify whether you want this to be an external network using the External Network toggle switch.

Step 6 From the **Extend Network** drop down, select L2 or L3, depending on your requirement. If you select L2, 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 between1 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 77. To extend an L3 network, choose L3 from the drop-down.

Step 7 Select the Zone.

Step 8 If the network is not external, enter the Static VNI. This can be an integer between 4096 and 65535.

Step 9 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).

Attaching a Port

To do a port attach:

Step 1 Go to **Overlay** > **Network**

Step 2 Click **Port Attach**, then click the add icon. The Attach Port popup appears.

Step 3 Specify whether it is a Physical Device or a Virtual Device, by selecting the appropriate radio button.
For Virtual Devices, the MAC address is required. For physical devices, it is optional.

You can use Static VLAN button if you want to specify a certain VLAN at the device level to be used for port attach. By default, VTS allocates a free VLAN from its pool of VLANs, but you may specify one if you wish to.

Step 4 Enable Tagging by selecting the **Tagging** check box.

Step 5 Select the **Device** from the drop-down list.

Step 6 Specify the Interface. Select one of the following:

- Ethernet
- Port Channel
- vPC
- ESI

It lists the available interfaces, based on your selection.

Note
- If the device is a Cisco Nexus 7000 series switch, the Interface VLAN column appears in the table. You can provide the interface VLAN you have specified in the Resources > Device > Interfaces page. If you do not specify an interface VLAN, Cisco VTS automatically chooses a VLAN from its pool.
- After importing ESI inventory you must manually create a device group on the VLAN Pool page from VTS GUI, and add the group of ESI devices into that device group. Also make sure a corresponding VLAN pool gets created for the ESI device group.
- Add the ESI device group to L2 and L3 GW groups in Admin Domain.
- ARP suppression needs to be disabled for ESI. It can be done from VTS GUI for L2 GW group in Admin Domain.

Step 7 Choose the desired interface, and click **OK**.

Creating Router using Cisco VTS GUI

To create a router using Cisco VTS GUI:

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, on page 62 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 will be automatically generated.

- The custom VRF name accepts up to 24 characters.
- If there is no input for custom VRF name, a default VRF name will be 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 the 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 will be 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 dropdown, 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 t1-rout
  address-family ipv4 unicast
   0.0.0.0/0 Null0 254
  exit
exit
exit
router bgp 100
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.

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 + icon. The Add Network page appears.

Step 3 Enter the Network name.

Step 4 Check the External Network check box.

Step 5 Click the + 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 + icon. The Add Router page appears.

Step 8 Click the + 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 0 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 In Overlay > Network screen, on the External Network toggle switch.

Step 5 Select L2 from the drop-down.

Step 6 Select 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 will be 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, on page 26 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, on page 50.

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 200
   bgp router-id 18.18.18.18
   .
   .
   .
   neighbor 5.1.1.1
     remote-as 400
     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 will be 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 than attach the router interfaces to the external network. See Creating a Network using Cisco VTS GUI, on page 73 and Creating Router using Cisco VTS GUI, on page 75 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, on page 53 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.

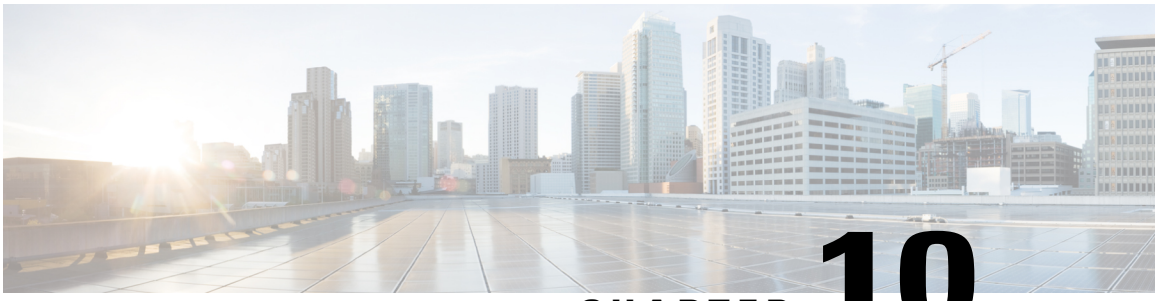

CHAPTER 10

Viewing Overlay Details

You can view details about the network at the device, VM, baremetal, network, and router levels. The following sections provide details:

- Viewing Device Details, page 81
- Viewing Virtual Machine Details, page 82
- Viewing Baremetal Details, page 82
- Viewing Network Details, page 83
- Viewing Router Details, page 84

Viewing Device Details

To view details at device level:

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

Step 2 Select the **Device tab**.
The following details are displayed:

- Device Name - Click the I icon to get the device overview. It displays the following information:
  - Ports
  - Host IP
  - Admin status
  - Oper Status
  - vPC Enabled
  - Number of VMs - Click on the **i** icon to get more details on the VM.
- Device IP
- Auth Group

- Device Type
- VLANs - Click the **i** icon to view the VLAN information.
- L3 VNI - Click the **i** icon to view VNI information.

Viewing Virtual Machine Details

To view Virtual Machine details:

Go to **Overlay** > **Virtual Machines**. The Overlay / Virtual Machines window appears.
The following details are displayed:

- VM Name
- IP Address
- Status
- Admin State
- MAC Address
- Network Name
- Host IP - Click the I icon for host details.
- VLAN
- UUID

Viewing Baremetal Details

To view baremetal server details:

Go to **Overlay** > **Baremetal**. The Overlay / Baremetal window appears.
The following details are displayed:

- UUID
- Name
- Status

- Admin State
- Network Name
- IP Address
- Binding Host
- VLAN Number
- Tenant Name

Viewing Network Details

To view details about the network:

Go to **Overlay** > **Network**. The Overlay / Network window appears.
Choose the source from the Source drop down. Then select the tenant from the Tenant drop down. The following details are displayed for the tenant you select:

- Network Name—Networks for the selected tenant. If there are shared networks, those are also displayed.
- Tenant—The tenant that owns the network.
- Zone
- Admin State
- Status
- Subnets—Count of subnet. Click on the link to get subnet details.
- Network Type
- External Network—Whether the network is an external network.
- Shared Network—Whether the network is a shared network.
- VNI
- Extended
- Multi VMM Operations

You may filter the to display all networks or only shared networks using the Select Netwrok drop down.

To edit the network details, select the Network Name check box and click **Edit** icon.

Note When you select Edit, from the context of the tenant that owns the network, every field is editable. Port attach will be enabled for that is the whose context you are in (That is, the tenant that you chose from the Tenant drop down.). If you edit a shared network, you can only do a port attach. Other fields are not editable. Port attach can be done to any of the zones of the tenant whose context you are in. You must save the port attach for the selected zone, before you proceed to do a port attach for another zone.

To delete the network details, select the Network Name check box and click **Delete (X)** icon.

Viewing Router Details

To view details about the routers in the network:

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

Step 2 Select the source from the Select Source drop down, then select the tenant from the Select Tenant drop down..
The following details are displayed:

- Name
- Zone
- UUID
- Gateway IP
- Status
- Connected Networks
- VNI
- VRF
- Provider Router
- Attached Templates—Templates that are attached to the router. Click on the link to view details of the template.

To edit the router details, select the Router Name check box and click **Edit** icon.

To delete the router details, select the Router Name check box and click **Delete (X)** icon.

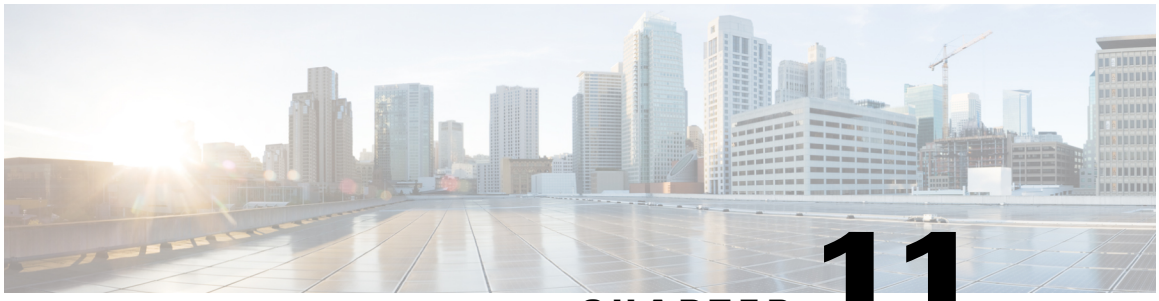

CHAPTER 11

Administering Cisco VTS

This chapter has the following topics:

- Setting up System, page 85
- Modifying Login Banners, page 86
- Setting Global Route Reflector, page 87
- Registering the Virtual Machine Manager using GUI, page 87
- Integrating Cisco VTS with Multiple Virtual Machine Managers, page 89
- Backing up the Database, page 96
- Restoring the Database, page 97
- Viewing HA Status, page 97
- Enabling External Authentication and Authorization, page 98
- Enabling Accounting and Logging, page 99

Setting up System

To set up the system:

Step 1 Go to **Administration** > **System Settings**.
The System Settings page appears.

Step 2 Enter the **DHCP Server IPv4** address. This can be a valid IPv4 address.

Step 3 Enter the **DHCP Server IPv6** address. This can be a valid IPv6 address.
You must ensure that the DHCP server is reachable from tenant leaves. The addresses need to be on the underlay side, not a management IP.

Step 4 Enter the **AnyCast GW Mac**. This is mandatory. Click **?** for information about the format.

Step 5 Choose the VTF Mode you want to use. You have the following options:

- VTEP—On OpenStack, VTF as VTEP works only in vhost mode.

- L2—If you want to use VTF as an L2 switch. VTF as L2 switch or in VTEP works only in the vhost mode.
- None—No VTF involved.

Step 6 Specify the **Out-of-Sync Commit** behavior to control the Check Sync feature. See Synchronizing Configuration, on page 26 for details about the synchronizing configuration using the Config Sync feature. Choose one of the following: Choose:

- Accept— Check sync feature in network inventory will be disabled.
- Reject— Check sync feature in network inventory will be enabled.

Step 7 Enable / disable Device South Bound Lock—Device southbound lock is enabled by default. When VTS has a redundant pair or group, it is possible for a transaction to succeed even when one or more of the redundant members are down, as long as one device is up. When the transaction comes, VTS checks the connectivity to the redundant devices and if it can not reach one of the devices, the admin state of the device will be changed to southbound-locked and the transaction configuration will only be pushed to the active devices. In order for the southbound lock feature to work, you must create a umap and provide the credentials that NSO will use, in the authgroup "vts-default" . This feature currently supports the following redundant groups:

- VPC Pair
- ESI Group
- Static Multi-Homed devices
- DCI
- VTSR

Step 8 Click **Submit**.

Modifying Login Banners

The Login Banners page lets you modify the text that appears on the VTS login page and Home page.

Step 1 Go to **Administration** > **Login Banners**. The Login Banners page appears.
Step 2 Modify the text in the Before login Text text box, to update the text that appears on the VTS login screen.
Step 3 Modify the text in the After login Text text box, to update the text that appears on the Home page after you log in.
Step 4 Click **Submit**.

Setting Global Route Reflector

You have the option to either use an inline route reflector, or global route reflector.

To set the global route reflector:

Step 1 Go to **Administration** > **Route Reflector**.

Step 2 Use the toggle switch to choose Global.

Note The Spine has to be selected as route reflector under global RR so that it is available for all other devices. This should be done before you create the admin domain.

Step 3 Select the device.

Step 4 Click **Save**.

Registering the Virtual Machine Manager using GUI

You can register the VMM using the VTS GUI. You can also specify whether the VMM you register is a trusted or an untrusted VMM.

Note For cluster-based deployments, you must install the plugin on each node.

To do this:

Step 1 Go to **Administration** > **Virtual Machine Manager**.

Step 2 Click the Add (+) button.
The Register VMM page is displayed.

Step 3 Enter the VMM Details:

- Name—Name of the VMM.
- Version —Specify the version from the drop-down.
- Mode—Whether the VMM has been registered as Trusted or Untrusted.
- VTF Deployment Mode—If you select openstack-liberty-rhel as the OpenStack version, this option appears. To install VTS in vhost mode, choose vtf-vhost. To install VTF as a VM, choose vtf-as-vm.
- API Endpoint Details—The fields differ based on the VMM you choose.
  - API Endpoint Details for OpenStack
    - API Protocol:IP Address:Port—VMM service endpoint's IPv4/IP6 address and port. Make sure you use the same IP address format (IPv4/IPv6) for all IP address fields. Mixed mode is not supported.
    - Keystone Protocol:IP Address:Port—Keystone protocol, IP address and port for OpenStack.

- Openstack Admin Project—Tenant with Administrator privileges in OpenStack. This can be any tenant with Administrator privileges. Any change to this tenant name, username, and passphrase needs to be updated in Cisco VTS for Multi-VMM operations to work properly.
- Admin User Name—admin user for the admin project in OpenStack.
- Admin Passphrase—Password of the admin user.

- API Endpoint Details for vCenter. This is optional.
  - API Protocol:IP Address:Port—VMM service endpoint's IPv4/IP6 address and port. Make sure you use the same IP address format (IPv4/IPv6) for all IP address fields. Mixed mode is not supported.
  - Datacenter—The name of the datacenter for which Cisco VTS acts as the controller.
  - Admin User Name—Username of the vCenter VMM.
  - Admin Passphrase —Password of the vCenter VMM.

Step 4 Click **Register**.
After the VMM is registered successfully, the Plugin sections opens up.

Step 5 **For OpenStack:**

a) Select the desired radio button to specify whether you want to Install plug in with Red Hat OSP Director or not. If you select Yes, enter the following details:

- OSP Director IP Address
- OSP Director User name
- OSP Director Passphrase

b) Click **Save**. The Neutron Servers section opens up.
c) Click Add (+) to add a Neutron Server. The Add Neutron Server popup is displayed.
d) Enter the Server IP Address and the Server User Name
e) Click Save and Install Plugin. You may add more Neutron Servers using the Add (+) option, if you have multiple controllers (HA Mode). The Server Plugin Installation status shows whether the installation was a success.

Note If you had opted not to use OSP Director, you will need to enter the password for the Neutron servers while adding the servers.

In case the Plugin Installation Status in the Virtual Machine Manager page shows the failure icon, you may choose to edit the VMM using the Edit option and rectify the error. Click the Server Plugin Status icon to view details of the error.

For vCenter:

a) Enter the following in the Plugin details section:

Note If you had entered the API endpoint details, the Plugin details will get populated automatically.

- IP Address : Port
- Admin User Name
- Admin Passphrase

Integrating Cisco VTS with Multiple Virtual Machine Managers

You can integrate Cisco VTS with multiple Virtual Machine Managers while managing a single data center fabric.

Note We recommend that you use an external DHCP server for your Multi VMM (MVMM) setup.

Cisco VTS, which manages hardware and software overlays, registers to multiple VMMs and enables:

- Tenant, router and network in Cisco VTS to be provisioned via Openstack or vCenter
- Cisco VTS to provision the same Tenant/Router/Network across different VMMs

The MVMM feature is supported on:

- vCenter 5.5 and vCenter 6.0
- Openstack Liberty and Newton

VMM Registration Modes

When you register a VMM with Cisco VTS, you can specify whether the VMM is a trusted VMM or an untrusted VMM. For information about registering VMMs, see Registering the Virtual Machine Manager using GUI, on page 87

Trusted VMM

A trusted VMM is one where the VMM administrator initiates service creation, and this gets reflected in VTC and the fabric. From trusted VMMs, Cisco VTS learns/discovers networks and auto-creates a network object in Cisco VTS.

In trusted mode:

- Cisco VTS registers with multiple VMMs and installs the appropriate plugins on the VMMs.
- Cisco VTS trusts the VMMs and accepts the tenant/network information published by VMM to Cisco VTS.
- VMM publishes the network information using the VTS plugin and the REST APIs exposed by Cisco VTS.

Cisco VTS supports the following variants in trusted mode:

- **Same Tenant / Disjoint Networks—**In this variant, Cisco VTS integrates with two or more VMMs, and
    - Allows the VMMs to share the same tenant, but work on disjoint networks.

- In case two or more VMMs need to share the same tenant, the operators of the VMMs have to co-ordinate on the names before sending the network information to Cisco VTS. Cisco VTS uses the tenant name and the network name to identify the tenant and network.
- Allows each VMM to create its own network to attach their respective workloads.
- Cisco VTS admin provisions an overlay router using the VTS GUI to bring the networks together by L3 routing.
- Cisco VTS admin can add an external network to the overlay router created above so that the VRF corresponding to overlay router can be extended to the DCI to facilitate MPLS L3VPN or internet connectivity.

- **Same Tenant / Same Network—**In this variant, Cisco VTS integrates with two or more VMMs, and
  - Allows the VMMs to share the same tenant, and also share the same networks, in order to attach their respective workloads.
  - In case two or more VMMs need to share the same tenant, the operators of the VMMs have to co-ordinate on the names before sending the network information to Cisco VTS.

Untrusted VMM

An untrusted VMM is one where the VMM administrator cannot create tenant/router/network service. Instead, the Cisco VTS administrator is the one who creates these services on these VMMs. Cisco VTS rejects any service creation call from an untrusted VMM.

In untrusted mode, Cisco VTS:

- Registers with multiple VMMs and installs its plugin on the VMMs.
- Does not trust the VMMs and reject the tenant/network information published by VMMs to VTS.
- Can publish the Tenant/Network information to the VMMs.

Cisco VTS supports the following variants in the untrusted mode:

- **Same Tenant / Disjoint Networks—**In this variant, Cisco VTS integrates with two or more VMMs, and
  - Allows the VMMs to share the same tenant, but work on disjoint networks.
  - In case Cisco VTS needs two or more VMMs to share the same tenant, VTS admin publishes the network information to the VMMs. VMMs sync the tenant information with Cisco VTS using the VTS plugin and the REST APIs exposed by VTS.
  - Creates disjoint networks for each of the VMMs and publishes it individually to the VMMs. VTS allows each VMM to create its own network to attach their respective workloads.
  - Cisco VTS admin provisions an overlay router using the VTS GUI to bring the networks together by L3 routing
  - Cisco VTS admin can add an external network to the Overlay router created above so that the VRF corresponding to overlay router can be extended to DCI to facilitate MPLS L3VPN or internet connectivity

- **Same Tenant / Same Network—**In this variant, VTS integrates with two or more VMMs, and

- Allows the VMMs to share the same tenant, and also the networks.
- Enables VMMs to share the same tenant. VTS admin publishes the tenant information individually to each VMM. VMM syncs the tenant information with Cisco VTS using the VTS plugin and the REST APIs exposed by Cisco VTS.
- Creates networks and publish it individually to the VMMs. Cisco VTS allows each VMM to attach their workloads to the networks

Workflows in MVMM mode of Operation

To support the above modes, Cisco VTS:

- Enables you to merge the private L2 networks on different VMMs to create a Multi VMM L2 network. The private L2 networks are created by the individual VMMs and the merge operation is controlled by the Cisco VTS administrator. Cisco VTS' involvement is to coalesce two or more network objects in the VTS database into one. After a successful merge operation, all the networks would be tied together by a unique L2 VNID. This means that the VLAN allocation scheme to VMM private L2 network remains intact. Even if there are workloads belonging to two different VMMs are placed on the same leaf node, there could be two different VLAN allocations, but the same VNI allocation. Traffic between the two workloads will go through VXLAN encap/decap. The normal mode of VNI allocation in Cisco VTS is ‘dynamic’ (per admin domain) and is assigned per private L2 network.
- Learns L2 networks from trusted VMMs and publishes these to other untrusted VMM under the control of VTS admin. The Cisco VTS GUI is used to create these networks and publish to untrusted VMMs. Cisco VTS can reuse the VNI that was assigned to the originating VMM and push that to the other untrusted VMMs. If there is no originating VMM (and VTS is the originator), then VNI allocation can happen freely in Cisco VTS.
- Creates a router that can interconnect L2 networks across multiple VMMs.

Merge and Publish Operations

Note Cisco VTS 2.5.1 supports only Merge operations. Publish operations are not supported in this release.

The VTS administrator is responsible for deciding which networks need to be merged and which networks need to be published to other VMMs.

- **Merge operation—**Cisco VTS learns and auto-creates an Mulit VMM L2 network by combining private L2 networks from multiple trusted VMMs . For a successful merge operation, the tenant name, network name, subnet name, subnet CIDR, and underlay multicast address must match. You can select one/multi/all tenants and networks within a source VMM, and then choose a list of VMMs within which the merge would be in effect. Both the source and destination VMMs need to be trusted.
- **Publish operation—**Cisco VTS initiates the creation of an Multi VMM L2 network on untrusted/trusted VMMs. This decides which network (regardless of the source) needs to be published to a list of VMMs. The VMMs can either be trusted or untrusted. Publish operation automatically pushes tenant and network information on the target VMM.

Note A merged netwrok cannot be published. To publish, you need to remove the merge definition, and the do the publish operation.

Important

- Upon publishing, Cisco VTS does not create the users for a tenant that it creates in OpenStack. To view the tenant project, user has to be assigned to the project. The OpenStack user has to attach a user to the tenant.
- Cisco VTS publishes networks to OpenStack as network type = vxlan. Before performing a publish operation, make sure that the plugin.ini, which is located at /etc/neutron/plugin.ini, has the following properties with network type vxlan as one of the values, for example:

```
type_drivers = vxlan, <network_type2>, <network_type3> ... <network_type_n> [comma separated list of network types]
tenant_network_types = vxlan, <network_type2>, <network_type3> ....<network_type_n> [comma separated list of network types]
```

Deleting Merged Networks
Individual VMMs can delete the merged networks from the VMMs as long as there is no workload attached to it. Cisco VTS will keep that network until the last VMM integrated with it deletes the network.

Deleting Published Networks All operations on published networks can be initiated only from the VTS GUI. If the network was published from VMM1 to VMM 2, then VMM 1 can remove the network, but the published network will still exist on VMM2. If VTS published the network to VMM2, then if VMM2 deletes the network, Cisco VTS will not allow to delete the network as long as the publish definition exists.
In order to delete a published network/subnet, you have to first unpublish the network, and then perform the delete operation. To unpublish a published network you need to remove the publish definition before you delete the network. To do this go to the source VMM, view the publish definition and deselect the network which you want to unpublish.

Note If there is a network which has already been published in the reverse direction, that is, from the current target to the source as per the UI, then, to unpublish it, you need to go to that target VMM, view the publish definition, and uncheck the check box for the network.

Performing Merge Operation from VMM

To initiate an merge operation from the Virtual Machine Manager page:

Step 1 Go to **Administration** > **Virtual Machine Manager**.

Step 2 Select the Source VMM and click on the Merge icon under the Multi VMM Operations column..
Merge window opens. The Source VMM is the one from which the operation is initiated. It will be selected and highlighted by default.

Step 3 Click the radio button corresponding to the Target VMM,.

The Tenants from Source column lists the tenants that are available. You can use the + New button to add a new tenant name. To add a new tenant, enter the tenant name in the text box, and click the tick icon. Click X to delete.

Note This will take effect only after the tenant is actually created.

You can use the filter to view the available tenants or selected tenants. By default, it shows all tenants.

You may use the select all button to select all tenants. If you use the select all option, you can set the Include Tenants that will be created in future automatically toggle switch to Yes.

Step 4 Select the desired tenant(s).
The Networks from Source column lists the Networks available in the source VMM, for that tenant. You can use the + New button to add a new network name. To add a new network, enter the network name in the text box, and click the tick icon. Click X to delete.

Note This will take effect only after the network is actually created.

You can use the filter to view the available networks or selected networks. By default, it shows all networks.

You may use the select all button to select all networks . If you use the select all option, you can set the Include Networks that will be created in future automatically toggle switch to Yes. This toggle switch will be set to Yes, also if you had set Include Tenants that will be created in future automatically toggle switch to Yes.

Step 5 Select the desired networks. Click **Save**.

Performing Merge Operation from Tenant

Step 1 Go to **Tenants**> **Tenant Management**.

Step 2 Select the VMM from the drop down. The tenants for the VMM are displayed.

Step 3 Click the Merge icon under the Multi VMM Operations column for the desired tenant.

Step 4 Click the radio button to select the Target VMM.

Step 5 Select the Networks from Source to be merged.
You can use the + New button to add a new network name. To add a new network, enter the network name in the text box, and click the tick icon. Click X to delete.

Note This will take effect only after the network is actually created.

You can use the filter to view the available networks or selected networks. By default, it shows all networks.

You may use the select all button to select all network. If you use the select all option, you can set the Include Networks that will be created in future automatically toggle switch to Yes.

Step 6 Click Save.

Performing Merge Operation from Network

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

Step 2 Select the source from the Select Source drop down list.

Step 3 Select the tenant from the Select Tenant drop down list.

Step 4 Click the Merge icon for the desired network.

Step 5 Select the target VMM.

Step 6 Select the network from the Network from Source column.

Step 7 Click **Save**.

Performing Publish Operation from VMM

To publish from VMM

Step 1 Go to **Administration** > **Virtual Machine Manager**.

Step 2 Select the Source VMM and click the Publish icon under the Multi VMM Operations column..
The Publish window opens. The Source VMM is the one from which the operation is initiated. It will be selected and highlighted by default.

Step 3 Click the radio button corresponding to the Target VMM,.
The Tenants from Source column lists the tenants that are available. You can use the + New button to add a new tenant name. To add a new tenant, enter the tenant name in the text box, and click the tick icon. Click X to delete.
Note This will take effect only after the tenant is actually created.

You can use the filter to view the available tenants or selected tenants. By default, it shows all tenants.

Step 4 Select the desired tenant(s).
The Networks from Source column lists the Networks available in the source VMM, for that tenant. You can use the + New button to add a new network name. To add a new network, enter the network name in the text box, and click the tick icon. Click X to delete.
Note This will take effect only after the network is actually created.

You can use the filter to view the available networks or selected networks. By default, it shows all networks.

You may use the select all button to select all networks.

Step 5 Select the desired networks. Click **Save**.

Performing Publish Operation from Tenant

To initiate a publish operation from Tenant:

Step 1 Go to **Tenants**> **Tenant Management**.

Step 2 Select the VMM from the drop down. The tenants for the VMM are displayed.

Step 3 Click the Publish icon under the Multi VMM Operations column for the desired tenant.

Step 4 Click the radio button to select the Target VMM.

Step 5 Select the Networks from Source to be merged.
You can use the + New button to add a new network name. To add a new network, enter the network name in the text box, and click the tick icon. Click X to delete.
Note This will take effect only after the network is actually created.

You can use the filter to view the available networks or selected networks. By default, it shows all networks.

Step 6 Click **Save**.

Performing Publish Operation from Network

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

Step 2 Select the source from the Select Source drop down list.

Step 3 Select the tenant from the Select Tenant drop down list.

Step 4 Click the Publish icon for the desired network.

Step 5 Select the target VMM.

Step 6 Select the network from the Network from Source column.

Step 7 Click **Save**.

Performing Publish Operation from VTS

To publish from VTS

Step 1 Go to **Administration** > **Virtual Machine Manager**.

Step 2 Click the publish from VTS icon.

The Publish window opens with the source as VTS. It is selected and highlighted by default.

Step 3 Click the radio button corresponding to the Target VMM,.
The Tenants from Source column lists the tenants that are available. You can use the + New button to add a new tenant name. To add a new tenant, enter the tenant name in the text box, and click the tick icon.
Note This will take effect only after the tenant is actually created.

Click X to delete a tenant you do not want to publish from VTS.
You can use the filter to view the available tenants or selected tenants. By default, it shows all tenants.

Step 4 Select the desired tenant(s).
The Networks from Source column lists the Networks available in the source VMM, for that tenant. You can use the + New button to add a new network name. To add a new network, enter the network name in the text box, and click the tick icon. Click X to delete.
Note This will take effect only after the network is actually created.

You can use the filter to view the available networks or selected networks. By default, it shows all networks.

You may use the select all button to select all networks.

Step 5 Select the desired networks. Click **Save**.

Backing up the Database

Perform the following tasks to backup the database.

For an HA deployment, do this on the Master.

Step 1 $ sudo su

Step 2 $ crm configure property maintenance-mode=true

Step 3 $ source /etc/profile.d/ncs.sh

Step 4 $ service nso status

```
root@vtc-2:/opt/vts/bin# service nso status
* nso.service - File for starting NCS
   Loaded: loaded (/opt/vts/etc/systemd/nso.service; enabled; vendor preset: enabled)
   Active: active (running) since Sat 2017-02-11 06:51:29 UTC; 11h ago
  Process: 4551 ExecStop=/bin/sh -lc /opt/vts/etc/systemd/ncs stop (code=exited, status=0/SUCCESS)
```

Step 5 $ service nso stop

```
ncs stop/waiting
```

Step 6 $ ncs-backup --install-dir /opt/nso

```
INFO Backup
                        /var/opt/ncs/backups/ncs-4.1.1.3@2016-11-09T16:57:07.backup created
                        successfully
```

Step 7 $ service nso start

```
ncs start/running
```

Step 8 $ crm configure property maintenance-mode=false

Restoring the Database

Do the following to restore the database.

For an HA deployment, do these only on the Master.

Note Make sure that both VTC 1 and VTC 2 passwords match with the one in the backup file.

In an L3 HA setup, if VTC 1 was the master while you had taken the backup, and at a later point if you had made VTC 2 the Master, you must still do the restore on VTC1 (which was the earlier Master), after making VTC1 the current Master.

Step 1 $ sudo su

Step 2 $ source /etc/profile.d/ncs.sh
crm configure property maintenance-mode=true

Step 3 $ service nso stop
service nso status

Step 4 $ ncs-backup --install-dir /opt/nso --restore <path of the *ncs-filename.backup.gz*>

```
root@vtc-1:~# ncs-backup --install-dir /opt/nso --restore
/home/admin/ncs-4.3.0.2@2017-02-11T18\:10\:40.backup.gz
Restore /opt/vts/etc/nso from the backup (y/n)? y
Restore /opt/vts/run/nso from the backup (y/n)? y
INFO  Restore completed successfully
```

Step 5 $ service nso start
crm configure property maintenance-mode=false
service nso status

Viewing HA Status

The High Availability page lets you view the status of nodes part of the high availability setup.

Go to **Administration** > **High Availabiity**.
The table displays the following details:

- IP Address—The IP address of the node.
- Role—The current role.

• Status—The current status.

Enabling External Authentication and Authorization

Cisco VTS allows you to integrate with a remote authentication and authorization server for user authentication and authorization. In this release, Cisco VTS supports external authentication and authorization via TACACS+ server.

You can add multiple TACACS+ servers. TACACS+ server are randomly chosen for authentication from the list of configured servers. See TACACS+ documentation for installing and configuring the TACACS+ server on IPv4/IPv6 network.

For a user logging into VTS to be able to authenticate via TACACS+ server, the VTS admin needs to setup the external authorization servers. A TACACS+ user has to be added to the user group and that user group has to be mapped to a VTS user role, which is the administrator and operator. To do this, you need to modify the TACACS+ configuration file and add users and groups to map with the VTS user role. The user group names that you need to use while you create users in TACACS+ server are:

- Administrator
- Operator

See Setting up Remote Authentication Server, on page 99 for details.

Important Notes:

- If the same username is present in both local (Cisco VTS) database and TACACS+, then the user will be first authenticated using local server. If the username is not present in the local database, or if local authentication fails due to a password mismatch, then the system tries to authenticate the user from the TACACS+server.
- Cisco VTS users and groups should be consistent across all the participating TACACS + servers.
- If the same username is configured in both local and TACACS+ server, you need to make sure the roles assigned are identical at both the places. We recommend that you have unique users in the local database and TACACS+ servers.
- If an AAA user is not assigned to any of the Cisco VTS groups in TACACS+ server, the user authentication will fail.
- AAA users, even AAA admin users, will not be able to disable AAA, but still will be able to add/delete AAA configuration.
- AAA username with special characters are not supported.
- We recommend that you use the *vts-default* authorization group while adding devices into network inventory. This is a system defined authorization group, available in Cisco VTS. If you are not using the *vts-default* authorization group, you need to ensure that you create an auth group which has AAA user added as the VTC Admin User Name.

If a TACACS+ server is unavailable, then the next server is contacted for authentication and so on till all the servers are exhausted. This process is repeated thrice. If the user cannot be authenticated or authorized all the three times, then the authentication for the external user fails.

Setting up Remote Authentication Server

To enable remote user authentication, you must configure the system to use an external authentication server. Before you begin, review the Enabling External Authentication and Authorization, on page 98 section.

Step 1 Go to **Administration** > **Remote Authentication Settings**.
The Remote Authentication Settings page appears.

Step 2 Use the **Enable Protocol** toggle button to enable the desired protocol. You must add at least one server for the selected protocol. Currently only TACACS+ is supported.

Step 3 Click **Add (+)**. The Configure TACACS + popup window appears.

Step 4 Enter the IP Address/Host Name, and the port details.

Step 5 Enter the secret key in the Key field. This can have 128 characters.

Step 6 Click **Logging** toggle button to enable the accounting.
For more details about accounting and logging, see Setting up Accounting and Enabling Accounting and Logging.

To delete a TACACS+ server, select the check box corresponding to the server, click delete **(X)**, and then click **Save**.

Enabling Accounting and Logging

The admin can select one of the TACACS+ Server as a logging server. Audit logs are sent to that server. In addition to that server, the audit logs will also be logged to the local log file (present in Cisco VTS).

On the TACACS+ server where you have enabled logging, you can find the log files at */var/log/tac_plus.acct*.

The Cisco VTS location where you can find the log file is */opt/vts/log/nso/vts-accounting.log*.

Logs are collected every 120 seconds (default setting).

- Logs are collected every 120 seconds(default setting)
  Following are the fields that can be found in the log:
  - Client IP—Client IP from where the request was made
  - Server IP—VTS server IP
  - User Name—User who performs the transaction
  - Message—The model change in the transaction or the REST API url
  - Date/Time—The time when the change was made
  - Application Name—VTS (static value)
  - Operation Type—Derived from the change, could be CREATE, UPDATE or DELETE

◦ Status—Success or Error (static value)

Setting up Accounting

To set up accounting, you must add one of the TACACS+ servers that are registered with Cisco VTS as the logging server. You can do this while you add the remote authorization servers. If you have already added remote authentication servers, you can select a server and edit it to make it the logging server.

Note You can have only one TACACS+ server as the logging server at a time.

Step 1 Go to **Administration** > **Remote Authentication Settings**.
The Remote Authentication Settings page appears.

Step 2 Use the **Enable Protocol** toggle button to enable the desired protocol. You must add at least one configuration instance for the selected protocol. Currently only TACACS+ is supported.

Step 3 Click **Add (+)**. The Configure TACACS+ popup window appears.

Step 4 Enter the IP Address/Host Name, and the port details.

Note Cisco VTS supports IPv4 and IPv6 addresses.

Step 5 Enter the secret key in the Key field. This can have 128 characters.

Step 6 Click **Logging** toggle button to enable the accounting.

Step 7 Click **Add**.

Step 8 Click **Save**.
The logs get saved in the local VTS server and TACACS server.
In Cisco VTS, you can see all the logs in vts-accounting.log, which has details like the Username, Date/Time, Application Name, Operation Type, Status, Sever IP, Client IP address, and the exact message about the transaction. Similarly, in the TACACS server also you can see all the logs for the transactions.

Logs are collected every 120 seconds(default setting), and pushed to TACACS+ accounting server(for example, tac_plus.acct) and to VTC(vts-accounting.log).

The log file will be rotated once it reaches 100MB in size. The backup exists for 10 rotations, then gets deleted.

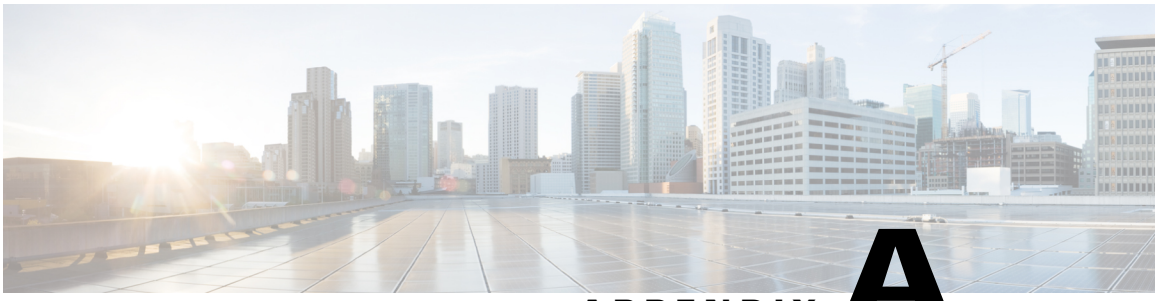

APPENDIX A

Service Extension Templates-Supported Configuration

This appendix provides information about supported service extension template configuration.

- Supported Service Extension Template Configuration Examples for Cisco Nexus 7000 Series Switches, page 101
- Supported Service Extension Template Configuration for Cisco Nexus 9000 Series Switches, page 104
- Supported Service Extension Template Configuration for Cisco ASR 9000 Series Routers , page 105

Supported Service Extension Template Configuration Examples for Cisco Nexus 7000 Series Switches

This section provides sample service extension configuration that Cisco VTS supports for Cisco 7000 series switches.

Interface Configuration

```
interface port-channel 6.1001

            description towards PE
            encapsulation dot1q 1001
            vrf member vrf1
            bfd interval 50 min_rx 50 multiplier 3
            no ip redirects
            ip address 10.23.65.1/31
            no ipv6 redirects
            no shutdown

interface port-channel 6.1041

            description towards PE
            encapsulation dot1q 1041
            vrf member ED
            bfd interval 50 min_rx 50 multiplier 3
            no ip redirects
            ipv6 address 2009:2009:3112::1/126
            no ipv6 redirects
            no shutdown
interface vlan 2471
  no shut
```

```
  vrf member WSN_LOOP
  ip address 69.83.136.133/28
  vrrpv3 1 address-family ipv4
    address 69.83.136.132
```

Note This configuration works for IPv4 and IPv6.

```
interface vlan 2475
  no shut
  vrf member WSN_LOOP
  ipv6 address 2001:4888:39:3080:308:201:0:1/64
  vrrpv3 1 address-family ipv6
    address fe80::1 primary
    address 2001:4888:39:3080:308:201::
```

Note This configuration works for IPv4 and IPv6.

```
interface vlan 2400
  no shut
  vrf member WSN_LOOP
  ip address 15.15.15.1/24
  ipv6 address xxxx::x/64 ! dual stack or create a difference interface
```

Note This configuration works for IPv4 and IPv6.

VRF Configuration

```
vrf context vrf1

            ip route 10.32.10.0/24 10.42.10.4
            ip route 10.52.10.0/24 10.42.10.4
            ipv6 route 2001:db8::5/128 2001:db8:0:1:2a0:a502:0:19da
vrf context WSN_MOBILE
  ip route 0.0.0.0/32 Ethernet2/5 10.2.56.6 track 10
  rd auto
  address-family ipv4 unicast
    route-target both auto
    route-target both auto evpn
  address-family ipv6 unicast
    route-target both auto
    route-target both auto evpn
```

Note This works configuration for IPv4 and IPv6.

```
vrf context WSN
  vni 50001
  ip route static bfd Vlan1050 68.50.50.50
  ip route 210.0.0.1/32 Vlan1050 68.50.50.50
```

Note This works configuration for IPv4 and IPv6.

Note User should enter the vlan value other than 1000-2000, which is reserved for bridge domain. If the range is configured differently, make sure the values within that range is not used.

```
vrf context WSN_VPER
    ip route 0.0.0.0/32 port-channel 110.2513 69.83.32.37 track 1
```

```
    ip route 0.0.0.0/32 vlan 2500 16.16.16.2 track 2 200
    ipv6 route ::/128 port-channel 110.2577 2001:4888:16:2079:1e1:2a1:: track 1
    ipv6 route ::/128 vlan 2500 <v6 address of SVI on other BL> track 2 200
```

Note This configuration works for IPv4 and IPv6.

Router BGP Configuration

```
router bgp 6167

vrf vrf1

       local-as 65004
       address-family ipv4 unicast
           network 10.32.10.0/24 route-map vrf1_LOCAL_COMMUNITIES
           network 10.52.10.0/24 route-map vrf1_LOCAL_COMMUNITIES
           advertise l2vpn evpn
       neighbor 10.23.65.0 remote-as 22394
           bfd
           password 3 XXXX
           description towards PE
           address-family ipv4 unicast
                  send-community
                  route-map vrf1_ROUTE_POLICY in
                  route-map vrf1_LOCAL_ROUTE_POLICY out
router bgp 65004
  Vrf WSN
    router-id 192.168.0.25
    address-family ipv4 unicast
      network 150.0.0.1/32 route-map ONLY_FABRIC
      advertise l2vpn evpn
     redistribute direct route-map vts-subnet-policy augumentation and deviation
      redistribute static route-map staticMap
      maximum-paths 32
      maximum-paths ibgp 32
    address-family ipv6 unicast
      advertise l2vpn evpn
      redistribute direct route-map vts-subnet-policy
      maximum-paths 32
      maximum-paths ibgp 32
    neighbor 68.50.50.50
      bfd
      remote-as 6167
      address-family ipv4 unicast
        send-community
        send-community extended
    neighbor 210.0.0.1
      bfd
      remote-as 6167
      update-source loopback150
      ebgp-multihop 255
      address-family ipv4 unicast
        send-community
        send-commuqnity extended
```

Note This configuration works for IPv4 only.

ICMP v6 Configuration

```
ip sla 11
icmp-echo 2009:2009:2009:10:1:56:0:5
vrf WSN_LOOP
threshold 500
timeout 500
frequency 1
ip sla schedule 11 life forever start-time now
```

Note This configuration works for IPv6 only.

Supported Service Extension Template Configuration for Cisco Nexus 9000 Series Switches

This section provides sample service extension configuration Cisco VTS supports for Cisco 9000 series switches.

Interface Configuration

```
int vlan 2471
  no shut
  vrf member WSN_LOOP
  ip address 69.83.136.133/28
  vrrpv3 1 address-family ipv4
    address 69.83.136.132
```

Note This configuration works for IPv4 and IPv6.

```
int vlan 2475
  no shut
  vrf member WSN_LOOP
  ipv6 address 2001:4888:39:3080:308:201:0:1/64
  vrrpv3 1 address-family ipv6
    address fe80::1 primary
    address 2001:4888:39:3080:308:201::
```

Note This configuration works for IPv4 and IPv6.

```
interface vlan 2400
  no shut
  vrf member WSN_LOOP
  ip address 15.15.15.1/24
  ipv6 address xxxx::x/64 ! dual stack or create a difference interface
```

VRF Configuration

```
vrf context WSN_MOBILE
  ip route 0.0.0.0/32 Ethernet2/5 10.2.56.6 track 10
  rd auto
  address-family ipv4 unicast
    route-target both auto
    route-target both auto evpn
  address-family ipv6 unicast
    route-target both auto
    route-target both auto evpn
```

Note This works configuration for IPv4 and IPv6.

```
vrf context WSN
  vni 50001
  ip route static bfd Vlan1050 68.50.50.50
  ip route 210.0.0.1/32 Vlan1050 68.50.50.50
```

Note This works configuration for IPv4 and IPv6.

```
vrf context WSN_VPER
    ip route 0.0.0.0/32 port-channel 110.2513 69.83.32.37 track 1
    ip route 0.0.0.0/32 vlan 2500 16.16.16.2 track 2 200
    ipv6 route ::/128 port-channel 110.2577 2001:4888:16:2079:1e1:2a1:: track 1
    ipv6 route ::/128 vlan 2500 <v6 address of SVI on other BL> track 2 200
```

Note This configuration works for IPv4 and IPv6.

Router BGP Configuration

```
router bgp 100
Vrf WSN
    router-id 192.168.0.25
    address-family ipv4 unicast
      network 150.0.0.1/32 route-map ONLY_FABRIC
      advertise l2vpn evpn
     redistribute direct route-map vts-subnet-policy
      redistribute static route-map staticMap
      maximum-paths 32
      maximum-paths ibgp 32
    address-family ipv6 unicast
      advertise l2vpn evpn
      redistribute direct route-map vts-subnet-policy
      maximum-paths 32
      maximum-paths ibgp 32
    neighbor 68.50.50.50

      bfd
      remote-as 6167
      address-family ipv4 unicast
        send-community
        send-community extended
    neighbor 210.0.0.1
      bfd
      remote-as 6167
      update-source loopback150
      ebgp-multihop 255
      address-family ipv4 unicast
        send-community
        send-commuqnity extended
```

Note This configuration works for IPv4 only.

Supported Service Extension Template Configuration for Cisco ASR 9000 Series Routers

This section provides information about the service extension configuration Cisco VTS supports for Cisco ASR 9000 series routers.

Router OSPF Configuration

```
router ospf 700
  log adjacency changed detail
  router-id 16.16.16.16
  timers throttle lsa all 0 20 5000
```

```
  timers throttle spf 50 100 5000
  timers lsa min-arrival 15
  auto-cost reference-bandwidth 80000
  area 0
    network point-to-point
    interface GigabitEthernet0/0/0/2
      authentication
      message-digest-key 1 md5 encrypted 07982c55db2b9985d3391f02e639db9c
      network point-to-point
      passive enable
     !
  !
  vrf admin-r-1
  !
!
```

Router Static Configuration

```
router static
  address-family ipv4 unicast
    0.0.0.0/0 172.20.100.1
  !
!
```

Router BGP Configuration

```
router bgp 200
  bgp router-id 49.1.1.1
  address-family ipv4 unicast
    maximum-paths ebgp 2
    maximum-paths ibgp 2
    !
  neighbor-group ng1
    remote-as 100
    password encrypted 07982c55db2b9985d3391f02e639db9c
    update-source Loopback0
    address-family ipv4 unicast
      next-hop-self
    !
  !
  vrf admin-r-1
  rd auto
  bgp router-id 49.1.1.1
  address-family ipv4 unicast
  !
  neighbor 13.1.1.8
    remote-as 100
    address-family ipv4 unicast
      route-policy vts-route-policy in
      default-originate
    !
   !
  !
!
```

VRF Configuration

```
vrf admin-r-1
  address-family ipv4 unicast
  !
!
```

Interface/{*any*}-subinterface Configuration

```
interface GigabitEthernet0/0/0/1.1
  vrf admin-r-1
  ipv4 address 200.200.200.2 255.255.255.0
  encapsulation dot1q 1002
!
```

We support the following subinterfaces:

```
TenGigE-subinterface
FortyGigE-subinterface
HundredGigE-subinterface
FastEthernet-subinterface
GigabitEthernet-subinterface
Bundle-Ether-subinterface
```

Interface BVI Configuration

```
interface BVI 1003
  service-policy input bvi-policymap
  vrf admin-router
  !
!
```

Interface NVE Configuration

```
interface nve1
  description desc123
  vrf admin-router
  shutdown
  !
!
```

l2vpn Configuration

```
l2vpn
  bridge group bg-name123
    bridge-domain-name
      interface GigabitEthernet
    !
  !
!
Any interface:
Subinterfaces:
TenGigE
FortyGigE
HundredGigE
FastEthernet
GigabitEthernet
Bundle-Ether
```

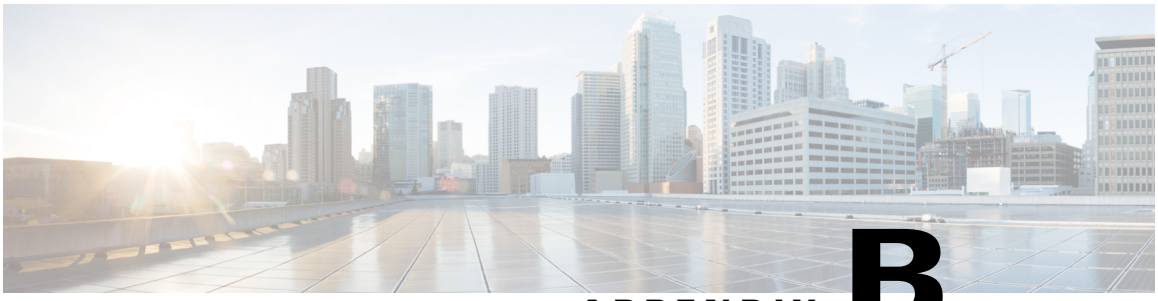

APPENDIX B

Supported Underlay Configuration

This appendix provides information about supported underlay template configuration.

- Supported Underlay Configuration, page 109

Supported Underlay Configuration

The following table lists the supported underlay configuration:

Configuration Area	Sample Configuration	Device Type	Device Role

Note Prerequisite: You should enable the following configuration for the device configuration to work.	`feature telnet`
	`feature nxapi`
	`feature ospf`
	`feature bgp`
	`feature pim`
	`feature udld`
	`feature interface-vlan`
	`feature vn-segment-vlan-based`
	`feature hsrp`
	`feature lacp`
	`feature vpc`
	`feature lldp`
	`feature nv overlay`
	`feature pbr`
	`feature sla sender`
	`feature sla responder`
	`feature vrrpv3`
	`feature bfd`		

UNDERLAY IGP ROUTING OSPF routing process	`router ospf 10`	N9K	Leaf
	`router ospf 10`	ASR9K	DCI
OSPF Area	`interface ethernet 1/5` `  ip ospf router 10 area 0.0.0.0`	N9K	Leaf
	`interface ethernet 1/5` `  ip ospf router 10 area 0.0.0.0`	N9K	Leaf
	`router ospf 10` `  area 0`	ASR9K	DCI

OSPF router-id	`router ospf 10` `  router-id 10.218.20.15`	N9K	Leaf
	`router ospf 10` `  router-id 10.218.20.15`	ASR9K	DCI

OSPF auto-cost reference	`router ospf 10`
	`!`
	`  auto-cost reference-bandwidth 800000`
		ASR9K	DCI

OSPF Network type	`interface ethernet1/5` `  ip ospf network point-to-point`	N9K	Leaf
	`interface vlan10` `  ip ospf network point-to-point`	N9K	Leaf

	`router ospf 10`
	`  area 0`
	`  interface GigabitEthernet0/0/1/3`
	`  network point-to-point`
		ASR9K	DCI

OSPF Authentication	`interface Ethernet1/5` `  ip ospf authentication message-digest`	N9K	Leaf
	`interface Ethernet1/5` `  ip ospf message-digest-key 1 md5 0 xxx`	N9K	Leaf

	`router ospf 10`
	`  area 0`
	`  interface <Fabric Interface>`
	`  authentication message-digest`
	`  message-digest-key 1 md5 encrypted`
	`202cb962ac59075b964b07152d234b70`
		ASR9K	DCI

OSPF Passive-interface	`interface loopback3` `  ip router ospf 100 area 0.0.0.0`	N9K	Leaf

	`router ospf 10`
	`  area 0`
	`  interface Loopback10`
	`  passive enable`
		ASR9K	DCI

OSPF Convergence	`router ospf 10`
	`  timers lsa arrival 15`
	`  timers throttle lsa 0 20 5000`
	`  timers throttle spf 50 100 5000`
		N9K	Leaf

	`router ospf 10`
	`  timers throttle lsa all 0 20 5000`
	`  timers throttle spf 50 100 5000`
	`  timers lsa min-arrival 15`
		ASR9K	DCI

OSPF BFD (per-link)	`feature bfd`
	`router ospf 10`
	`  bfd`
		N9K	Leaf

	`interface Ethernet1/5` `  no ip redirects`	N9K	Leaf

	`router ospf 10`
	`  bfd minimum-interval 150`
	`  bfd multiplier 3`
	`  area 0`
	`  interface TenGigE0/0/2/1`
	`  bfd fast-detect`
		ASR9K	DCI

	`interface vlan 10` `  no bfd echo`	N9K	Leaf

Multicast Routing	`feature pim`	N9K	Leaf

	`interface loopback1`
	`  ip address 10.218.20.250/32`
	`  ip router ospf 10 area 0.0.0.0`
	`  ip pim sparse-mode`
		N9K	Spine

	`ip pim rp-address 10.218.20.250 group-list 239.255.0.0/16 override`	N9K	Spine
	`ip pim anycast-rp 10.218.20.250 10.218.20.249` `ip pim anycast-rp 10.218.20.250 10.218.20.248`	N9K	Spine
	`feature pim`	N9K	Leaf
	`ip pim rp-address 10.218.20.250 group-list 239.255.0.0/16 override`	N9K	Leaf
	`interface Vlan10` `  ip pim sparse-mode`	N9K	Leaf
	`interface loopback0` `  ip pim sparse-mode`	N9K	Leaf
	`interface Ethernet2/1` `  ip pim sparse-mode`	N9K	Leaf
	`interface Ethernet2/2` `  ip pim sparse-mode`	N9K	Leaf
L2 Technologies	`interface Ethernet 1/10` `  switchport mode trunk`	N9K	Leaf
	`interface Ethernet 1/10` `  switchport trunk allowed vlan none`	N9K	Leaf
	`interface Ethernet 1/10` `  spanning-tree port type edge trunk`	N9K	Leaf
	`interface Ethernet 1/10` `  spanning-tree bpduguard enable`	N9K	Leaf
	`interface Ethernet 1/10` `  spanning-tree bpdufilter enable`	N9K	Leaf
	`interface Ethernet 1/10` `  storm-control broadcast level 20.0`	N9K	Leaf
	`interface Ethernet 1/10` `  storm-control multicast level 30.0`	N9K	Leaf
	`interface Ethernet 1/10` `  storm-control unicast level 50.0`	N9K	Leaf
	`interface Ethernet 1/10` `  storm-control action shutdown`	N9K	Spine
vPC Role and Priority	`vpc domain 1` `  role priority 100`	N9K	Leaf
	`vpc domain 1` `  role priority 200`	N9K	Leaf
vPC Peer Keep-alive Link	`vrf context management`	N9K	Leaf
	`interface mgmt 0` `  vrf member management`	N9K	Leaf

	`interface mgmt 0`
	`  ip address 172.20.118.119/25`
	`  no shutdown`
		N9K	Leaf

	`vpc domain 1` `  peer-keepalive destination 172.20.118.20`	N9K	Leaf

vPC Peer-Link	`interface Ethernet 1/1`
	`  spanning-tree port type network`
	`  channel-group 1 mode active no shutdown`
		N9K	Leaf

	`interface Ethernet 1/2`
	`  spanning-tree port type network`
	`  channel-group 1 mode active no shutdown`
		N9K	Leaf

	`interface port-channel1`
	`  switchport`
	`  switchport mode trunk`
	`  spanning-tree port type network`
	`  vpc peer-link`
		N9K	Leaf

vPC Port	`interface Ethernet 2/9` `  channel-group mode active id 51`	N9K	Leaf
	`interface port-channel 51` `  switchport`	N9K	Leaf

	`interface port-channel 51`
	`  switchport`
	`  vpc 51`
		N9K	Spine

vPC Peer-Switch Option	`vpc domain 1` `  peer-switch`	N9K	Leaf
vPC ARP Synchronization	`vpc domain 1` `  ip arp synchronize`	N9K	Leaf

vPC in VXLAN environment adjustment	`vpc domain 10`
	`  peer-switch`
	`  system-priority 100 ( could not find this option)`
	
	`  peer-keepalive destination 172.20.118.120`
	`  delay restore 200`
	`  peer-gateway`
	`  ip arp synchronize`
		N9K	Leaf

	`interface port-channel1` `description vPC peer-link`	N9K	Spine
	`interface port-channel1` `  description vPC switchport mode trunk`	N9K	Leaf
	`interface port-channel 1` `  description vPC switchport mode trunk`	N9K	
	`interface port-channel 1` `  description vPC spanning-tree port type network`	N9K	
	`interface port-channel 1` `  vpc peer-link`	N9K	
	`interface port-channel 10` `  switchport trunk allowed vlan`	N9K	
	`interface port-channel 10` `  spanning-tree port type edge trunk`	N9K	
	`interface port-channel 10` `  spanning-tree bpdufilter enable`	N9K	
	`interface port-channel 10` `  spanning-tree bpduguard enable`	N9K	
	`interface port-channel 10` `  vpc 10`	N9K	Spine
	`interface Ethernet 1/10` `  switchport trunk allowed vlan none`	N9K	Spine
	`interface Ethernet 1/10` `  spanning-tree port type edge trunk`	N9K	Leaf
	`interface Ethernet 1/10` `  spanning-tree bpduguard enable`	N9K	Leaf
	`interface Ethernet 1/10` `  spanning-tree bpdufilter enable`	N9K	Leaf
	`interface Ethernet 1/10` `  channel-group 10 mode active`	N9K	Leaf
	`interface loopback 0` `  ip address 10.218.20.1/32`	N9K	Leaf
	`interface loopback 0` `  ip address 10.218.20.1/32 secondary`	N9K	Leaf
	`interface loopback 0` `  ip router ospf 100 area 0.0.0.0`	N9K	Leaf
	`interface loopback 0` `  ip pim sparse-mode`	N9K	Leaf
	`interface Vlan 10` `  ip address 10.218.18.0/31`	N9K	Leaf

	interface Vlan 10
	  description Underlay vPC Backup link
	  no shutdown
	  no bfd echo
		N9K	Spine

	interface Vlan 10 ip ospf network point-to-point	N9K	Leaf
	interface Vlan 10 ip router ospf 100 area 0.0.0.0	N9K	Leaf
	interface Vlan10 ip pim sparse-mode	N9K	Leaf

STP	interface Ethernet 1/10 switchport mode trunk	N9K	Leaf
	interface Ethernet 1/10 switchport mode trunk allowed vlan 10	N9K	Leaf
	interface Ethernet 1/10 spanning-tree port type edge trunk	N9K	Leaf
	interface Ethernet 1/10 spanning-tree bpduguard enable	N9K	Leaf
	nx:interface Ethernet 1/10 spanning-tree bpdufilter enable	N9K	Leaf
	interface Ethernet 1/10 no shutdown	N9K	
	interface port-channel 10 switchport mode trunk	N9K	Leaf
	interface port-channel 10 switchport mode trunk trunk allowed vlan ids 1	N9K	Leaf
	interface port-channel 10 spanning-tree port type edge	N9K	Leaf
	interface port-channel 10 spanning-tree bpduguard enable	N9K	Leaf
	interface port-channel 10 spanning-tree bpdufilter enable	N9K	Leaf
	interface port-channel 10 no shutdown	N9K	Leaf
	interface port-channel 10 vpc port-channel-number 10	N9K	Leaf
	interface Ethernet 1/10 switchport mode trunk	N9K	Leaf
	interface Ethernet 1/10 switchport mode trunk allowed vlan 10	N9K	Leaf
	interface Ethernet 1/10 spanning-tree port type edge trunk	N9K	Leaf
	interface Ethernet 1/10 spanning-tre guard root	N9K	Leaf
	interface Ethernet 1/10 no shutdown	N9K	Leaf
	interface ethernet <xxxx> description <leaf/Spine Fabric> ip address	N9K	Leaf
	interface Ethernet 1/10 description leaf mtu 9216	N9K	Leaf
	interface Vlan 1 description <attachment/border facing intf>	N9K	Leaf
	interface Vlan 1 description ip address <addr>	N9K	Leaf
	interface Vlan 1 description ip address mtu 1500	N9K	Leaf
	interface <To Spine> mtu 9214	N9K	Leaf
	interface <To Border Leaf> mtu 1518	N9K	Leaf

	interface GigabitEthernet0/0/1/5 mtu 9214	ASR9K	DCI
Nexus 9500 QOS	system qos service-policy type queuing output default-out-policy	N9K	Leaf
	policy-map type network-qos Jumbo-nq-policy class type network-qos c-nq3	N9K	Leaf
	policy-map type network-qos Jumbo-nq-policy class type network-qos c-nq3 match qos-group 3	N9K	Leaf
	policy-map type network-qos Jumbo-nq-policy class type network-qos c-nq3 mtu 9216	N9K	Leaf

	class type network-qos c-nq3
	  match qos-group 3
	  mtu 9216
	class type network-qos c-nq2
	  match qos-group 2
	  mtu 9216
	class type network-qos c-nq1
	  match qos-group 1
	  mtu 9216
	class type network-qos c-nq-default
	  match qos-group 0
	  mtu 9216
		N9K	Leaf

	system qos service-policy type network-qos Jumbo-nq-policy	N9K	Leaf
QoS Hardware resources configuration	"hardware access-list tcam region racl 0 hardware access-list tcam region e-racl 0 hardware access-list tcam region span 0 hardware access-list tcam region vqos 256 hardware access-list tcam region e-qos 256 hardware access-list tcam region arp-ether 256"	N9K	Leaf
N 9500 QoS	system qos service-policy type queuing output default-out-policy	N9K	Leaf
	system qos service-policy type network-qos Jumbo-nq-policy	N9K	Leaf

N 9500 QoS Queuing policy	policy-map type queuing default-out-policy
	class type queuing c-out-q3
	  priority level 1
	class type queuing c-out-q2
	  bandwidth remaining percent 0
	class type queuing c-out-q1
	  bandwidth remaining percent 0
	class type queuing c-out-q-default
	bandwidth remaining percent 100
		N9K	Leaf

	System qos Service-policy type queuing out default-out-policy	N9K	Leaf

N 9500 QoS Queuing Policy	policy-map type queuing default-out-policy class type queuing c-out-q3 priority level 1 class type queuing c-out-q2 bandwidth remaining percent 0 class type queuing c-out-q1 bandwidth remaining percent 0 class type queuing c-out-q-default bandwidth remaining percent 100	N9K	Leaf
	System qos Service-policy type queuing out default-out-policy	N9K	Leaf
Network Management Ethernet (Mgmt0)	interface mgmt0 ip address 10.218.22.10/24	N9K	Leaf
	vrf context management ip route 0.0.0.0/0 10.218.23.254	N9K	Leaf
Configuring Hostname on Nexus 9000	hostname nw_lf_cnx9_001.41gebz_o01_s01	N9K	Leaf
Time Zone and day-light saving	clock timezone EET 2 0 clock summer-time EEST 4 Sunday March 02:00 4 Sunday October 03:00 60	N9K	Leaf
DNS	ip domain-name <cust_name> no ip domain-lookup	N9K	Leaf
SNMP	snmp-server contact <contact_name> snmp-server location <location_name>	N9K	Leaf
	snmp-server host 85.29.26.36 traps version 2c <SNMP_Community_1> snmp-server host 85.29.56.136 traps version 2c <SNMP_Community_1> snmp-server host 85.29.60.191 traps version 2c <SNMP_Community_1> snmp-server host 85.29.60.235 traps version 2c <SNMP_Community_1> snmp-server host 213.74.189.232 traps version 2c <SNMP_Community_1> snmp-server host 213.74.189.233 traps version 2c <SNMP_Community_1>	N9K	Leaf
	snmp-server host 85.29.26.36 use-vrf management snmp-server host 85.29.56.136 use-vrf management snmp-server host 85.29.60.191 use-vrf management snmp-server host 85.29.60.235 use-vrf management snmp-server host 213.74.189.232 use-vrf management snmp-server host 213.74.189.233 use-vrf management	N9K	Leaf
	snmp-server source-interface trap mgmt0	N9K	Leaf
	snmp-server community <community> group network-admin	N9K	Leaf
	15 permit ip host 213.74.197.43 any ... 390 permit ip host 176.43.250.25 any	N9K	Leaf

LLDP on Nexus 9000	`feature lldp`	N9K	Leaf

Network Security Disable IP Redirects	`interface Ethernet slot#/port#`
	`  no ip redirects`
	`  no ipv6 redirects`
		N9K	Leaf

Device Access Security	`NX-OS(config)#no ssh server enable`
	`NX-OS(config)#ssh key {dsa [force] | rsa [bits [force]]}`
	`NX-OS(config)#ssh server enable`
	`NX-OS#show ssh key`
	`**********`
	`rsa Keys generated:Fri Apr 10 20:13:21 2010`
	`<clipped> !`
		N9K	Leaf

AAA-N	`NX-OS(config)#feature tacacs+`
	`NX-OS(config)#tacacs-server host {ipv4-address | ipv6-address | host-name}`
	`NX-OS(config)#tacacs-server key [0 | 7] key-value`
	`NX-OS(config)#aaa group server tacacs+ group-name`
	
	`server {ipv4-address | ipv6-address | host-name}`
	`deadtime minutes`
	`use-vrf vrf-name`
	`NX-OS(config)#tacacs-server timeout seconds`
	`NX-OS(config)#tacacs-server host {ipv4-address | ipv6-address | host-name} port tcp-port`
	`NX-OS(config)#tacacs-server deadtime minutes`
		N9K	Leaf

	`feature tacacs+`
	`aaa group server tacacs+ TacacsGroup`
	`  use-vrf management`
	`  server 10.35.175.1`
	`aaa authentication login console group TacacsGroup`
	`aaa authentication login default group TacacsGroup`
	`aaa authentication login error-enable`
	`!`
	`tacacs-server host 10.35.175.1 key <shared-key> port 49`
	`tacacs-server directed-request`
	`ip tacacs source-interface mgmt 0`
	`!`
	`! Device Login Authorisation with AAA`
	`!`
		N9K	Leaf

	`aaa authorization config-commands default group TacacsGroup local`
	`aaa authorization commands default group TacacsGroup local`
	`!`
	`! Device Login Accounting with AAA`
	`!`
	`aaa accounting default group TacacsGroup`
	`!`
	`! Local User Configuration`
	`!`
	`username admin Pword <Pword> role network-admin`
		N9K	Leaf

Device Hardening 3.9.9.4 COPP policy and class maps	policy-map type control-plane
	copp-system-p-policy-strict
	  class copp-system-p-class-l3uc-data
	  set cos 1
	  police cir 250 pps bc 32 packets conform
	transmit violate drop
		N9K	Leaf

	class copp-system-p-class-critical
	  set cos 7
	  police cir 19000 pps bc 128 packets conform
	transmit violate drop
		N9K	Leaf

	class copp-system-p-class-important
	  set cos 6
	  police cir 3000 pps bc 128 packets conform
	transmit violate drop
		N9K	Leaf

	class copp-system-p-class-multicast-router
	  set cos 6
	  police cir 3000 pps bc 128 packets conform
	transmit violate drop
		N9K	Leaf

	class copp-system-p-class-management
	  set cos 2
	  police cir 3000 pps bc 32 packets conform
	transmit violate drop
		N9K	Leaf

	class copp-system-p-class-multicast-host
	  set cos 1
	  police cir 2000 pps bc 128 packets conform
	transmit violate drop
		N9K	Leaf

	class copp-system-p-class-l3mc-data
	  set cos 1
	  police cir 3000 pps bc 32 packets conform
	transmit violate drop
		N9K	Leaf

	class copp-system-p-class-normal
	  set cos 1
	  police cir 1500 pps bc 32 packets conform
	transmit violate drop
		N9K	Leaf

	class copp-system-p-class-ndp
	  set cos 6
	  police cir 1500 pps bc 32 packets conform
	transmit violate drop
		N9K	Leaf

	class copp-system-p-class-normal-dhcp
	  set cos 1
	  police cir 300 pps bc 32 packets conform
	transmit violate drop
		N9K	Leaf

	class
	copp-system-p-class-normal-dhcp-relay-response
	  set cos 1
	  police cir 400 pps bc 64 packets conform
	transmit violate drop
		N9K	Leaf

	class copp-system-p-class-normal-igmp
	  set cos 3
	  police cir 6000 pps bc 64 packets conform
	transmit violate drop
		N9K	Leaf

	class copp-system-p-class-redirect
	  set cos 1
	  police cir 1500 pps bc 32 packets conform
	transmit violate drop
		N9K	Leaf

	class copp-system-p-class-exception
	  set cos 1
	  police cir 50 pps bc 32 packets conform transmit
	violate drop
		N9K	Leaf

	class copp-system-p-class-exception-diag
	  set cos 1
	  police cir 50 pps bc 32 packets conform transmit
	 violate drop
		N9K	Leaf

	class copp-system-p-class-monitoring
	  set cos 1
	  police cir 300 pps bc 128 packets conform
	transmit violate drop
		N9K	Leaf

	class copp-system-p-class-l2-unpoliced
	  set cos 7
	  police cir 20000 pps bc 8192 packets conform
	transmit violate drop
		N9K	Leaf

	class copp-system-p-class-undesirable
	  set cos 0
	  police cir 15 pps bc 32 packets conform transmit
	 violate drop
		N9K	Leaf

	class copp-system-p-class-fcoe
	  set cos 6
	  police cir 1500 pps bc 128 packets conform
	transmit violate drop
		N9K	Leaf

	class copp-system-p-class-nat-flow
	  set cos 7
	  police cir 100 pps bc 64 packets conform
	transmit violate drop
		N9K	Leaf

	class copp-system-p-class-l2-default
	  set cos 0
	  police cir 50 pps bc 32 packets conform transmit
	 violate drop
		N9K	Leaf

	class class-default
	  set cos 0
	  police cir 50 pps bc 32 packets conform transmit
	 violate drop
		N9K	Leaf

	N9k-ST-Leaf-01# sh copp status Last Config Operation: None Last Config Operation Timestamp: None Last Config Operation Status: None Policy-map attached to the control-plane: copp-system-p-policy-strict	N9K	Leaf
	N9k-ST-Leaf-01# sh copp profile ? dense Display dense profile lenient Display lenient profile moderate Display moderate profile strict Display strict profile	N9K	Leaf
BFD	feature bfd bfd interval 50 min_rx 50 multiplier 3	N9K	Leaf
	router ospf UNDERLAY bfd	N9K	Leaf
	router bgp 100 vrf vrf1 address-family ipv4 unicast	N9K	Leaf

	router bgp 100
	  vrf vrf1
	  local-as 65004
		N9K	Leaf

	router bgp 100
	vrf vrf1
	  neighbor 10.23.65.0 remote-as 22394
	  bfd
		N9K	Leaf

OSPF Routing Process	`feature ospf`
	`!`
	`router ospf UNDERLAY`
	`  log-adjacency-changes detail`
	`  bfd`
		N9K	Leaf

OSPF Router ID	`router ospf UNDERLAY`
	`  log-adjacency-changes detail`
	`  bfd`
	`  router-id <loopback17-ip-address>`
		N9K	Leaf

Enabling OSPF on interfaces	`router ospf UNDERLAY` `  passive-interface default`	N9K	Leaf

	`continue from the above...`
	`interface Ethernet1/5`
	`  ip router ospf UNDERLAY area 0.0.0.1`
	`  ip ospf bfd`
	`  ip ospf network point-to-point`
	`  no ip ospf passive-interface`
		N9K	Leaf

	`interface loopback<id>` `  ip router ospf UNDERLAY area 0.0.0.1`	N9K	Leaf

OSPF Authentication	`interface eth <slot>/<port>`
	`  ip ospf authentication message-digest`
	`  ip ospf message-digest-key <key-id> md5 0 <clear-text-key>`
		N9K	Leaf

OSPF Reference-Bandwidth	`router ospf UNDERLAY` `  auto-cost reference bandwidth 100Gbps`	N9K	Leaf
Underlay OSPF Configuration on Leaf Underlay OSPF Configuration on Spine	`interface loopback17` `  ip router ospf UNDERLAY area 0.0.0.1`	N9K	Leaf

	`interface eth<slot>/<port>`
	`  ip router ospf UNDERLAY area 0.0.0.1`
	`  ip ospf network point-to-point`
	`  no ip ospf passive-interface`
	`  ip ospf bfd`
	`  ip ospf authentication message-digest`
	`  ip ospf message-digest-key <key-id> md5 0 <clear-text-key>`
		N9K	Leaf / Spine

Enabling Multicast Routing - PIM	`feature pim`	N9K	Leaf
	`ip pim long-neighbor-changes`	N9K	Spine
	`interface ethernet 1/10` `  ip pim sparse-mode`	N9K	Spine
	`interface ethernet 1/10` `  ip pim bfd-instance`	N9K	Spine
	`interface loopback<id>` `  ip pim sparse-mode`	N9K	Leaf

Mapping Layer 2 VNI VXLAN segment to ASM group	interface nve<id>
	  member vni <L2-VNID>
	  mcast-group 239.239.0.1
	  member vni <L2-VNID>
	  mcast-group 239.239.0.2
		N9K	Leaf

PIM Anycast RP (RFC 4610)	interface loopback18 ip pim sparse-mode	N9K	Leaf
	interface loopback17 ip pim sparse-mode	N9K	Leaf
	ip pim rp-address <loopback18> group-list 239.239.0.0/16	N9K	Leaf
Multicast configuration for Leaf	ip pim rp-address <anycast-loopback> group-list 239.239.0.0/16	N9K	Leaf
	feature pim ip pim log-neighbor-changes	N9K	Leaf
	interface loopback17 ip pim sparse-mode	N9K	Leaf

	interface ethernet<slot>/<port>
	  ip pim sparse-mode
	  ip pim bfd-instance
		N9K	Leaf

	interface nve1
	member vni <L2-VNID>
	  mcast-group 239.64.64.1
	member vni <L2-VNID>
	  mcast-group 239.64.64.2
		N9K	Leaf

	ip pim rp-address <anycast-loopback> group-list 239.239.0.0/16	N9K	Leaf
Multicast configuration for Spine	feature pim ip pim log-neighbor-changes	N9K	Spine
	interface ethernet 1/10 ip pim sparse-mode	N9K	Spine
	interface ethernet 1/10 ip pim bfd-instance	N9K	Spine
	interface loopback17 ip pim sparse-mode	N9K	Spine
	interface loopback18 ip pim sparse-mode	N9K	Spine
	ip pim rp-address <loopback18> group-list 239.239.0.0/16	N9K	Spine
	ip pim anycast-rp <loopback18> <loopback17>	N9K	Spine
Service Extensions for OSPF routing	vlan 17 vn-segment 10019	N9K	
	interface Vlan17 mtu 9216 vrf member VRF-MEMBER-1 ip ospf cost 10 ip ospf passive-interface ip router ospf 1 area 0.0.0.0	N9K	

Service Extensions for Static routing	`vrf context VRF-MEMBER-1`
	
	`ip route 0.0.0.0/0 Vlan1605 11.0.23.30`
		N9K	

Service Extension for default route injection on N9K BL/redistribute mode.	`router bgp 65002`
	`vrf VRF-MEMBER-1`
	`address-family ipv4 unicast`
	`network 0.0.0.0/0`
		N9K	

route-map	`route-map RM-IN-S2 permit 10`
	`match tag 1000`
	`route-map RM-IN-S3 permit 10`
	`match tag 1000`
		N9K	

	`route-map RM-S-to-O permit 10`
	`match tag 131 132 133 139 134 135`
	`set metric-type type-1`
		N9K	

vrf context RM-IN-S2	`vrf context RM-IN-S2` `ip route 9.59.207.0/24 Vlan1603 11.0.34.30 name qianyi-henfang tag 1000 50`	N9K	
	`vrf context RM-IN-S2` `ip route 9.59.207.0/24 Ethernet1/46.2 11.0.40.142 name qianyi-henfang tag 1000 10`	N9K	
	`vrf context RM-IN-S2` `ip route 10.0.0.0/12 Vlan1603 11.0.34.30 tag 1000 50`	N9K	
	`vrf context RM-IN-S2` `ip route 10.0.0.0/12 Ethernet1/46.2 11.0.40.142 tag 1000 10`	N9K	
	`vrf context RM-IN-S2` `ip route 10.2.52.0/24 Vlan6 10.2.42.3 tag 1000`	N9K	

	`vrf context RM-IN-S2`
	`ip route 192.168.0.0/16 Vlan1603 11.0.34.30`
	`name beimaluceshiwang tag 1000`
	`rd auto`
	`address-family ipv4 unicast`
	`route-target both auto`
	`route-target both auto evpn`
		N9K	

vrf context RM-OUT	`vrf context RM-OUT` `ip route 10.2.0.0/19 Vlan1607 11.0.34.14 tag 131 50`	N9K	
	`vrf context RM-OUT` `ip route 10.2.0.0/19 Ethernet1/45.1 11.0.40.145 tag 131 10`	N9K	
	`vrf context RM-OUT` `ip route 10.2.96.0/19 Vlan3203 11.0.39.14 tag 134`	N9K	

interface Vlan1601	interface Vlan1601
	no shutdown
	vrf member RM-IN-S3
	no ip redirects
	  ip address 11.0.34.34/28
	no ipv6 redirects
	hsrp version 2
	hsrp 1601
	    preempt
	   priority 110
	   ip 11.0.34.33
		N9K	

interface Vlanxx	interface Vlan1602 no shutdown vrf member VRF-MEMBER-IN-S2 no ip redirects no ipv6 redirects ip ospf cost 10 ip ospf passive-interface ip router ospf 100 area 0.0.0.0	N9K	

interface Ethernet
(IPv4 and IPv6)
	interface ex/y
	  mac aaaa.bbbb.cccc
	  vrf member WSN_LOOP
	  ip address x.x.x.x/31
	  ipv6 address x:x:x::x
	  ip policy route-map TO_VPER_OR_FW
	  ipv6 policy route-map TO_VPER_OR_FW_v6
	  no shut
		N7K	

interface Ethernet1/46.1	interface Ethernet1/36.1 mtu 1500	N9K	
	interface Ethernet1/36.1 encapsulation dot1q 1602 mac-address 0000.0000.2222 vrf member RM-IN-S3 no ip redirects ip address 11.0.40.145/30	N9K	
interface Ethernet1/47.1	interface Ethernet1/37.1 mtu 1500 encapsulation dot1q 1608 vrf member RM-OUT no ip redirects ip address 11.0.72.42/30 ip ospf dead-interval 20 ip ospf hello-interval 5 ip ospf network point-to-point ip router ospf 100 area 0.0.0.0	N9K	

router ospf 1	router ospf 1
	vrf RM-IN-S2
	   router-id 55.2.32.5
	vrf RM-IN-S3
	   router-id 55.2.32.5
	vrf RM-OUT
	   router-id 55.2.32.5
	   redistribute static route-map RM-S-to-O
		N9K	

router bgp 65003	router bgp 65003
	vrf RM-IN-S2
	   address-family ipv4 unicast
	     advertise l2vpn evpn
	     redistribute direct route-map
	vts-subnet-policy
	     redistribute static route-map RM-IN-S2
		N9K	

	nv overlay evpn	N9K	
	clock protocol ntp vdc 1	N9K	

role name nsdcheck	role name nsdcheck rule 4 permit command show * rule 3 permit command terminal length * rule 2 permit command ping * rule 1 permit read	N9K	
	role name devcheck rule 8 permit command tac-pac * rule 7 permit command dir * rule 6 permit command ssh * rule 5 permit command traceroute * rule 4 permit command ping *	N9K	
	role name devopera rule 1 permit read-write	N9K	
ip name-server 55.6.8.73 55.22.8.3	ip name-server 55.6.8.73 55.22.8.3	N9K	
username	username user password 5 1lDuqR.60$eNzZ5I22WxJT58gdEm88N0 role network-operator	N9K	
	username vtsadmin password 5 5MmpswImI$vbZhP/52dNjHY5KWj4yBvmiDvuOZZ9gd2vo2oZc61b4 role network-admin	N9K	
	username nsdcheck password 5 5dpIXMjZs$jDIZVf6grMu1yq79vTts2mcgPlt0QWp5z3tDnw3N5W8 role nsdcheck	N9K	
snmp-server	snmp-server source-interface trap loopback1	N9K	
	snmp-server user user network-operator auth md5 0x3eaa4221f6bbf8722cbdea7ea6bf2f11 priv 0x3eaa4221f6bbf8722cbdea7ea6bf2f11 localizedkey	N9K	
	snmp-server host 55.6.8.1 traps version 2c COMMUNITY1 snmp-server host 55.6.8.1 use-vrf default	N9K	
	snmp-server enable traps bgp snmp-server enable traps ospf snmp-server enable traps callhome event-notify snmp-server enable traps callhome smtp-send-fail snmp-server enable traps cfs state-change-notif snmp-server enable traps lldp lldpRemTablesChange snmp-server enable traps aaa server-state-change snmp-server enable traps hsrp state-change snmp-server enable traps feature-control FeatureOpStatusChange snmp-server enable traps sysmgr cseFailSwCoreNotifyExtended snmp-server enable traps config ccmCLIRunningConfigChanged snmp-server enable traps snmp authentication snmp-server enable traps link cisco-xcvr-mon-status-chg snmp-server enable traps vtp notifs snmp-server enable traps vtp vlancreate snmp-server enable traps vtp vlandelete snmp-server enable traps bridge newroot snmp-server enable traps bridge topologychange snmp-server enable traps stpx inconsistency snmp-server enable traps stpx root-inconsistency snmp-server enable traps system Clock-change-notification snmp-server enable traps feature-control ciscoFeatOpStatusChange	N9K	
	snmp-server community COMMUNITY1 group network-operator	N9K	

ntp	ntp source-interface loopback0 ntp logging	N9K	
ip pim	ip pim ssm range 232.0.0.0/8	N9K	
spanning-tree	spanning-tree pathcost method long spanning-tree mst 1 priority 4096 spanning-tree mst configuration name CFG01 revision 1 instance 1 vlan 1-4094	N9K	
hardware	hardware access-list tcam region qos 0	N9K	
vpc domain	vpc domain 151 peer-keepalive destination 55.2.34.2 source 55.2.34.1 vrf default	N9K	
	vpc domain 151 auto-recovery	N9K	
interface vlan	interface Vlan1602 no shutdown vrf member VRF-MEMBER-IN-S2 no ip redirects fabric forwarding mode anycast-gateway	N9K	
interface port-channel	interface port-channel101 no switchport mtu 9216 no ip redirects ip address 55.2.36.2/30 ip ospf cost 10 ip ospf dead-interval 20 ip ospf hello-interval 5 ip ospf network point-to-point ip router ospf 1 area 0.0.0.0 ip pim sparse-mode	N9K	
interface Ethernet	interface Ethernet1/45 no switchport mtu 9216 mac-address 0000.0000.1111	N9K	
	interface Ethernet1/47 no switchport mtu 9216 udld enable	N9K	
	interface Ethernet2/5 switchport mode trunk switchport trunk allowed vlan 2-4094 channel-group 21 mode active	N9K	

interface mgmt0	`interface mgmt0`
	`no lldp transmit`
	`no lldp receive`
		N9K	

clock timezone	`clock timezone PRC 8 0`	N9K	
ip route	`ip route 0.0.0.0/0 Ethernet1/46.452 55.6.34.198 tag 1000 10` `ip route 0.0.0.0/0 Vlan3903 55.6.40.14 tag 1000 50`	N9K	
router ospf	`router ospf 1` `redistribute static route-map RM-S-to-O`	N9K	

router bgp	`router bgp 65003`
	`router-id 55.2.32.5`
	`address-family ipv4 unicast`
	`address-family l2vpn evpn`
	`neighbor 55.2.32.1`
	`remote-as 65003`
	`update-source loopback1`
	`address-family ipv4 unicast`
	`address-family l2vpn evpn`
	`send-community extended`
		N9K	

router bgp
(IPv4 only)
	`router bgp 65004`
	`router-id 192.168.0.25`
	`log-neighbor-changes`
	`address-family ipv4 unicast`
	`maximum-paths 32`
	`maximum-paths ibgp 32`
	`address-family ipv6 unicast`
	`maximum-paths 32`
	`maximum-paths ibgp 32`
	`address-family l2vpn evpn`
	`neighbor 192.168.0.3`
	`remote-as 65004`
	`password 3 2b7cf4643b66b222`
	`update-source loopback17`
	`address-family l2vpn evpn`
	`send-community`
	`send-community extended`
		N9K and N7K	

Event manager config
(IPv4 and IPv6)
	`event manager applet TRACK-PING-FOR-BGP-DOWN`
	`event track 1 state down`
	`action 1.0 syslog msg CANNOT PING FW. GOING TO SHUTDOWN BGP PEER`
	`action 2.0 cli config term`
	`action 3.0 cli router bgp 65004`
	`action 4.0 cli vrf WSN`
	`action 5.0 cli neighbor 175.175.175.175`
	`action 6.0 cli shutdown`
	`event manager applet TRACK-PING-FOR-BGP-UP`
	`event track 1 state up`
	`action 1.0 syslog msg CAN PING FW. GOING TO NO SHUTDOWN BGP PEER`
	`action 2.0 cli config term`
	`action 3.0 cli router bgp 65004`
	`action 4.0 cli vrf WSN`
	`action 5.0 cli neighbor 175.175.175.175`
	`action 6.0 cli no shutdown`
		N9K	

IP sla config (IPv4 only for N9K) (IPv4 and IPv6 for N7K)		N9K and N7K	

	!On BL-1 Track the local VPER-1
	ip sla 1
	  icmp-echo 69.83.32.36 source-interface vlan 2400
	
	  vrf WSN_LOOP ☐ forward reference to VRF
	  threshold 500
	  timeout 500
	  frequency 1
	! Start the SLAs
	ip sla schedule 1 life forever start-time now
	
	! Setup a track object for sla 1
	track 1 ip sla 1 reachability
	 delay up 180 down 3
	
	! Set up a track open that returns a DOWN only if
	 both objects 1 and 2 are down.
	track 111 list boolean or
	  object 1		

Track config (IPv4 and IPv6)	track 10 ip route 0.0.0.0/0 reachability vrf member WSN	N9K and N7K	

Interface port channel
(IPv4 and IPv6)
	interface port-channel 110.2511
	  encapsulation dot1q 2511
	  vrf member WSN_LOOP
	  ip address 69.83.28.32/31
	  no shut
	interface port-channel 110.2575
	  encapsulation dot1q 2575
	  vrf member WSN_LOOP
	  ipv6 address 2001:4888:16:2078:1e1:2a1::/64
		N9K and N7K	

	interface port-channel 110.2577
	  ip policy route-map FROM_VPER
	
	interface port-channel 110.2577
	  ipv6 policy route-map FROM_VPERv6
	! EEM to track both VPERs, when one is up restore
	 traffic
	event manager applet VPER_TRACK_UP
	  event track 111 state up
	  action 1.0 syslog msg "BOTH VPERS ARE UP.
	REMOVING BYPASS!"
	  action 2.0 cli command "config t"
	  action 3.0 cli command "route-map TO_VPER_OR_FW
	 permit 20"
	  action 4.0 cli command "no continue 30"
	  action 5.0 cli command "exit"
	  action 6.0 cli command "route-map TO_VPER_OR_FWv6
	 permit 20"
	  action 7.0 cli command "no continue 30"
	  action 8.0 cli command "exit"
	  action 9.0 cli command "route-map
	FROM_FW_TO_VPER_OR_MOBILE permit 10"
	  action 10.0 cli command "no continue 20"
	  action 11.0 cli command "end"
	  action 12.0 cli command "route-map
	FROM_FW_TO_VPER_OR_MOBILEv6 permit 10"
	  action 13.0 cli command "no continue 20"
	  action 14.0 cli command "end"
	  action 15.0 syslog msg "TRAFFIC HAS BEEN RESTORED
	 TO VPER"
		N7K	

		N9K	

	! EEM to track both VPERs, when one is up restore traffic
	event manager applet VPER_TRACK_UP
	  event track 111 state up
	  action 1.0 syslog msg BOTH VPERS ARE UP. REMOVING BYPASS
	  action 2.0 cli command "config t"
	  action 3.0 cli command "route-map TO_VPER_OR_FW permit 20"
	  action 4.0 cli command "no continue 30"
	  action 5.0 cli command "exit"
	  action 6.0 cli command "route-map TO_VPER_OR_FWv6 permit 20"
	  action 7.0 cli command "no continue 30"
	  action 8.0 cli command "exit"
	  action 9.0 cli command "route-map FROM_FW_TO_VPER_OR_MOBILE permit 10"
	  action 10.0 cli command "no continue 20"
	  action 11.0 cli command "end"
	  action 12.0 cli command "route-map FROM_FW_TO_VPER_OR_MOBILEv6 permit 10"
	  action 13.0 cli command "no continue 20"
	  action 14.0 cli command "end"
	  action 15.0 syslog msg TRAFFIC HAS BEEN RESTORED TO VPER		

IP access list
(IPv4 and IPv6)
	ip access-list ALL_POOLS
	  10 permit ip 1.0.0.0/8 any
	  20 permit ip any 1.0.0.0/8
	  30 permit ip 2.0.0.0/8 any
	  40 permit ip any 2.0.0.0/8
	
	! Need to configure a ACL for all All POOLS
	ipv6 access-list ALL_POOLSv6
	  10 permit ipv6 2001:1::/32 any
	  20 permit ipv6 any 2001:1::/32
	  30 permit ipv6 2001:2::/32 any
	  40 permit ipv6 any 2001:2::/32
		N9K and N7K	

Route-map
(IPv4 and IPv6)
	set ip next-hop verify-availability 69.83.32.35 track 2
	route-map TO_VPER_OR_FW permit 30
	  match ip address ALL_POOLS
	  ! Set the ip next-hop to the FW VIP
	  set ip next-hop 69.83.136.129
	
	route-map TO_VPER_OR_FW_v6 permit 10
	! Leave room here for the pilot packets
	route-map TO_VPER_OR_FW_v6 permit 20
	  match ipv6 address VPER_POOLSv6
	  set ipv6 next-hop verify-availability 2001:4888:16:2078:1e1:210:: track 1
	  set ipv6 next-hop verify-availability 2001:4888:16:207a:1e1:210:: track 2
	route-map TO_VPER_OR_FW_v6 permit 30
	  match ipv6 address ALL_POOLSv6
	  ! Set the ipv6 next-hop to the FW VIP
	  set ipv6 next-hop 2001:4888:39:3080:308:25::
		N9K and N7K	

	route-map FROM_FW_TO_VPER_OR_MOBILE permit 10
	  match ip address VPER_POOLS
	  set vrf WSN_VPER
		N7K	

Monitor erspan	monitor session 1 type erspan-source
	  erspan-id 5
	  vrf TEST
	  ip ttl 25
	  ip dscp 42
	
	monitor erspan origin ip-address 10.0.0.1 global
		N9K and N7K	

QOS- class-map	class-map type qos match-any TEST1 match packet length 5	N9K and N7K	

QOS class-map policy-map	class-map type control-plane match-any vzw-copp-system-p-class-exception
	  match exception ip option
	  match exception ip icmp unreachable
	  match exception ipv6 option
	  match exception ipv6 icmp unreachable
	class-map type control-plane match-any vzw-copp-system-p-class-fcoe
	  match access-group name vzw-copp-system-p-acl-mac-fcoe
	
	policy-map type control-plane vzw-copp-system-p-policy-strict
	  class vzw-copp-system-p-class-exception
	  set cos 1
	  police cir 360 kbps bc 250 ms conform transmit violate drop
	  class vzw-copp-system-p-class-fcoe
	  set cos 6
	  police cir 1060 kbps bc 1000 ms conform transmit violate drop		

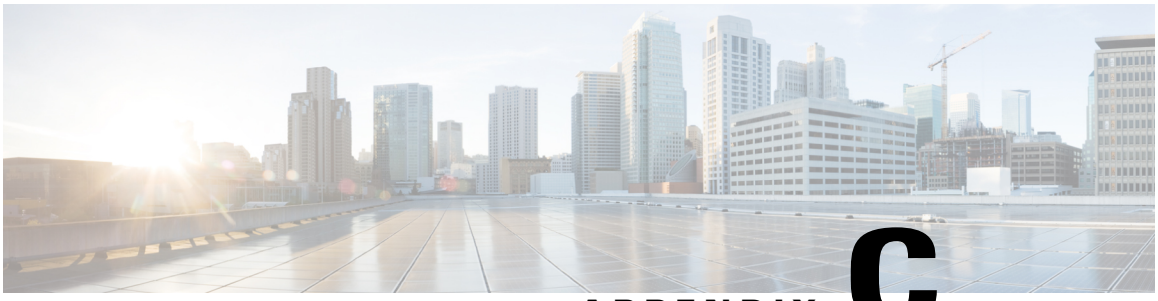

APPENDIX C

OpenStack Configuration for SR-IOV Support

You need to follow RedHat OpenStack documentation for detailed instructions for SR-IOV configuration. The details given below may be referred to in addition to this.

On Controller Node

1 In /etc/neutron/plugins/ml2/ml2_conf.ini file

```
[ml2]

tenant_network_types = vxlan, vlan

# Add support for vlan, vxlan and flat type drivers

type_drivers = vlan, vxlan, flat

# Add support for sriov and vts mechanism drivers, in the following sequence

mechanism_drivers = sriovnicswitch, cisco_vts
```

2 In /etc/neutron/plugins/ml2/ml2_conf.ini file

```
 [ml2_type_vlan]

# Add all the physnet names that will be used on your compute hosts along with the VLAN
 ranges

# reserved for the provider networks for those physnets

network_vlan_ranges = physnet1:2000:2100, physnet2:2500:2600

# Note: network vlan range should be in between the range specified in VTS.


ml2_type_flat]

# List of physical_network names with which flat networks can be created.

# Use default '*' to allow flat networks with arbitrary physical_network names.

flat_networks = *
```

3 If your deployment has a /etc/neutron/plugins/ml2/ml2_conf_sriov.ini file, include the following section in your ml2_conf_sriov.ini file.

```
[ml2_sriov]

# The default supported_pci_vendor_devs value for the installation may have the value
only for the PFs

# in your compute. Use lspci —nn | grep Ethernet to find the ids for the Virtual
functions and add that
```

```
# as well in this.

supported_pci_vendor_devs = 8086:10fb,8086:10ed

agent_required = True
```

Update ExecStart section in the file /usr/lib/systemd/system/neutron-server.service to include ml2_conf_sriov.ini config file

```
ExecStart=/usr/bin/neutron-server --config-file /usr/share/neutron/neutron-dist.conf
--config-dir /usr/share/neutron/server --config-file /etc/neutron/neutron.conf
--config-file /etc/neutron/plugin.ini ––config-file
/etc/neutron/plugins/ml2/ml2_conf_sriov.ini --config-dir /etc/neutron/conf.d/common
--config-dir /etc/neutron/conf.d/neutron-server --log-file /var/log/neutron/server.lo
```

4. If the ml2_conf_sriov.inif file is not present, then add the ml2_sriov section to /etc/neutron/plugins/ml2/ml2_conf.ini file.
5. Restart neutron service using "systemctl restart neutron-server.service"

On each compute hosts' /etc/nova/nova.conf file, define pci_passthroughs_whitelist.

```
pci_passthrough_whitelist = [ {"devname": “eth4", "physical_network":"physnet1"}, {"devname":
 “eth5", "physical_network":"physnet2"}]
```

For multiple SRIOV nics , there should be a mapping entry per physnet/NIC card

Enable the OpenStack Networking SR-IOV agent

If not present already as part of the OSPD/Packstack installation (check the neutron agent-list on your director/controller node), then you will need to install "openstack-neutron-sriov-nic-agent" on your compute hosts and start that agent/service. After you have installed, follow RedHat documentation. It is important to have the "physical_device_mappings section" in the /etc/neutron/plugins/ml2/sriov_agent.ini file. You can leave the exclude_devices section blank.

Then proceed to create the SR-IOV port instances.

Note

- SRIOV support is present only for network types vlan and flat.
- While creating a provider network, when choosing vlan network type, choose a segmentation ID (vlan) from within the physnet range.
- While creating a provider network of type flat, segmentation id can be any value or null. VTS will provide native vlan for flat networks.
- While creating a port, as described in RedHat OpenStack documentation, choosing binding:vnic_type = direct is a must for SRIOV.
- Cirros image is not supported for creating instances using SRIOV ports.