# Cisco Virtual Topology System (VTS) 2.4.2 User Guide

**First Published:** 2017-02-19

**Last Modified:** 2017-02-21

**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 2

Cisco Virtual Topology Forwarder 3

Virtual Topology System High Availability 3

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 8

Inventory 9

Using the Expert Mode—Quick Guide 10

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 19

Performing Auto Discovery 20

Importing Inventory using CSV File 22

Viewing the Network Topology 23

Viewing Network Inventory 24
Adding Fabric Connection 25
Synchronizing Configuration 25
Viewing Host Inventory 26
Adding a new Host 27
Installing OpenStack Host Agent and VTF using GUI 28
Viewing the IOS XRv to VTF Mapping 29
Migrating from vPC to ESI 29
Redeploying Device Inventory 32
Enabling Static Multi Homing 33
Enabling Static Multi Homing on Cisco Nexus 7000 33
Enabling Static Multi Homing on Cisco Nexus 9000 33

CHAPTER 5 Managing Resources 35
Specifying Global VNI Range 36
Specifying VLAN Range 36
Specifying Device VLAN Range 37
Specifying Group VLAN Range 38
Specifying Multicast IP Pool 38
Specifying Interface VLAN Range 39
Creating Interface Groups 39

CHAPTER 6 Creating and Managing Admin Domains 43
Admin Domain Overview 43
Viewing Admin Domain 44
Creating an Admin Domain 44
Creating DCI Interconnect Profiles 48

CHAPTER 7 Managing Templates 51
Creating Route Templates 53
Adding Static Routes 53
Adding Route Targets 53
Adding Fabric Internal Route Targets 54
Adding Fabric External Route Targets 54
Disabling Auto Route Target Configuration 55

Creating L3 Extension Templates 56
Supported System Variables 56
Editing Templates 57
Copying a Template 58
Deleting Templates 58
Attaching Templates to Tenants 58
Attaching Template while Adding Tenants 59
Attaching Template while Editing a Tenant 59
Detaching Template from Tenants 59
Attaching Templates to Routers 60
Attaching Templates while Adding Routers 60
Attaching Template while Editing a Router 61
Creating Underlay Templates 61
Attaching Devices to Underlay Template 62

CHAPTER 8 Managing Tenants 63
Adding Tenants 63
Editing Tenants 64

CHAPTER 9 Provisioning Overlay Networks 65
Provisioning Overlay Networks Using Cisco Virtual Topology System 65
Creating Overlays 66
Using OpenStack 67
Using VMware 67
Using Cisco VTS GUI 67
Creating Network using VMware 67
Creating Subnetwork using VMware 68
Creating Routers using VMware 68
Attaching Network to Router 68
Attaching a Virtual Machine to Network 69
Creating a Network using Cisco VTS GUI 69
Creating a Subnetwork 70
Attaching a Port 70
Creating Router using Cisco VTS GUI 71
Assigning BVI Interface IP Address 72

Extending Layer 2 Network Across Data Centers 73
Enabling Global Route Leaking Service 73

CHAPTER 10 Viewing Overlay Details 77
Viewing Device Details 77
Viewing Virtual Machine Details 78
Viewing Baremetal Details 78
Viewing Network Details 79
Viewing Router Details 80

CHAPTER 11 Administering Cisco VTS 81
Setting up System 81
Modifying Login Banners 82
Setting Global Route Reflector 82
Registering the Virtual Machine Manager using GUI 82
Backing up the Database 83
Restoring the Database 84
Viewing HA Status 85
Enabling External Authentication and Authorization 85
Setting up Remote Authentication Server 86
Enabling Accounting and Logging 87
Setting up Accounting 87

APPENDIX A Service Extension Templates-Supported Configuration 89
Supported Service Extension Template Configuration for Cisco Nexus 7000 Series Switches 89
Supported Service Extension Template Configuration for Cisco ASR 9000 Series Routers 90

APPENDIX B Supported Underlay Configuration 93
Supported Underlay Configuration 93

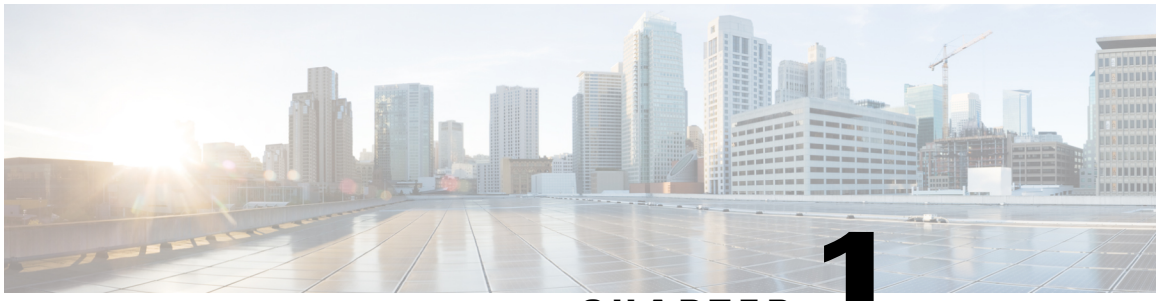

CHAPTER 1

# 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 2
- Cisco Virtual Topology Forwarder, page 3
- Virtual Topology System High Availability, page 3


## 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 2.3.1 Installation Guide.
- Software forwarder (Virtual Topology Forwarder [VTF])

# 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.

***Figure 1: Cisco VTS Architecture***

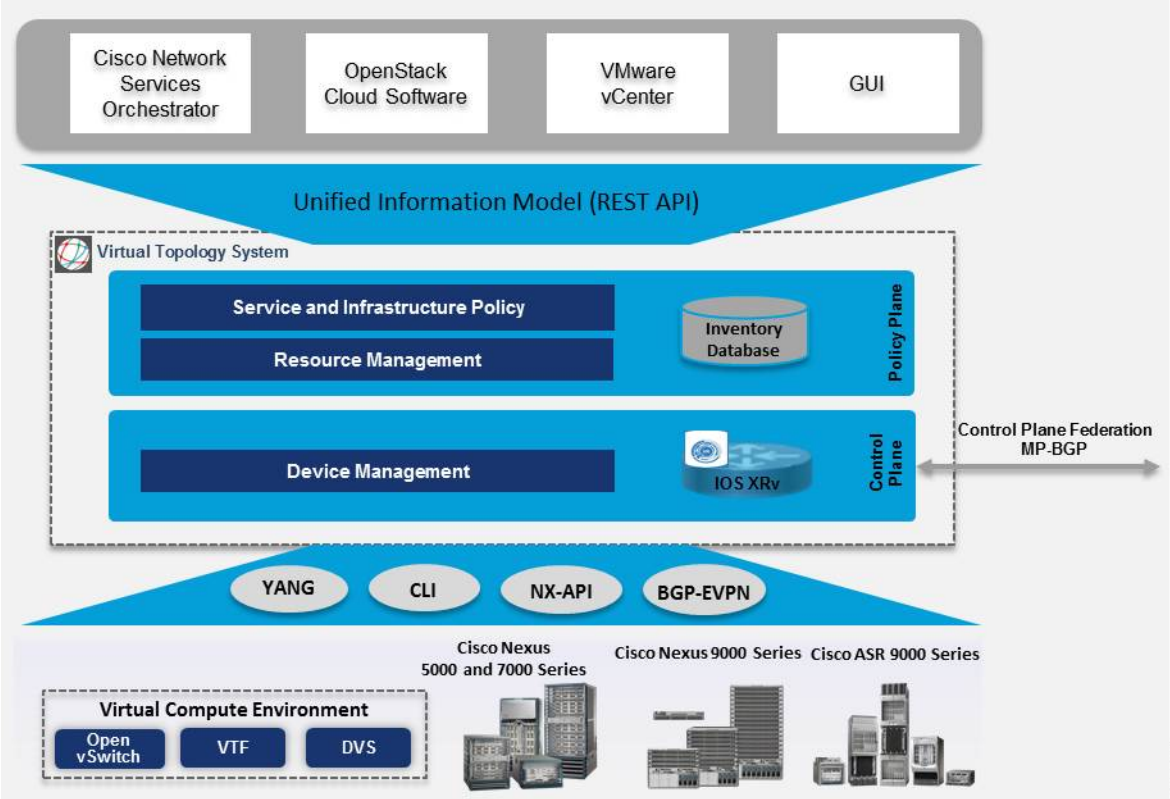


- **Policy Plane**: The policy plane enables Cisco VTS to implement a declarative policy model designed to capture user intent and render it into specific device-level constructs. The solution exposes a set of modular policy constructs that can be flexibly organized into user-defined services for use 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 the Cisco IOS XRv Software instance that provides route peering capabilities between the DCI gateways or to a BGP route reflector. Cisco IOS XRv is the virtualized version of Cisco IOS XR Software. The control plane enables an MP-BGP EVPN-based control plane for VXLAN overlays originating from leafs or software VXLAN tunnel endpoints (VTEPs).

  The device management module enables device configuration and management capabilities within Virtual Topology System, with multiprotocol support to support a multivendor environment.

# Cisco Virtual Topology Forwarder

Cisco VTS solution provides a l2/l3 software switch that can act as a software VTEP and is called Virtual Topology forwarder (VTF). Cisco VTS can be deployed with a Virtual Topology Forwarder (VTF). VTF is a lightweight, multitenant software data plane designed for high performance packet processing on x86 servers. VTF uses Vector Packet Processing (VPP). VPP is a full-featured networking stack with a software forwarding engine. VTF leverages VPP technology and Intel Data Path Development Kit (DPDK) for high performance Layer 2 (L2), Layer 3 (L3), and VXLAN packet forwarding.

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 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.

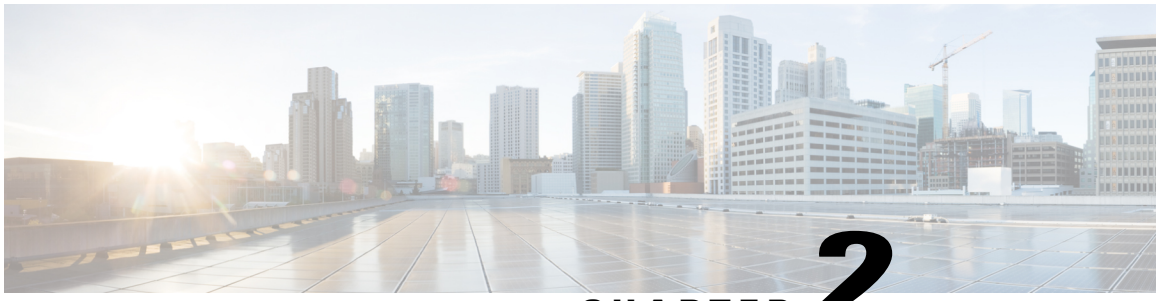

CHAPTER 2

# 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 10
- 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.

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

**Step 2** Enter the following details:

- VMM Type—Specify the VMM type. Choose vcenter or openstack from the drop-down list.
- Version Name—Specify the version details. See Note below for details.
- Description—Enter a description for the VMM.
- IP Address-Port—If you choose vcenter, enter the IP address and the port. The default port is 443. If you choose openstack, enter the IP address.
- User Name—Enter the VMM username.
- Passphrase—Enter the VMM passphrase.

**Note**

- **1** If VMM Type is **openstack** and Version Name is **openstack-liberty-centos**, a new field VTF Deployment Mode appears in the Add VMM popup.
- **2** VTF Deployment Mode field had two modes: **vtf-vhost** and**vtf-as-vm**.
- **3** openstack-liberty-rhel, openstack-icehouse, openstack-juno, and openstack-kilo support only VM mode.
- **4** Openstack-liberty-centos supports both vhost and VM modes.
- **5** In vHost mode VTF installation has to be done manually.
- **6** VTF deployment mode can be selected at the time of VMM registration and this cannot be changed later.
- **7** All the computes in a data center should be either in vhost mode (VTF running as a process) or in vNode mode (VTF running as a VM). Currently VTS does not support mixed mode.

**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.
The Authorization Group screen is displayed.

# 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 IoS XRv. 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.

  The Host Inventory has the **Install Capabilities** button. You can select the host from the list, and click **Install Capabilities**. If the host is connected to a physical VTEP (ToR), it installs the host agent in case of an OpenStack environment. If the capability is Virtual, clicking Install Capabilities installs the VTF VM on that host. See Installing OpenStack Host Agent and VTF using GUI, on page 28.

  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<br>• Virtual Machine Manager | • Setting up System, on page 81<br>• Registering the Virtual Machine Manager using GUI, on page 82 |
| Set up Inventory | • Discover Devices | • Managing Inventory, on page 19 |
| Admin Domains | Create Admin Domains | Creating an Admin Domain, on page 44 |
| Set up Tenants | Add Tenants | Adding Tenants, on page 63 |
| Set up Overlay | • Add Network<br>• Define Baremetal<br>• Add Virtual Machines<br>• Add Routers | Provisioning Overlay Networks, on page 65 |

# 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<br>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. |

| Sequence | Task | Navigation in VTS GUI / User Guide Section | Additional Notes |
|---|---|---|---|
| 2 | Discover the Topology of all the leafs, spine, border-leafs, and DCI | Discovery > Topology Discovery<br>For more information about adding devices and host information, see Performing Auto Discovery, on page 20 | VTFs and IOS XRv are not detected in topology discovery. |
| 3 | Import the devices after adding the auth group | Inventory > Import Inventory<br>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<br>For more information about adding devices and host information, see Administering Cisco VTS, on page 81. | |
| 5 | Perform IOS XRv and VTF Registration.<br>**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*. | IOS XRv takes a few minutes to boot. Register the VTF only after you see the VFG group on the GUI.<br>To verify that the VFG group is created, go to Inventory > Virtual Forwarding Groups. |

| Sequence | Task | Navigation in VTS GUI / User Guide Section | Additional Notes |
|---|---|---|---|
| 6 | Login into the IOS XRv and verify that the network-controller config is pushed<br>**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<br>For more information, see Viewing Network Inventory, on page 24. | |
| | Create resource pools | • Resource Pools > Global VNI Pool<br>• Resource Pools > Device Specific VLAN Pools<br>• Resource Pools > Multicast IP Pool<br>For more information about creating an admin domain, see Managing Resources, on page 35 | |

| Sequence | Task | Navigation in VTS GUI / User Guide Section | Additional Notes |
|---|---|---|---|
| 8 | Create an Admin Domain | Admin Domains > Domains<br>For more information about creating an admin domain, see Creating an Admin Domain, on page 44 | Properties for the L2/L3 Gateway Group are as follows:<br>• Control Protocol: BGP-EVPN<br>• Replication Modes : Multicast and Ingress<br>• Distribution Mode: Decentralized |
| 9 | Add the devices to the Gateway Group | Admin Domains > Domains<br>For more information about creating an admin domain, see Creating an Admin Domain, on page 44 | See the *Supported Platforms* section in the *Cisco VTS Installation Guide* for details about devices support for different roles. |
| 10 | Add the ToR and IOS XRv to the L2 and L3 Gateway Group | Admin Domains > Domains<br>For more information about creating an admin domain, see Creating an Admin Domain, on page 44 | |

| Sequence | Task | Navigation in VTS GUI / User Guide Section | Additional Notes |
|---|---|---|---|
| 11 | Save the Admin Domain you created | Admin Domains > Domains<br>For more information about creating an admin domain, see Creating an Admin Domain, on page 44 | |

# 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.

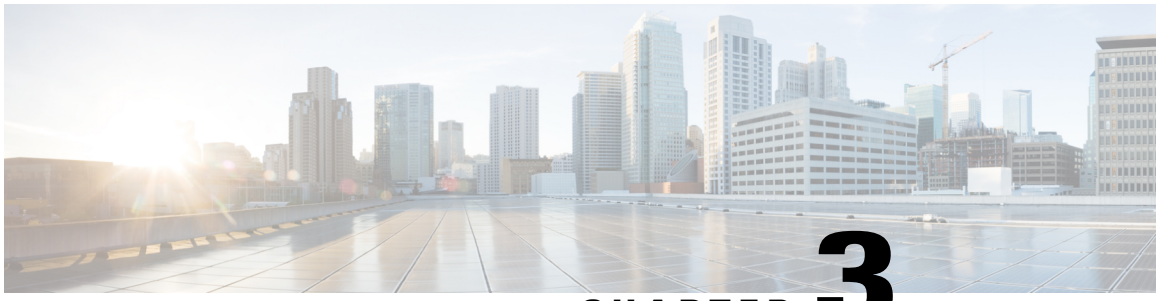

CHAPTER 3

# 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.

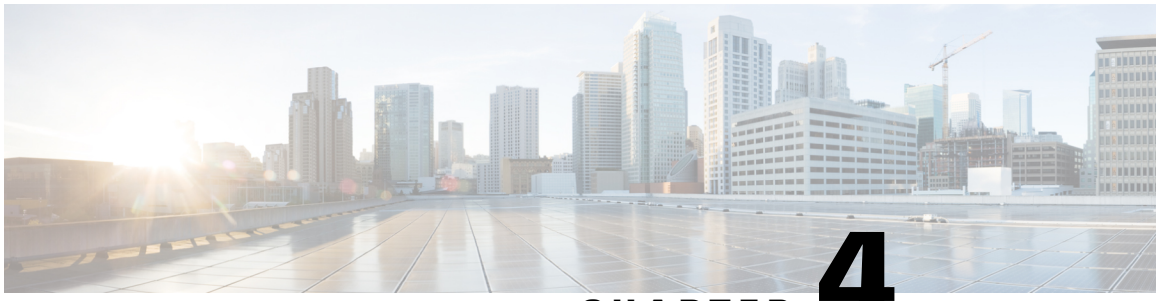

CHAPTER 4

# 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 19
- Performing Auto Discovery, page 20
- Importing Inventory using CSV File, page 22
- Viewing the Network Topology, page 23
- Viewing Network Inventory, page 24
- Viewing Host Inventory, page 26
- Viewing the IOS XRv to VTF Mapping, page 29
- Migrating from vPC to ESI, page 29
- Redeploying Device Inventory, page 32
- Enabling Static Multi Homing, page 33


# 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.
- 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.

**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.

**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.

# 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 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.

Different types of links will be represented as shown in legends.

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

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

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 IoS XRv. 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.

The Host Inventory has the **Install Capabilities** button. You can select the host from the list, and click **Install Capabilities**. If the host is connected to a physical VTEP (ToR), it installs the host agent in case of an OpenStack environment. If the capability is Virtual, clicking the Install Capabilities button installs the VTF VM on that host. Repeat the install capabilities procedure for each of the host/device in the list on Host Inventory. See Installing OpenStack Host Agent and VTF using GUI, on page 28.

To view host inventory details:

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

**Step 2** Select the device from the Select Device drop-down list. The following details are displayed:

- Host Name
- Host Type
- IPV4 Address
- IPV6 Address —Applicable only for baremetal.
- Device Port Name
- Installation Status—Shows the installation status.

You can click the info icon against each host to view additional host details like Host Interface. User Name, Connection ID, VMM Type, and Capability.

**Note** If you have converted a virtual-server to baremetal, before you upgrade, you will not be seeing details like Installation Status and Capabilities, after the upgrade.

To edit a host from the table, select the Host Name check box corresponding to the device and click the **Edit** icon.

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.

# Adding a new Host

To add a new host:

**Step 1** Click the **Add (+)** icon. The Add New Host popup window appears.

**Step 2** Enter the following details:

- Host Name—This is mandatory.
- Host Type—baremetal or virtual-server.
- Host Interface—This is mandatory.
- Host IPV4 Address
- Host IPV6 Address
- Device Port Name
- Capability—Applicable only if host type is virtual-switch. The host capability - no-virtual switch or virtual-switch.
- User
- Passphrase
- VTF Details—These are displayed only if the host capability is virtual-switch.
  - VTF IP—The IP address of the VTF.
  - VTF Name
  - underlay VLAN ID
  - Subnet Mask
  - Gateway
  - Underlay Bridge—Name of the underlay network portgroup/bridge on the binding-host to which VTF is attached.
  - Internal Bridge—Name of the tenant network port group/bridge on the binding-host to which VTF is attached.
  - User

- Passphrase
- Confirm Passphrase

**Step 3** Click Save

# Installing OpenStack Host Agent and VTF using GUI

You can use the Install Capabilities button to install the Host Agent and VTF on a host based on whether it is a virtual or physical server.

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

**Step 2** Click + to add a host. You may also edit a host and modify the parameters to enable installation of physical or virtual capabilities.

- If you click + (Add) the Add New Host popup is displayed. Enter the following details.
- If you choose to edit an existing host, the following windows are displayed depending upon the host type:
  - Host Details:
    - Host Name
    - Host Type—baremetal or virtual-server.
    - Host Interface
    - Host IP Address—
    - Device Port Name
    - Capability—Specify whether it is a virtual-switch or not.
    - Vmm IP Address
    - User
    - Passphrase
  - VTF Details—These are displayed only if the host capability is virtual-switch.
    - VTF IP—The IP address of the VTF.
    - Subnet Mask
    - Gateway
    - Underlay Bridge—Name of the underlay network portgroup/bridge on the binding-host to which VTF is attached.
    - Internal Bridge—Name of the tenant network port group/bridge on the binding-host to which VTF is attached.

- User
- Passphrase

Ensure that you review the tooltips for important information about the entries.

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.

**Step 3** Click **Install Capabilities**. Based on the host type, it installs the host agent / VTF on the host.
See the Status column for the installation status. The VMM type is also shown once the capabilities get installed.

**Note** While installing capabilities for a host, if the selected VMM is **openstack-liberty-centos** and the mode for VTF deployment is **vtf-vhost**, VTF installation has to be done manually.

# Viewing the IOS XRv 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 the IOS XRvs . The table on the right hand side shows the VTFs.

# 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 IOS XRv HA, bring down the IOS XRv.

**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`<br>`  no ip address 44.44.44.44/32 secondary` |
| Enable ESI | `evpn esi multihoming` |
| Create nve | `interface nve1`<br>`  no shutdown`<br>`  source-interface loopback0`<br>`  host-reachability protocol bgp` |
| Enable core links | `interface Ethernet1/35`<br>`  Description " Connected with Spine"`<br>`  no switchport`<br>`  evpn multihoming core-tracking  <<< Add here`<br>`  ip address 16.1.1.2/24`<br>`  ip router ospf 100 area 0.0.0.0`<br>`  ip pim sparse-mode`<br>`  no shutdown` |
| Add Ethernet-segment and system-mac address in the port-channel | `interface port-channel220`<br>`  switchport mode trunk`<br>`  switchport trunk allowed vlan none`<br>`  ethernet-segment 220`<br>`  system-mac eeee.1111.2222` |
| Apply the channel group to the TORs interface which are connected to compute. | `interface Ethernet1/5`<br>`  switchport trunk allowed vlan none`<br>`  channel-group 220 mode active` |
| Verify whether the ESI is up. | `tor1# show nve ethernet-segment`<br><br>`ESI Database`<br>`----------------------------------------`<br>`ESI: 03aa.bbcc.ddee.ee00.002d,`<br>`  Parent interface: port-channel30,`<br>`  ES State: Up`<br>`  Port-channel state: U`<br>`  NVE Interface: nve1`<br>`  NVE State: Up`<br>`  Host Learning Mode: control-plane`<br>`  Active Vlans: 1001`<br>`  DF Vlans: 0-4095`<br>`  Active VNIs: 30001`<br>`  Number of ES  members: 1`<br>`  My ordinal: 0`<br>`  DF timer start time: 00:00:00`<br>`  Config State: config-applied`<br>`  DF List: 1.1.1.1`<br>`  ES route added to L2RIB: True`<br>`  EAD routes added to L2RIB: True`<br>`----------------------------------------` |

**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 32 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 IOS XRv 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 25 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.

# 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.

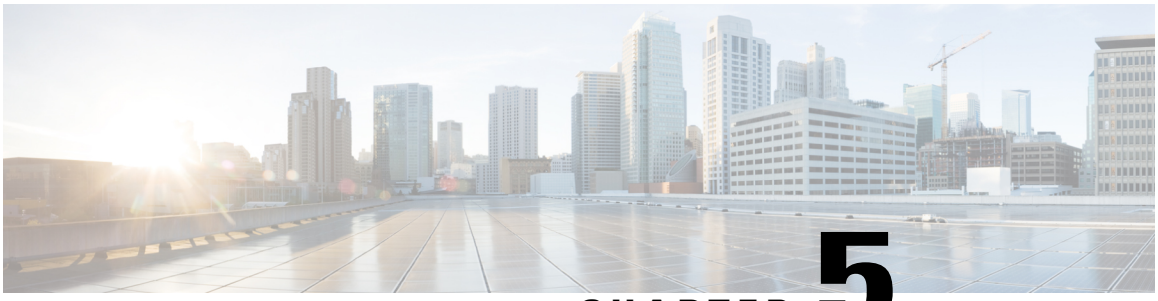

CHAPTER 5

# 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 36
- Specifying VLAN Range, page 36
- Specifying Multicast IP Pool, page 38
- Specifying Interface VLAN Range, page 39
- Creating Interface Groups, page 39

# 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.

**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 37
- Specifying Group VLAN Range, on page 38
- Specifying Interface VLAN Range, on page 39
- Creating Interface Groups, on page 39


# 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)**.

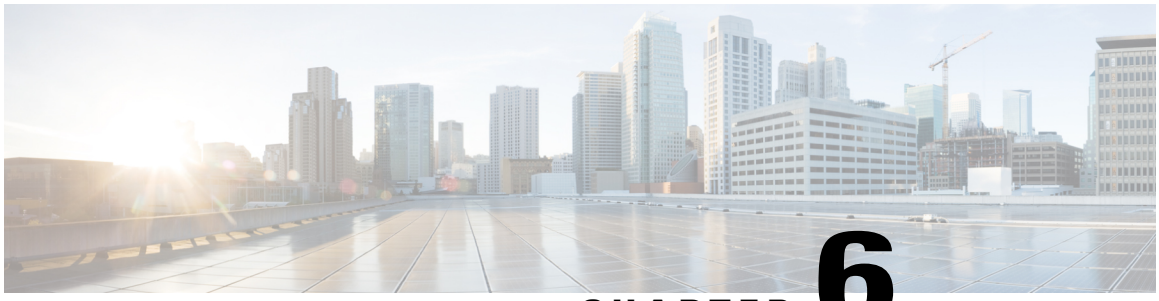

CHAPTER 6

# Creating and Managing Admin Domains

This chapter has the following sections:


- Admin Domain Overview, page 43
- Viewing Admin Domain, page 44
- Creating an Admin Domain, page 44
- Creating DCI Interconnect Profiles, page 48


## 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.<br>**Note** If it is a DCI mode, then you need to add DCI device to both the DC GW and DCI.<br>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.<br>**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.<br>**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.

| Functional Group | Property |
|---|---|
| DCI | Specify:<br>• Whether it is a New or Shared DCI.<br>• The Redundancy / Availability settings:<br>  • Enable/Disable Redundancy using the toggle switch.<br>  • ICCP—VXLAN/fabric ICCP group number. Valid range is 1 to 4294967295. MPLS/core ICCP group number. Valid range is 1 to 4294967295<br>  • 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.<br><br>Click Stitching Profile and choose the required profile. |
| DC GW | Specify:<br>• Whether it is a New or Shared DC GW.<br>• The Control Protocol—BGP EVPN. |

| Functional Group | Property |
|---|---|
| L3 GW | Specify:<br>• Whether it is a New and Shared L3 GW.<br>• The Control Protocol—BGP EVPN.<br>• 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.<br>**Note** • Cisco Nexus 5600 supports Multicast replication mode only.<br>• VTF supports Ingress mode only.<br>• Cisco Nexus 9000 supports both modes.<br>• Distribution Mode—Decentralized.<br>**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.<br>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.<br>• ARP Suppression—On. |
| L2 GW | Specify:<br>• Whether it is a New and Shared L2 GW.<br>• The Control Protocol—BGP EVPN.<br>• The Replication Mode—Multicast or Ingress.<br>• The Distribution Mode—Decentralized.<br>• 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 73 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**.

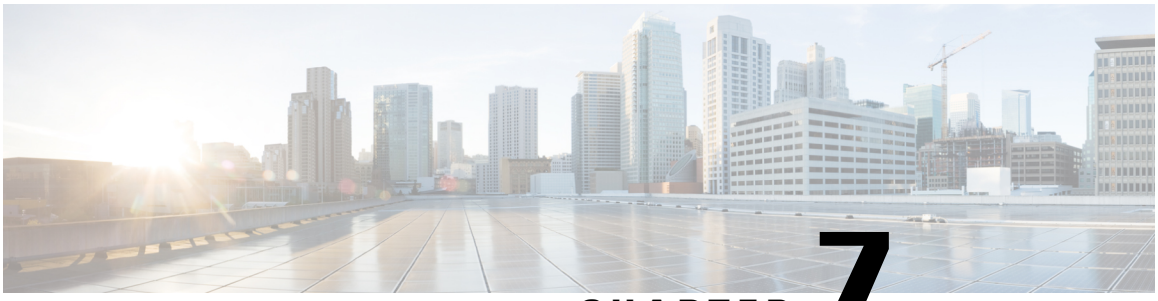

CHAPTER 7

# 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.

- **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 53
- Creating L3 Extension Templates, page 56
- Editing Templates, page 57
- Copying a Template, page 58
- Deleting Templates, page 58
- Attaching Templates to Tenants, page 58
- Attaching Templates to Routers, page 60
- Creating Underlay Templates, page 61

# 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 55 for details.

**Step 9** Add Static Routes. See Adding Static Routes, on page 53 for details.

**Step 10** Add Route Targets. See Adding Route Targets, on page 53 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 56 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 9000 Series

- Tenant name
- VRF name
- VLAN ID

Supported variables in Cisco Nexus 7000 Series

-
-

Supported variables in Cisco ASR 9000 Series

- VRF Name
- Bridge Group Name
- Bridge Domain Name
- BGP No Instance Name
- Subinterface ID

# 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 59
- Attaching Template while Editing a Tenant, on page 59

# 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.
If you want to create a new template and attach that, click **Create 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 60
- Attaching Template while Editing a Router , on page 61

**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 69 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 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.

# Attaching 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.

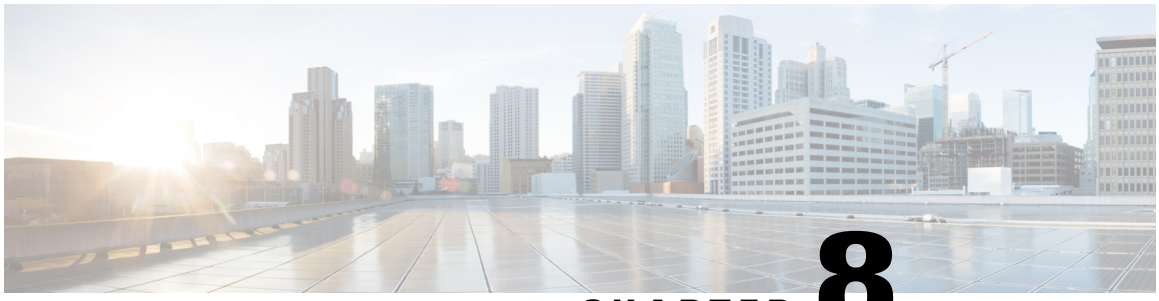

CHAPTER 8

# 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:


- Adding Tenants, page 63
- Editing Tenants, page 64


## 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 58 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 58 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**.

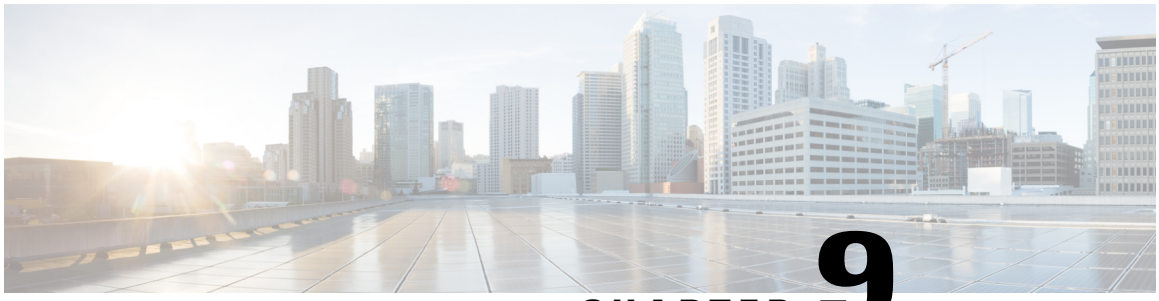

CHAPTER 9

# Provisioning Overlay Networks

This chapter has the following sections:


- Provisioning Overlay Networks Using Cisco Virtual Topology System, page 65
- Creating Overlays, page 66
- Creating Network using VMware, page 67
- Creating Subnetwork using VMware, page 68
- Creating Routers using VMware, page 68
- Attaching Network to Router, page 68
- Attaching a Virtual Machine to Network, page 69
- Creating a Network using Cisco VTS GUI, page 69
- Creating Router using Cisco VTS GUI, page 71
- Assigning BVI Interface IP Address, page 72
- Extending Layer 2 Network Across Data Centers, page 73
- Enabling Global Route Leaking Service, page 73


# 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, Cisco IOS XRv route reflectors, 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.

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 68
- Creating Network using VMware, on page 67
- Creating Subnetwork using VMware, on page 68
- Creating Routers using VMware, on page 68
- Attaching a Virtual Machine to Network, on page 69

## 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 69
- Creating Router using Cisco VTS GUI, on page 71

# 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. See Extending Layer 2 Network Across Data Centers, on page 73. 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.

**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 60 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 Cisco IOS XRv 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 IOS XRv configuration, a similar error is displayed if the default VRF name exceeds 24 characters.
- 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.

**Step 10** Click **Add (+)** icon. The Add Interface popup appears.

**Step 11** Select the subnet from the drop down list, and click **OK**.

**Step 12** 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 25 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 48.

**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.

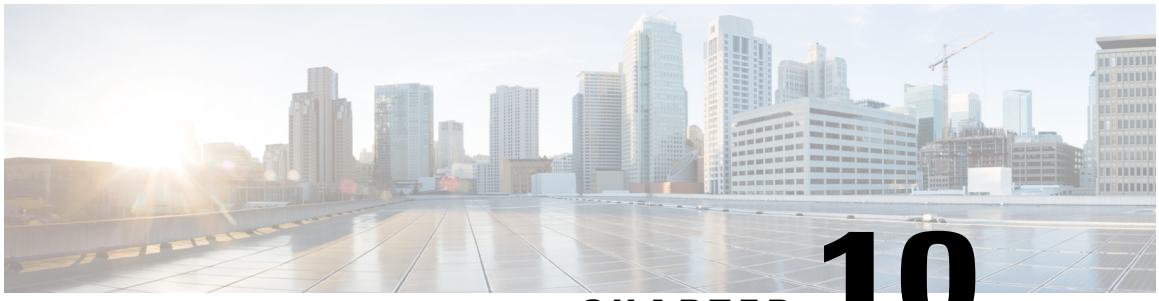

CHAPTER 10

# 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 77
- Viewing Virtual Machine Details, page 78
- Viewing Baremetal Details, page 78
- Viewing Network Details, page 79
- Viewing Router Details, page 80


## 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.
Select the tenant from the dropdown. The following details are displayed:

- Network Name
- Zone
- Admin State
- Status
- Subnets—Count of subnet. Click on the link to get subnet details.
- Network Type
- External Network
- VNI
- Extended

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

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 tenant.
The following details are displayed:

- Name
- Zone
- UUID
- Gateway IP
- Status
- Connected Networks
- VNI
- VRF
- 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.

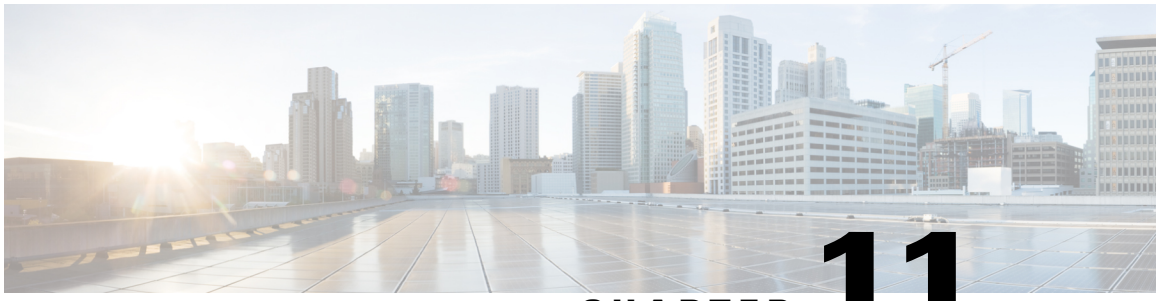

CHAPTER 11

# Administering Cisco VTS

This chapter has the following topics:


- Setting up System, page 81
- Modifying Login Banners, page 82
- Setting Global Route Reflector, page 82
- Registering the Virtual Machine Manager using GUI, page 82
- Backing up the Database, page 83
- Restoring the Database, page 84
- Viewing HA Status, page 85
- Enabling External Authentication and Authorization, page 85
- Enabling Accounting and Logging, page 87


## 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** Specify the **Out-of-Sync Commit** behavior to control the Check Sync feature. See Synchronizing Configuration, on page 25 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 6** 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. Spine 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. To do this:

**Note** To install the Host Agent in case of an OpenStack environment, you can use the Install Capabilities feature.

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

**Step 2** Click the Add (+) button.
The Add Virtual Machine Manager popup is displayed.

**Step 3** Enter the following details:

- VMM Type—Specify the VMM type. Choose vcenter or openstack from the drop-down list.
- Version Name—Specify the version details. See Note below for details.
- Description—Enter a description for the VMM.
- IP Address-Port—If you choose vcenter, enter the IP address and the port. The default port is 443. If you choose openstack, enter the IP address.
- User Name—Enter the VMM username.
- Passphrase—Enter the VMM passphrase.

**Note**

- **1** If VMM Type is **openstack** and Version Name is **openstack-liberty-centos**, a new field VTF Deployment Mode appears in the Add VMM popup.
  **2** VTF Deployment Mode field had two modes: **vtf-vhost** and**vtf-as-vm**.
  **3** openstack-liberty-rhel, openstack-icehouse, openstack-juno, and openstack-kilo support only VM mode.
  **4** Openstack-liberty-centos supports both vhost and VM modes.
  **5** In vhost mode VTF installation has to be done manually.
  **6** VTF deployment mode can be selected at the time of VMM registration and this cannot be changed later.
  **7** All the computes in a data center should be either in vhost mode (VTF running as a process) or in vNode mode (VTF running as a VM). Currently VTS does not support mixed mode.

**Step 4** Click **Save**.
The VMM you added is listed in the Virtual Machine Manager screen.

You can check the status of VMM registration in the Status column.
To delete a VMM, select the VMM and click **X** (delete).

# 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.

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 86 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 *default-map* 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 *default-map* 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 85 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.

**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.

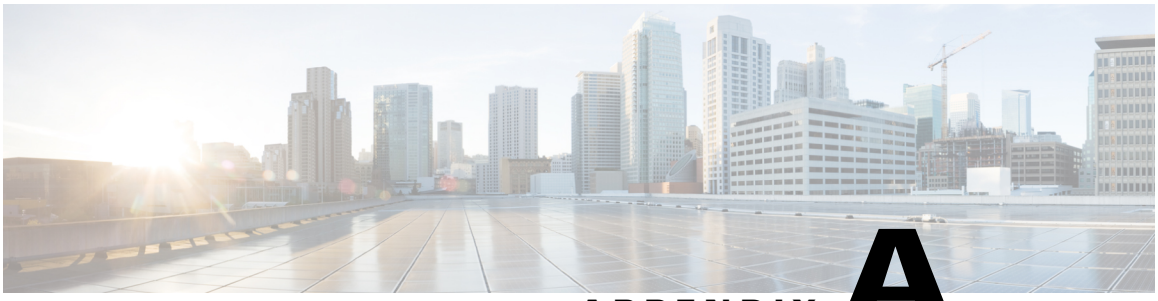

APPENDIX A

# Service Extension Templates-Supported Configuration

This appendix provides information about supported service extension template configuration.


- Supported Service Extension Template Configuration for Cisco Nexus 7000 Series Switches, page 89
- Supported Service Extension Template Configuration for Cisco ASR 9000 Series Routers , page 90


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

This section provides information about the service extension configuration 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
```

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
```

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
```

# 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
```

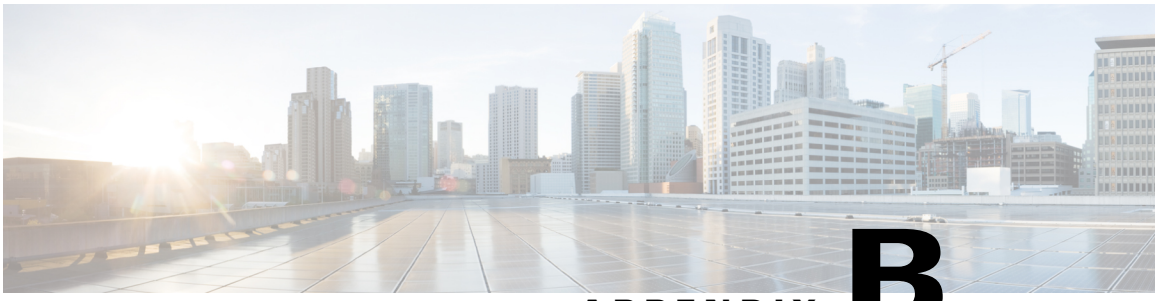

APPENDIX B

# Supported Underlay Configuration

This appendix provides information about supported underlay template configuration.


- Supported Underlay Configuration, page 93


## Supported Underlay Configuration

The following table lists the supported underlay configuration:

| Configuration Area | Sample Configuration | Device Type | Device Role |
|---|---|---|---|
| UNDERLAY IGP ROUTING OSPF routing process | `router ospf 10` | N9K | Leaf |
| | `router ospf 10` | ASR9K | DCI |
| OSPF Area | `interface ethernet 1/5`<br>`  ip ospf router 10 area 0.0.0.0` | N9K | Leaf |
| | `interface ethernet 1/5`<br>`  ip ospf router 10 area 0.0.0.0` | N9K | Leaf |
| | `router ospf 10`<br>`  area 0` | ASR9K | DCI |
| OSPF router-id | `router ospf 10`<br>`  router-id 10.218.20.15` | N9K | Leaf |
| | `router ospf 10`<br>`  router-id 10.218.20.15` | ASR9K | DCI |

| | | | |
|---|---|---|---|
| OSPF auto-cost reference | router ospf 10<br>!<br>auto-cost reference-bandwidth 800000 | ASR9K | DCI |
| OSPF Network type | interface ethernet1/5<br>ip ospf network point-to-point | N9K | Leaf |
| | interface vlan10<br>ip ospf network point-to-point | N9K | Leaf |
| | router ospf 10<br>area 0<br>interface GigabitEthernet0/0/1/3<br>network point-to-point | ASR9K | DCI |
| OSPF Authentication | interface Ethernet1/5<br>ip ospf authentication message-digest | N9K | Leaf |
| | interface Ethernet1/5<br>ip ospf message-digest-key 1 md5 0 xxx | N9K | Leaf |
| | router ospf 10<br>area 0<br>interface <Fabric Interface><br>authentication message-digest<br>message-digest-key 1 md5 encrypted<br>202cb962ac59075b964b07152d234b70 | ASR9K | DCI |
| OSPF Passive-interface | interface loopback3<br>ip router ospf 100 area 0.0.0.0 | N9K | Leaf |
| | router ospf 10<br>area 0<br>interface Loopback10<br>passive enable | ASR9K | DCI |
| OSPF Convergence | router ospf 10<br>timers lsa arrival 15<br>timers throttle lsa 0 20 5000<br>timers throttle spf 50 100 5000 | N9K | Leaf |
| | router ospf 10<br>timers throttle lsa all 0 20 5000<br>timers throttle spf 50 100 5000<br>timers lsa min-arrival 15 | ASR9K | DCI |
| OSPF BFD (per-link) | feature bfd<br>router ospf 10<br>bfd | N9K | Leaf |
| | interface Ethernet1/5<br>no ip redirects | N9K | Leaf |
| | router ospf 10<br>bfd minimum-interval 150<br>bfd multiplier 3<br>area 0<br>interface TenGigE0/0/2/1<br>bfd fast-detect | ASR9K | DCI |
| | interface vlan 10<br>no bfd echo | N9K | Leaf |

| | | | |
|---|---|---|---|
| Multicast Routing | `feature pim` | N9K | Leaf |
| | `interface loopback1`<br>`  ip address 10.218.20.250/32`<br>`  ip router ospf 10 area 0.0.0.0`<br>`  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`<br>`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`<br>`  ip pim sparse-mode` | N9K | Leaf |
| | `interface loopback0`<br>`  ip pim sparse-mode` | N9K | Leaf |
| | `interface Ethernet2/1`<br>`  ip pim sparse-mode` | N9K | Leaf |
| | `interface Ethernet2/2`<br>`  ip pim sparse-mode` | N9K | Leaf |
| L2 Technologies | `interface Ethernet 1/10`<br>`  switchport mode trunk` | N9K | Leaf |
| | `interface Ethernet 1/10`<br>`  switchport trunk allowed vlan none` | N9K | Leaf |
| | `interface Ethernet 1/10`<br>`  spanning-tree port type edge trunk` | N9K | Leaf |
| | `interface Ethernet 1/10`<br>`  spanning-tree bpduguard enable` | N9K | Leaf |
| | `interface Ethernet 1/10`<br>`  spanning-tree bpdufilter enable` | N9K | Leaf |
| | `interface Ethernet 1/10`<br>`  storm-control broadcast level 20.0` | N9K | Leaf |
| | `interface Ethernet 1/10`<br>`  storm-control multicast level 30.0` | N9K | Leaf |
| | `interface Ethernet 1/10`<br>`  storm-control unicast level 50.0` | N9K | Leaf |
| | `interface Ethernet 1/10`<br>`  storm-control action shutdown` | N9K | Spine |
| vPC Role and Priority | `vpc domain 1`<br>`  role priority 100` | N9K | Leaf |
| | `vpc domain 1`<br>`  role priority 200` | N9K | Leaf |
| vPC Peer Keep-alive Link | `vrf context management` | N9K | Leaf |
| | `interface mgmt 0`<br>`  vrf member management` | N9K | Leaf |
| | `interface mgmt 0`<br>`  ip address 172.20.118.119/25`<br>`  no shutdown` | N9K | Leaf |
| | `vpc domain 1`<br>`  peer-keepalive destination 172.20.118.20` | N9K | Leaf |

| | | | |
|---|---|---|---|
| vPC Peer-Link | `interface Ethernet 1/1`<br>`  spanning-tree port type network`<br>`  channel-group 1 mode active no shutdown` | N9K | Leaf |
| | `interface Ethernet 1/2`<br>`  spanning-tree port type network`<br>`  channel-group 1 mode active no shutdown` | N9K | Leaf |
| | `interface port-channel1`<br>`  switchport`<br>`  switchport mode trunk`<br>`  spanning-tree port type network`<br>`  vpc peer-link` | N9K | Leaf |
| vPC Port | `interface Ethernet 2/9`<br>`  channel-group mode active id 51` | N9K | Leaf |
| | `interface port-channel 51`<br>`  switchport` | N9K | Leaf |
| | `interface port-channel 51`<br>`  switchport`<br>`  vpc 51` | N9K | Spine |

| | | | |
|---|---|---|---|
| vPC Peer-Switch Option | `vpc domain 1`<br>`  peer-switch` | N9K | Leaf |
| vPC ARP Synchronization | `vpc domain 1`<br>`  ip arp synchronize` | N9K | Leaf |
| vPC in VXLAN environment adjustment | `vpc domain 10`<br>`   peer-switch`<br>`   system-priority 100 ( could not find this option)`<br><br>`   peer-keepalive destination 172.20.118.120`<br>`   delay restore 200`<br>`   peer-gateway`<br>`   ip arp synchronize` | N9K | Leaf |
| | `interface port-channel1`<br>`description vPC peer-link` | N9K | Spine |
| | `interface port-channel1`<br>`   description vPC switchport mode trunk` | N9K | Leaf |
| | `interface port-channel 1`<br>`  description vPC switchport mode trunk` | N9K | |
| | `interface port-channel 1`<br>`  description vPC spanning-tree port type network` | N9K | |
| | `interface port-channel 1`<br>`  vpc peer-link` | N9K | |
| | `interface port-channel 10`<br>`  switchport trunk allowed vlan` | N9K | |
| | `interface port-channel 10`<br>`  spanning-tree port type edge trunk` | N9K | |
| | `interface port-channel 10`<br>`  spanning-tree bpdufilter enable` | N9K | |
| | `interface port-channel 10`<br>`  spanning-tree bpduguard enable` | N9K | |
| | `interface port-channel 10`<br>`  vpc 10` | N9K | Spine |
| | `interface Ethernet 1/10`<br>`  switchport trunk allowed vlan none` | N9K | Spine |
| | `interface Ethernet 1/10`<br>`  spanning-tree port type edge trunk` | N9K | Leaf |
| | `interface Ethernet 1/10`<br>`  spanning-tree bpduguard enable` | N9K | Leaf |
| | `interface Ethernet 1/10`<br>`  spanning-tree bpdufilter enable` | N9K | Leaf |
| | `interface Ethernet 1/10`<br>`  channel-group 10 mode active` | N9K | Leaf |
| | `interface loopback 0`<br>`  ip address 10.218.20.1/32` | N9K | Leaf |
| | `interface loopback 0`<br>`  ip address 10.218.20.1/32 secondary` | N9K | Leaf |
| | `interface loopback 0`<br>`  ip router ospf 100 area 0.0.0.0` | N9K | Leaf |
| | `interface loopback 0`<br>`  ip pim sparse-mode` | N9K | Leaf |
| | `interface Vlan 10`<br>`  ip address 10.218.18.0/31` | N9K | Leaf |
| | | | |

| | | | |
|---|---|---|---|
| | interface Vlan 10<br>  description Underlay vPC Backup link<br>  no shutdown<br>  no bfd echo | N9K | Spine |
| | interface Vlan 10<br>  ip ospf network point-to-point | N9K | Leaf |
| | interface Vlan 10<br>  ip router ospf 100 area 0.0.0.0 | N9K | Leaf |
| | interface Vlan10<br>  ip pim sparse-mode | N9K | Leaf |

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

| | | | |
|---|---|---|---|
| | interface GigabitEthernet0/0/1/5<br>  mtu 9214 | ASR9K | DCI |
| Nexus 9500 QOS | system qos<br>  service-policy type queuing output<br>default-out-policy | N9K | Leaf |
| | policy-map type network-qos Jumbo-nq-policy<br>  class type network-qos c-nq3 | N9K | Leaf |
| | policy-map type network-qos Jumbo-nq-policy<br>class type network-qos c-nq3<br>  match qos-group 3 | N9K | Leaf |
| | policy-map type network-qos Jumbo-nq-policy<br>class type network-qos c-nq3<br>  mtu 9216 | N9K | Leaf |
| | class type network-qos c-nq3<br>  match qos-group 3<br>  mtu 9216<br>class type network-qos c-nq2<br>  match qos-group 2<br>  mtu 9216<br>class type network-qos c-nq1<br>  match qos-group 1<br>  mtu 9216<br>class type network-qos c-nq-default<br>  match qos-group 0<br>  mtu 9216 | N9K | Leaf |
| | system qos<br>service-policy type network-qos Jumbo-nq-policy | N9K | Leaf |
| QoS Hardware resources configuration | "hardware access-list tcam region racl 0<br>hardware access-list tcam region e-racl 0<br>hardware access-list tcam region span 0<br>hardware access-list tcam region vqos 256<br>hardware access-list tcam region e-qos 256<br>hardware access-list tcam region arp-ether 256" | N9K | Leaf |
| N 9500 QoS | system qos<br>service-policy type queuing output<br>default-out-policy | N9K | Leaf |
| | system qos<br>service-policy type network-qos Jumbo-nq-policy | N9K | Leaf |
| N 9500 QoS Queuing policy | policy-map type queuing default-out-policy<br>class type queuing c-out-q3<br>  priority level 1<br>class type queuing c-out-q2<br>  bandwidth remaining percent 0<br>class type queuing c-out-q1<br>  bandwidth remaining percent 0<br>class type queuing c-out-q-default<br>bandwidth remaining percent 100 | N9K | Leaf |
| | System qos<br>  Service-policy type queuing out<br>default-out-policy | N9K | Leaf |

| | | | |
|---|---|---|---|
| N 9500 QoS Queuing Policy | policy-map type queuing default-out-policy<br>class type queuing c-out-q3<br>priority level 1<br>class type queuing c-out-q2<br>bandwidth remaining percent 0<br>class type queuing c-out-q1<br>bandwidth remaining percent 0<br>class type queuing c-out-q-default<br>bandwidth remaining percent 100 | N9K | Leaf |
| | System qos<br>Service-policy type queuing out<br>default-out-policy | N9K | Leaf |
| Network Management Ethernet (Mgmt0) | interface mgmt0<br>ip address 10.218.22.10/24 | N9K | Leaf |
| | vrf context management<br>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<br>clock summer-time EEST 4 Sunday March 02:00 4 Sunday October 03:00 60 | N9K | Leaf |
| DNS | ip domain-name <cust_name><br>no ip domain-lookup | N9K | Leaf |
| SNMP | snmp-server contact <contact_name><br>snmp-server location <location_name> | N9K | Leaf |
| | snmp-server host 85.29.26.36 traps version 2c <SNMP_Community_1><br>snmp-server host 85.29.56.136 traps version 2c <SNMP_Community_1><br>snmp-server host 85.29.60.191 traps version 2c <SNMP_Community_1><br>snmp-server host 85.29.60.235 traps version 2c <SNMP_Community_1><br>snmp-server host 213.74.189.232 traps version 2c <SNMP_Community_1><br>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<br>snmp-server host 85.29.56.136 use-vrf management<br>snmp-server host 85.29.60.191 use-vrf management<br>snmp-server host 85.29.60.235 use-vrf management<br>snmp-server host 213.74.189.232 use-vrf management<br>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<br>...<br>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#`<br>`  no ip redirects`<br>`  no ipv6 redirects` | N9K | Leaf |
| Device Access Security | `NX-OS(config)#no ssh server enable`<br>`NX-OS(config)#ssh key {dsa [force] \| rsa [bits [force]]}`<br>`NX-OS(config)#ssh server enable`<br>`NX-OS#show ssh key`<br>`**********`<br>`rsa Keys generated:Fri Apr 10 20:13:21 2010`<br>`<clipped> !` | N9K | Leaf |
| AAA-N | `NX-OS(config)#feature tacacs+`<br>`NX-OS(config)#tacacs-server host {ipv4-address \| ipv6-address \| host-name}`<br>`NX-OS(config)#tacacs-server key [0 \| 7] key-value`<br>`NX-OS(config)#aaa group server tacacs+ group-name`<br><br>`  server {ipv4-address \| ipv6-address \| host-name}`<br>`deadtime minutes`<br>`  use-vrf vrf-name`<br>`NX-OS(config)#tacacs-server timeout seconds`<br>`NX-OS(config)#tacacs-server host {ipv4-address \| ipv6-address \| host-name} port tcp-port`<br>`NX-OS(config)#tacacs-server deadtime minutes` | N9K | Leaf |
| | `feature tacacs+`<br>`aaa group server tacacs+ TacacsGroup`<br>`  use-vrf management`<br>`  server 10.35.175.1`<br>`aaa authentication login console group TacacsGroup`<br>`aaa authentication login default group TacacsGroup`<br>`aaa authentication login error-enable`<br>`!`<br>`tacacs-server host 10.35.175.1 key <shared-key> port 49`<br>`tacacs-server directed-request`<br>`ip tacacs source-interface mgmt 0`<br>`!`<br>`! Device Login Authorisation with AAA`<br>`!` | N9K | Leaf |
| | `aaa authorization config-commands default group TacacsGroup local`<br>`aaa authorization commands default group TacacsGroup local`<br>`!`<br>`! Device Login Accounting with AAA`<br>`!`<br>`aaa accounting default group TacacsGroup`<br>`!`<br>`! Local User Configuration`<br>`!`<br>`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<br>copp-system-p-policy-strict<br>  class copp-system-p-class-l3uc-data<br>  set cos 1<br>  police cir 250 pps bc 32 packets conform<br>transmit violate drop | N9K | Leaf |
| | class copp-system-p-class-critical<br>  set cos 7<br>  police cir 19000 pps bc 128 packets conform<br>transmit violate drop | N9K | Leaf |
| | class copp-system-p-class-important<br>  set cos 6<br>  police cir 3000 pps bc 128 packets conform<br>transmit violate drop | N9K | Leaf |
| | class copp-system-p-class-multicast-router<br>  set cos 6<br>  police cir 3000 pps bc 128 packets conform<br>transmit violate drop | N9K | Leaf |
| | class copp-system-p-class-management<br>  set cos 2<br>  police cir 3000 pps bc 32 packets conform<br>transmit violate drop | N9K | Leaf |
| | class copp-system-p-class-multicast-host<br>  set cos 1<br>  police cir 2000 pps bc 128 packets conform<br>transmit violate drop | N9K | Leaf |
| | class copp-system-p-class-l3mc-data<br>  set cos 1<br>  police cir 3000 pps bc 32 packets conform<br>transmit violate drop | N9K | Leaf |
| | class copp-system-p-class-normal<br>  set cos 1<br>  police cir 1500 pps bc 32 packets conform<br>transmit violate drop | N9K | Leaf |
| | class copp-system-p-class-ndp<br>  set cos 6<br>  police cir 1500 pps bc 32 packets conform<br>transmit violate drop | N9K | Leaf |
| | class copp-system-p-class-normal-dhcp<br>  set cos 1<br>  police cir 300 pps bc 32 packets conform<br>transmit violate drop | N9K | Leaf |
| | class<br>copp-system-p-class-normal-dhcp-relay-response<br>  set cos 1<br>  police cir 400 pps bc 64 packets conform<br>transmit violate drop | N9K | Leaf |
| | class copp-system-p-class-normal-igmp<br>  set cos 3<br>  police cir 6000 pps bc 64 packets conform<br>transmit violate drop | N9K | Leaf |
| | class copp-system-p-class-redirect<br>  set cos 1<br>  police cir 1500 pps bc 32 packets conform<br>transmit violate drop | N9K | Leaf |
| | class copp-system-p-class-exception<br>  set cos 1<br>  police cir 50 pps bc 32 packets conform transmit<br>violate drop | N9K | Leaf |
| | | | |

| | | | |
|---|---|---|---|
| | class copp-system-p-class-exception-diag<br>  set cos 1<br>  police cir 50 pps bc 32 packets conform transmit<br> violate drop | N9K | Leaf |
| | class copp-system-p-class-monitoring<br>  set cos 1<br>  police cir 300 pps bc 128 packets conform<br>transmit violate drop | N9K | Leaf |
| | class copp-system-p-class-l2-unpoliced<br>  set cos 7<br>  police cir 20000 pps bc 8192 packets conform<br>transmit violate drop | N9K | Leaf |
| | class copp-system-p-class-undesirable<br>  set cos 0<br>  police cir 15 pps bc 32 packets conform transmit<br> violate drop | N9K | Leaf |
| | class copp-system-p-class-fcoe<br>  set cos 6<br>  police cir 1500 pps bc 128 packets conform<br>transmit violate drop | N9K | Leaf |
| | class copp-system-p-class-nat-flow<br>  set cos 7<br>  police cir 100 pps bc 64 packets conform<br>transmit violate drop | N9K | Leaf |
| | class copp-system-p-class-l2-default<br>  set cos 0<br>  police cir 50 pps bc 32 packets conform transmit<br> violate drop | N9K | Leaf |
| | class class-default<br>  set cos 0<br>  police cir 50 pps bc 32 packets conform transmit<br> violate drop | N9K | Leaf |
| | N9k-ST-Leaf-01# sh copp status<br>Last Config Operation: None<br>Last Config Operation Timestamp: None<br>Last Config Operation Status: None<br>Policy-map attached to the control-plane:<br>copp-system-p-policy-strict | N9K | Leaf |
| | N9k-ST-Leaf-01# sh copp profile ?<br>dense  Display dense profile<br>lenient  Display lenient profile<br>moderate Display moderate profile<br>strict  Display strict profile | N9K | Leaf |
| BFD | feature bfd<br>bfd interval 50 min_rx 50 multiplier 3 | N9K | Leaf |
| | router ospf UNDERLAY<br>  bfd | N9K | Leaf |
| | router bgp 100<br>vrf vrf1<br>  address-family ipv4 unicast | N9K | Leaf |
| | router bgp 100<br>  vrf vrf1<br>  local-as 65004 | N9K | Leaf |
| | router bgp 100<br>vrf vrf1<br>  neighbor 10.23.65.0 remote-as 22394<br>  bfd | N9K | Leaf |

| | | | |
|---|---|---|---|
| OSPF Routing Process | feature ospf<br>!<br>router ospf UNDERLAY<br>log-adjacency-changes detail<br>bfd | N9K | Leaf |
| OSPF Router ID | router ospf UNDERLAY<br>log-adjacency-changes detail<br>bfd<br>router-id <loopback17-ip-address> | N9K | Leaf |
| Enabling OSPF on interfaces | router ospf UNDERLAY<br>passive-interface default | N9K | Leaf |
| | continue from the above...<br>interface Ethernet1/5<br>ip router ospf UNDERLAY area 0.0.0.1<br>ip ospf bfd<br>ip ospf network point-to-point<br>no ip ospf passive-interface | N9K | Leaf |
| | interface loopback<id><br>ip router ospf UNDERLAY area 0.0.0.1 | N9K | Leaf |
| OSPF Authentication | interface eth <slot>/<port><br>ip ospf authentication message-digest<br>ip ospf message-digest-key <key-id> md5 0<br><clear-text-key> | N9K | Leaf |
| OSPF Reference-Bandwidth | router ospf UNDERLAY<br>auto-cost reference bandwidth 100Gbps | N9K | Leaf |
| Underlay OSPF Configuration on Leaf<br>Underlay OSPF Configuration on Spine | interface loopback17<br>ip router ospf UNDERLAY area 0.0.0.1 | N9K | Leaf |
| | interface eth<slot>/<port><br>ip router ospf UNDERLAY area 0.0.0.1<br>ip ospf network point-to-point<br>no ip ospf passive-interface<br>ip ospf bfd<br>ip ospf authentication message-digest<br>ip ospf message-digest-key <key-id> md5 0<br><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<br>ip pim sparse-mode | N9K | Spine |
| | interface ethernet 1/10<br>ip pim bfd-instance | N9K | Spine |
| | interface loopback<id><br>ip pim sparse-mode | N9K | Leaf |

| | | | |
|---|---|---|---|
| Mapping Layer 2 VNI VXLAN segment to ASM group | interface nve<id><br>member vni <L2-VNID><br>mcast-group 239.239.0.1<br>member vni <L2-VNID><br>mcast-group 239.239.0.2 | N9K | Leaf |
| PIM Anycast RP (RFC 4610) | interface loopback18<br>ip pim sparse-mode | N9K | Leaf |
| | interface loopback17<br>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<br>ip pim log-neighbor-changes | N9K | Leaf |
| | interface loopback17<br>ip pim sparse-mode | N9K | Leaf |
| | interface ethernet<slot>/<port><br>ip pim sparse-mode<br>ip pim bfd-instance | N9K | Leaf |
| | interface nve1<br>member vni <L2-VNID><br>mcast-group 239.64.64.1<br>member vni <L2-VNID><br>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<br>ip pim log-neighbor-changes | N9K | Spine |
| | interface ethernet 1/10<br>ip pim sparse-mode | N9K | Spine |
| | interface ethernet 1/10<br>ip pim bfd-instance | N9K | Spine |
| | interface loopback17<br>ip pim sparse-mode | N9K | Spine |
| | interface loopback18<br>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<br>vn-segment 10019 | N9K | |
| | interface Vlan17<br><br>mtu 9216<br>vrf member VRF-MEMBER-1<br><br>ip ospf cost 10<br>ip ospf passive-interface<br>ip router ospf 1 area 0.0.0.0 | N9K | |

| | | | |
|---|---|---|---|
| Service Extensions for Static routing (EGB has detail in their slide deck. Which config we get have at Q3. Please see EGB Slide Deck.) | `vrf context VRF-MEMBER-1`<br><br>`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`<br>`vrf VRF-MEMBER-1`<br>`address-family ipv4 unicast`<br>`network 0.0.0.0/0` | N9K | |
| route-map | `route-map RM-IN-S2 permit 10`<br>`match tag 1000`<br>`route-map RM-IN-S3 permit 10`<br>`match tag 1000` | N9K | |
| | `route-map RM-S-to-O permit 10`<br>`match tag 131 132 133 139 134 135`<br>`set metric-type type-1` | N9K | |
| vrf context RM-IN-S2 | `vrf context RM-IN-S2`<br>`ip route 9.59.207.0/24 Vlan1603 11.0.34.30 name qianyi-henfang tag 1000 50` | N9K | |
| | `vrf context RM-IN-S2`<br>`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`<br>`ip route 10.0.0.0/12 Vlan1603 11.0.34.30 tag 1000 50` | N9K | |
| | `vrf context RM-IN-S2`<br>`ip route 10.0.0.0/12 Ethernet1/46.2 11.0.40.142 tag 1000 10` | N9K | |
| | `vrf context RM-IN-S2`<br>`ip route 10.2.52.0/24 Vlan6 10.2.42.3 tag 1000` | N9K | |
| | `vrf context RM-IN-S2`<br>`ip route 192.168.0.0/16 Vlan1603 11.0.34.30 name beimaluceshiwang tag 1000`<br>`rd auto`<br>`address-family ipv4 unicast`<br>`route-target both auto`<br>`route-target both auto evpn` | N9K | |
| vrf context RM-OUT | `vrf context RM-OUT`<br>`ip route 10.2.0.0/19 Vlan1607 11.0.34.14 tag 131 50` | N9K | |
| | `vrf context RM-OUT`<br>`ip route 10.2.0.0/19 Ethernet1/45.1 11.0.40.145 tag 131 10` | N9K | |
| | `vrf context RM-OUT`<br>`ip route 10.2.96.0/19 Vlan3203 11.0.39.14 tag 134` | N9K | |

| | | | |
|---|---|---|---|
| interface Vlan1601 | interface Vlan1601<br>no shutdown<br>vrf member RM-IN-S3<br>no ip redirects<br>  ip address 11.0.34.34/28<br>no ipv6 redirects<br>hsrp version 2<br>hsrp 1601<br>    preempt<br>   priority 110<br>   ip 11.0.34.33 | N9K | |
| interface Vlanxx | interface Vlan1602<br>no shutdown<br>vrf member VRF-MEMBER-IN-S2<br>no ip redirects<br>no ipv6 redirects<br>ip ospf cost 10<br>ip ospf passive-interface<br>ip router ospf 100 area 0.0.0.0 | N9K | |
| interface Ethernet1/46.1 | interface Ethernet1/36.1<br>mtu 1500 | N9K | |
| | interface Ethernet1/36.1<br>encapsulation dot1q 1602<br>mac-address 0000.0000.2222<br>vrf member RM-IN-S3<br>no ip redirects<br>ip address 11.0.40.145/30 | N9K | |

| | | | |
|---|---|---|---|
| interface Ethernet1/47.1 | interface Ethernet1/37.1<br>mtu 1500<br>encapsulation dot1q 1608<br>vrf member RM-OUT<br>no ip redirects<br>ip address 11.0.72.42/30<br>ip ospf dead-interval 20<br>ip ospf hello-interval 5<br>ip ospf network point-to-point<br>ip router ospf 100 area 0.0.0.0 | N9K | |
| router ospf 1 | router ospf 1<br>vrf RM-IN-S2<br>router-id 55.2.32.5<br>vrf RM-IN-S3<br>router-id 55.2.32.5<br>vrf RM-OUT<br>router-id 55.2.32.5<br>redistribute static route-map RM-S-to-O | N9K | |
| router bgp 65003 | router bgp 65003<br>vrf RM-IN-S2<br>address-family ipv4 unicast<br>advertise l2vpn evpn<br>redistribute direct route-map<br>vts-subnet-policy<br>redistribute static route-map RM-IN-S2 | N9K | |
| | feature telnet<br>feature nxapi<br>feature ospf<br>feature bgp<br>feature pim<br>feature udld<br>feature interface-vlan<br>feature vn-segment-vlan-based<br>feature hsrp<br>feature lacp<br>feature vpc<br>feature lldp<br>feature nv overlay | N9K | |
| | nv overlay evpn | N9K | |
| | clock protocol ntp vdc 1 | N9K | |
| role name nsdcheck | role name nsdcheck<br>rule 4 permit command show *<br>rule 3 permit command terminal length *<br>rule 2 permit command ping *<br>rule 1 permit read | N9K | |
| | role name devcheck<br>rule 8 permit command tac-pac *<br>rule 7 permit command dir *<br>rule 6 permit command ssh *<br>rule 5 permit command traceroute *<br>rule 4 permit command ping * | N9K | |
| | role name devopera<br>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 $1$lDuqR.60$eNzZ5I22WxJT58gdEm88N0 role network-operator` | N9K | |
| | `username vtsadmin password 5 $5$MmpswImI$vbZhP/52dNjHY5KWj4yBvmiDvuOZZ9gd2vo2oZc61b4 role network-admin` | N9K | |
| | `username nsdcheck password 5 $5$dpIXMjZs$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`<br>`snmp-server host 55.6.8.1 use-vrf default` | N9K | |
| | `snmp-server enable traps bgp`<br>`snmp-server enable traps ospf`<br>`snmp-server enable traps callhome event-notify`<br>`snmp-server enable traps callhome smtp-send-fail`<br>`snmp-server enable traps cfs state-change-notif`<br>`snmp-server enable traps lldp lldpRemTablesChange`<br>`snmp-server enable traps aaa server-state-change`<br>`snmp-server enable traps hsrp state-change`<br>`snmp-server enable traps feature-control FeatureOpStatusChange`<br>`snmp-server enable traps sysmgr cseFailSwCoreNotifyExtended`<br>`snmp-server enable traps config ccmCLIRunningConfigChanged`<br>`snmp-server enable traps snmp authentication`<br>`snmp-server enable traps link cisco-xcvr-mon-status-chg`<br>`snmp-server enable traps vtp notifs`<br>`snmp-server enable traps vtp vlancreate`<br>`snmp-server enable traps vtp vlandelete`<br>`snmp-server enable traps bridge newroot`<br>`snmp-server enable traps bridge topologychange`<br>`snmp-server enable traps stpx inconsistency`<br>`snmp-server enable traps stpx root-inconsistency`<br>`snmp-server enable traps system Clock-change-notification`<br>`snmp-server enable traps feature-control ciscoFeatOpStatusChange` | N9K | |
| | `snmp-server community COMMUNITY1 group network-operator` | N9K | |

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

| | | | |
|---|---|---|---|
| interface mgmt0 | `interface mgmt0`<br>`no lldp transmit`<br>`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`<br>`ip route 0.0.0.0/0 Vlan3903 55.6.40.14 tag 1000 50` | N9K | |
| router ospf | `router ospf 1`<br>`redistribute static route-map RM-S-to-O` | N9K | |
| router bgp | `router bgp 65003`<br>`router-id 55.2.32.5`<br>`address-family ipv4 unicast`<br>`address-family l2vpn evpn`<br>`neighbor 55.2.32.1`<br>`  remote-as 65003`<br>`  update-source loopback1`<br>`  address-family ipv4 unicast`<br>`  address-family l2vpn evpn`<br>`  send-community extended` | N9K | |