



Cisco Virtual Topology System (VTS) 2.6.6 User Guide

First Published: 2021-10-27

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

THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS IN THIS MANUAL ARE SUBJECT TO CHANGE WITHOUT NOTICE. ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS MANUAL ARE BELIEVED TO BE ACCURATE BUT ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. USERS MUST TAKE FULL RESPONSIBILITY FOR THEIR APPLICATION OF ANY PRODUCTS.

THE SOFTWARE LICENSE AND LIMITED WARRANTY FOR THE ACCOMPANYING PRODUCT ARE SET FORTH IN THE INFORMATION PACKET THAT SHIPPED WITH THE PRODUCT AND ARE INCORPORATED HEREIN BY THIS REFERENCE. IF YOU ARE UNABLE TO LOCATE THE SOFTWARE LICENSE OR LIMITED WARRANTY, CONTACT YOUR CISCO REPRESENTATIVE FOR A COPY.

The Cisco implementation of TCP header compression is an adaptation of a program developed by the University of California, Berkeley (UCB) as part of UCB's public domain version of the UNIX operating system. All rights reserved. Copyright © 1981, Regents of the University of California.

NOTWITHSTANDING ANY OTHER WARRANTY HEREIN, ALL DOCUMENT FILES AND SOFTWARE OF THESE SUPPLIERS ARE PROVIDED "AS IS" WITH ALL FAULTS. CISCO AND THE ABOVE-NAMED SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, THOSE OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NON-INFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE.

IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THIS MANUAL, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

Any Internet Protocol (IP) addresses and phone numbers used in this document are not intended to be actual addresses and phone numbers. Any examples, command display output, network topology diagrams, and other figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses or phone numbers in illustrative content is unintentional and coincidental.

All printed copies and duplicate soft copies of this document are considered uncontrolled. See the current online version for the latest version.

Cisco has more than 200 offices worldwide. Addresses and phone numbers are listed on the Cisco website at www.cisco.com/go/offices.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: <https://www.cisco.com/c/en/us/about/legal/trademarks.html>. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1721R)

© 2021 Cisco Systems, Inc. All rights reserved.



CONTENTS

CHAPTER 1

Introduction 1

- Understanding Cisco VTS 1
- Cisco VTS Architecture Overview 2
- Virtual Topology Forwarder 3
- Multi-Site Support 3
- Virtual Topology System High Availability 4

CHAPTER 2

Getting Started with Cisco Virtual Topology System 7

- Logging in 7
- Creating Sites 8
- Setting up Sites 8
- Deleting a Site 8
- Using the Quick Guide 9
- Initial Configuration Tasks 10
- Notes Regarding VMware vSphere Distributed Switch 14

CHAPTER 3

Global Settings 17

- Transaction Settings 17
- Modifying Login Banners 18
- Managing Users 18
- Enabling External Authentication and Authorization 19
 - Setting up Remote Authentication Server 20
- Enabling Accounting and Logging 21
 - Setting up Accounting 21
- Creating Authentication Groups 22
- Viewing HA Status 23

Enabling Auto Southbound Lock 24

CHAPTER 4

Administering Cisco VTS 25

Setting up Site 25

Setting Global Route Reflector 26

Registering the Virtual Machine Manager using GUI 26

Uninstalling the OpenStack Plugin 28

Integrating Cisco VTS with Multiple Virtual Machine Managers 28

Performing Merge Operation from VMM 32

Performing Merge Operation from Tenant 33

Performing Merge Operation from Network 34

Performing Publish Operation from VMM 34

Performing Publish Operation from Tenant 35

Performing Publish Operation from Network 35

Performing Publish Operation from VTS 35

Backing up the Database in non HA Mode 36

Restoring the Database in non HA Mode 37

Backing up the Database in HA Mode 38

Restoring the Database in HA Mode 39

Configuring Syslog for Monitoring Logs 41

Troubleshooting Syslog Issues 43

Collect Data Before Contacting Technical Support 44

Troubleshooting Command Examples and Log Files 46

Problem-Dependent Data Collection Examples 47

CHAPTER 5

Monitoring Cisco VTS 49

Monitoring Cisco VTS Infrastructure using Monit 49

Metrics Collected using Monit-D 51

Setting up Monit Credentials 53

Setting up Policy Plane Credentials 53

Setting up Control Plane Credentials 53

Setting up Data Plane Credentials 53

Viewing Metrics Collected by Monit 53

Viewing Policy Plane Metrics 53

Viewing Control Plane Metrics	54
Viewing Data Plane Metrics	54
Monitoring Cisco VTS Infrastructure using collectd	54
Setting up collectd Plugins	58
Setting up Policy Plane Plugins	58
Setting up Data Plane Plugins	58

CHAPTER 6
Managing Inventory 61

Importing Inventory using CSV File	62
Performing Auto Discovery	65
Enabling lldpd on Computes	66
Enabling lldpd Using Anisble	67
Enabling lldpd Manually	67
Working with Discovered Data	68
Important Notes	73
Viewing the Network Topology	73
Viewing Network Inventory	74
Adding Fabric Connection	75
Synchronizing Configuration	75
Important Notes	76
Viewing Host Inventory	77
Adding a new Host on Virtual Servers	77
Adding a new Host on Baremetal	79
Viewing the VTSR to VTF Mapping	80
SR-IOV Support	80
Trunk Port Support	81
Migrating from vPC to ESI	81
Redeploying Device Inventory	83
Enabling Static Multi Homing	84
Enabling Static Multi Homing on Cisco Nexus 7000	84
Enabling Static Multi Homing on Cisco Nexus 9000	84
Administrative State for Devices	85

CHAPTER 7
Managing Resources 87

Specifying Global Provider VLAN Range	88
Global Provider VLAN Tool	88
Specifying Global VNI Range	89
Specifying Global EVI Range	90
Specifying VLAN Range	90
Specifying Device VLAN Range	91
Specifying Group VLAN Range	92
Specifying Interface VLAN Range	92
Creating Interface Groups	93
Specifying Multicast IP Pool	95
Important Note	95
Resource Pool Use Cases	96
Static VLAN Options	97

CHAPTER 8 **Creating and Managing Admin Domains** **101**

Admin Domain Overview	101
Viewing Admin Domain	102
Creating an Admin Domain	102
Creating DCI Interconnect Profiles	105

CHAPTER 9 **Managing Templates and Device Objects** **107**

Creating Route Templates	109
Adding Route Targets	109
Adding Fabric Internal Route Targets	109
Adding Fabric External Route Targets	110
Disabling Auto Route Target Configuration	111
Creating L3 Extension Templates	111
Supported System Variables	112
Editing Templates	113
Copying a Template	114
Deleting Templates	114
Importing and Exporting Templates	115
Importing Templates	115
Exporting Templates	115

Attaching Templates to Routers	116
Attaching Templates while Adding Routers	116
Attaching Template while Editing a Router	117
Creating L2 Extension Templates	118
Important Notes—L2 QoS Template	119
Important Notes—VPLS Template	119
Attaching Devices to L2 Extension Templates	120
Detaching Devices from L2 Extension Templates	121
Creating Underlay Templates	122
Attaching Underlay Template to Devices	123
Detaching Underlay Template from Devices	124
Previewing Template Configuration	124
Preview Template Configuration Examples	126
Searching Template Content	127
Using Search Hints	128
Creating Device Objects	129
Editing Device Objects	129
Deleting Device Objects	130
Associating Device Objects to Devices	130
Editing Device Object Instances	131
Deleting Device Object Instances	131
Searching Device Object	132
Device Objects Notes and Caveats	133

CHAPTER 10
Managing Tenants 137

Viewing Tenant Details	137
Adding Tenants	137
Editing Tenants	138

CHAPTER 11
Deploying Security Groups 139

Security Group - Feature Scope	140
Support for Reflexive ACLs	142
Creating Security Groups from Cisco VTS GUI	142
Attaching Security Group to Baremetal Port	143

Detaching Security Group from Baremetal Port	144
Attaching Security Groups to OVS, VTF, and SR-IOV Ports	144
Detaching Security Groups from OVS, VTF, and SR-IOV Ports	144
Security Group - Examples	144
Creating Security Group to Restrict Access to a Given Application	145
Associating SRIOV port with Security Group	145

CHAPTER 12**Provisioning Overlay Networks 147**

Provisioning Overlay Networks Using Cisco VTS	147
Creating Overlays	148
Using OpenStack	149
Using VMware	149
Using Cisco VTS GUI	149
Creating Network using VMware	150
Creating Subnetwork using VMware	150
Creating Routers using VMware	150
Attaching Network to Router	150
Attaching a Virtual Machine to Network	151
Creating a Network using Cisco VTS GUI	151
Creating a Subnetwork	152
Creating Router using Cisco VTS GUI	152
Port Extensions Support	155
Creating a Port Extension	155
Editing a Port Extension	158
Attaching Port Extension to Baremetal Ports	158
Detaching Port Extension from Baremetal Ports	159
Attaching Port Extension to Virtual Machine Ports	160
Detaching Port Extension from Virtual Machine Ports	161
Assigning BVI Interface IP Address	161
Extending Layer 2 Network Across Data Centers	162
Enabling Global Route Leaking Service	162
Enabling L3VPN to EVPN Route Stitching	164
Adding Static Routes	164
Adding Fabric Static Routes	165

	Adding External Static Routes	165
	Adding Port Static Routes	166
	OpenStack Allowed Address Pairs Support	167
<hr/>		
CHAPTER 13	Viewing Overlay Details	169
	Viewing Device Details	169
	Viewing Virtual Machine Details	170
	Viewing Baremetal Port Details	170
	Viewing Network Details	171
	Viewing Router Details	172
<hr/>		
APPENDIX A	Service Extension Templates-Supported Configuration Examples	173
	Supported Service Extension Template Configuration Examples for Cisco Nexus 7000 Series Switches	173
	Supported Service Extension Template Configuration Examples for Cisco Nexus 9000 Series Switches	176
	Supported Service Extension Template Configuration Examples for Cisco ASR 9000 Series Routers	179
<hr/>		
APPENDIX B	Supported Underlay Configuration Examples	181
	Supported Underlay Configuration Examples	181
<hr/>		
APPENDIX C	OpenStack Configuration for SR-IOV Support	203
	Sample for SR-IOV Trunk (No-Bonding)	205
	Sample for SR-IOV Trunk (Bonding)	205
<hr/>		
APPENDIX D	collectd Plugin Configuration for VTC and VTF	209
	collectd Plugin Configuration	209
	Policy Plane (VTC) Plugin Configuration	209
	Data Plane (VTF)—Plugin Confgs	210
	Write_Http Plugin Format	211
<hr/>		
APPENDIX E	collectd Output JSON Examples	213
	Default Plugins—JSON Examples	214
	Custom Plugin—JSON Examples	214



CHAPTER 1

Introduction

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

- [Understanding Cisco VTS, on page 1](#)
- [Cisco VTS Architecture Overview, on page 2](#)
- [Virtual Topology Forwarder, on page 3](#)
- [Multi-Site Support, on page 3](#)
- [Virtual Topology System High Availability, on page 4](#)

Understanding Cisco VTS

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

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

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

Cisco VTS provides:

- Fabric automation
- Programmability
- Open, scalable, standards based solution
- Cisco Nexus 2000, 3000, 5000, 5500, 7000, and 9000 Series Switches. For more information, see Supported Platforms in *Cisco VTS Installation Guide*.
- Software forwarder (Virtual Topology Forwarder [VTF])

VTS performs the role of an overlay orchestrator in data-center networks. In this role, it manages configuration on the data center leaf and spine devices. The configuration of the devices is dependent on the type of overlay service that the Cisco VTS user intends to create. The Cisco VTS user in this context could either be manual

users interfacing via GUI or APIs, or could be virtual machine managers like OpenStack or vSphere. Since the device configuration is derived from overlay service instances, Cisco VTS holds the ‘desired’ device configuration in its database.

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

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

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



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

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

Cisco VTS Architecture Overview

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

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

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

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

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

- **Control Plane:** The control plane module serves as the SDN control subsystem that programs the various data planes including the VTFs residing on the x86 servers, hardware leafs, DCI gateways. The Control plane hosts Service Routing (SR) module, which provides routing services to Cisco VTS. The Service Routing (SR) module is responsible for calculating L2 and L3 tables and routes to provide connectivity between the different VMs for a given tenant and service chaining. The main components of this module are the VTSR and VTF. VTSR is the controller and Virtual topology forwarder (VTF) runs on each compute server hosting the tenant VMs.

Virtual Topology Forwarder

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

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

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

Multi-Site Support

With the Multi-site feature, a single Cisco VTS instance can manage multiple sites (within the scale limit).

We have tested 20 sites with three sites fully loaded including two VMM per site.

In order to support multiple sites, Cisco VTS introduces a new construct called “Site”.

Earlier, the Cisco VTS services were modeled for a single Site in an instance of Cisco VTS. Currently, this is enhanced to support multiple sites.

A Site is an abstraction which is introduced to provide namespace isolation (for tenant/network/router) and to manage resources like VNI across Admin domains. A Single VTS instance includes one redundant pair of VTC and multiple redundant pairs of VTSR.

A Site can have multiple admin domains based on the topology. The overlay objects are specific to a Site and they cannot be stretched across Sites.

A VTSR instance can manage multiple Virtual Forwarding Groups (VFG) and multiple Admin Domains within a site. The VFG is a grouping of homogenous VTFs. Multiple instances are allowed for scaling with increased number of sites. Floating Route Reflectors (RR) and DC GW can either be shared across Admin Domain or be dedicated to an Admin Domain. Computes may span across Admin Domains. Inter-Admin Domain traffic may require the DC GW.

Virtual Machine Managers (VMMs) are local to a site. Inventory discovery and management are local to a site. Also, Templates are local to a site.

If you click the Delete Site icon and if the site has workloads, depending upon the type of workload the following error message is displayed:

```
Error: com.tailf.maapi.MaapiException: Exception in callback: Corresponding Site
e67165b1-e352-4d57-8111-42a42f5ec857 can not be deleted, Please delete workload first.
```

If the site has devices associated with it (even though all other workload is removed), the following error message is displayed

```
Error: com.tailf.maapi.Maapi Exception: illegal reference /ncs:devices/device{device1}
/vts-device-meta-data/site-id <- this indicates that device is still associated with site
that's getting deleted.
```

Constraints

Multi-site feature has the following constraints:

- VTC HA pair is Global and reachable to all DC sites created in the fabric.
- VTSR HA pair is Site specific and currently supported as V-deployment in a single site.



Note Multiple V-deployment is not supported and you can have only one V site.

- VTF's are site specific and currently restricted only to the site which has VTSR registered to.
- Each VTS installation supports either VXLAN or SR fabrics, but not both. For example:

For VXLAN or SR deployments of VTS multi-fabric/sites is supported (VTS will support multiple sites of VXLAN or SR).

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 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, on page 7](#)
- [Creating Sites, on page 8](#)
- [Setting up Sites, on page 8](#)
- [Deleting a Site, on page 8](#)
- [Using the Quick Guide, on page 9](#)
- [Initial Configuration Tasks, on page 10](#)
- [Notes Regarding VMware vSphere Distributed Switch, on page 14](#)

Logging in



Note On Login page, button links namely Change Passphrase, Home, Monitor, and Global Settings are visible to user even before they login. If user has not logged in, clicking these button does not have any effect i.e., the buttons are not functional.

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.
- Step 2** Enter the username and passphrase, and click **Login**. The default username and passphrase is admin/admin.
The Change Password window appears.
You are required to change the passphrase for the admin account the first time you are logging in. Click the Passphrase guidelines link in the Cisco VTS GUI 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**.

To change your passphrase subsequently, click **Change Passphrase** on the top right settings button. See the *Changing Password for Cisco VTS from VTS GUI* in the *Cisco VTS Installation Guide*, for details.

The Cisco VTS welcome screen is displayed. You can create and set up Sites from this UI.

After you create Sites:

- You may use the Setup Wizard, which displays the tasks you need to complete in order to get started with the system.
- Proceed with the tasks on your own, if you are familiar with the Cisco VTS setup tasks. You can access the Quick Guide anytime from the Settings menu on the top right corner of the Cisco VTS GUI.

Creating Sites

After you log in to Cisco VTS, you must create and set up Sites.

To create a site:

-
- Step 1** Click the + icon in the Sites page. The **Inventory / Discovery** window appears.
- Step 2** Enter the **Site Name**. This is mandatory.
- Step 3** (Optional) Enter the **Site ID**. If you enter the site ID, it should be in the UUID format.
- Step 4** Enter the **Description**.
- Step 5** Specify the Fabric Data Plane type. It can be EVPN VXLAN or MPLS SR.
- Note** Creating second site will default to the same data plane as 1st (as co-exist of different data plane is not supported).
- Step 6** Click the tick icon. The site gets created.
- You can use the **Edit** icon in case you need to edit the site details. To cancel the changes, click the **X** icon.
-

Setting up Sites

System-level settings for Cisco VTS are available under Global Settings. Site-level settings reside under Administration.

See [Global Settings, on page 17](#) for the various global settings options.

See [Administering Cisco VTS, on page 25](#) for the site-level settings options.

Deleting a Site

You can click on Delete icon to delete a site:

- If you click the Delete Site icon and if the site has workloads, depending upon the type of workload the following error message is displayed:

```
Error: com.tailf.maapi.MaapiException: Exception in callback: Corresponding Site
e67165b1-e352-4d57-8111-42a42f5ec857 can not be deleted, Please delete workload first.
```

- If the site has devices associated with it (even though all other workload is removed), the following error message is displayed.

```
Error: com.tailf.maapi.Maapi Exception: illegal reference /ncs:devices/device{device1}
/vts-device-meta-data/site-id <- this indicates that device is still associated with
site that's getting deleted.
```

In that case, delete the corresponding workload before deleting site.

You cannot delete a site:

- If it has no inventory, but is registered to a VMM
- Resources defined.

To delete a site:

Click the delete icon in the Site tile.

If the site has workloads depending upon the type of workload the following error message is displayed:

```
Error: com.tailf.maapi.MaapiException: Exception in callback: Corresponding Site
e67165b1-e352-4d57-8111-42a42f5ec857 can not be deleted, Please delete workload first.
If the site has devices associated with it (even though all other workload is removed),
the following error message is displayed
Error: com.tailf.maapi.Maapi Exception:
illegal reference /ncs:devices/device{device1}
/vts-device-meta-data/site-id <- this indicates that device is still associated with site that's
getting deleted.
```

In this case you should delete the corresponding workload before deleting site.

Using the Quick Guide

On logging in for the first time, 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 10](#) 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	<ul style="list-style-type: none"> • Site Settings • Virtual Machine Manager 	<ul style="list-style-type: none"> • Setting up Site, on page 25 • Registering the Virtual Machine Manager using GUI, on page 26
Set up Inventory	<ul style="list-style-type: none"> • Discover Devices 	<ul style="list-style-type: none"> • Managing Inventory, on page 61
Admin Domains	Create Admin Domains	Creating an Admin Domain, on page 102
Set up Tenants	Add Tenants	Adding Tenants, on page 137
Set up Overlay	<ul style="list-style-type: none"> • Add Network • Define Baremetal • Add Virtual Machines • Add Routers 	Provisioning Overlay Networks, on page 147

Initial Configuration Tasks

After you create and set up Sites, you need 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	Discover the Topology of all the leafs, spine, border-leafs, and DCI	Discovery > Topology Discovery For more information about adding devices and host information, see Performing Auto Discovery, on page 65 .	VTFs are not detected in topology discovery.

Sequence	Task	Navigation in VTS GUI / User Guide Section	Additional Notes
2	Import the devices after adding the auth group	Inventory > Import Inventory For more information about adding devices and host information, see Importing Inventory using CSV File , on page 62.	
3	Add the DHCP Server IP and Anycast Gateway MAC	Administration > Site Settings For more information about adding devices and host information, see Administering Cisco VTS , on page 25.	
4	Perform VTSR and VTF Registration. First bring up VTSR and let it register with VTC. Then bring up VTFs. Note This step is required only if you have a VTF-based deployment.	See the <i>Installing the Virtual Topology Forwarder</i> section in the <i>Cisco VTS Installation Guide</i> .	To verify that the VFG group is created, go to Inventory > Virtual Forwarding Groups.
5	Update the BGP ASN information for the devices	Inventory > Network Inventory For more information, see Viewing Network Inventory , on page 74.	

Sequence	Task	Navigation in VTS GUI / User Guide Section	Additional Notes
6	Create resource pools	<ul style="list-style-type: none"> • For the site with VxLAN data plane, Resource Pools > Global VNI Pool. • For the site with MPLS data plane, Resource Pools > Global EVI. • Resource Pools > Device Specific VLAN Pools • Resource Pools > Multicast IP Pool <p>For more information about creating an admin domain, see Managing Resources, on page 87.</p>	

Sequence	Task	Navigation in VTS GUI / User Guide Section	Additional Notes
7	Set the Route Reflector	Administration > Route Reflector For more information about setting global route reflector, see Setting Global Route Reflector, on page 26 . Note: Cisco VTS supports both global and admin-domain scoped route-reflector (via route-reflector functional group). So, there are 2 possible ways to configure route-reflector.	
8	Create an Admin Domain	Admin Domains > Domains For more information about creating an admin domain, see Creating an Admin Domain, on page 102 .	Properties for the L2/L3 Gateway Group are as follows: <ul style="list-style-type: none"> • Control Protocol: BGP-EVPN • Replication Modes : Multicast and Ingress • Distribution Mode: Decentralized

Sequence	Task	Navigation in VTS GUI / User Guide Section	Additional Notes
9	Add the devices to the Gateway Group	Admin Domains > Domains For more information about creating an admin domain, see Creating an Admin Domain, on page 102.	See the <i>Supported Platforms</i> section in the <i>Cisco VTS Installation Guide</i> for details about devices support for different roles.
10	Add the ToR and VTSR to the L2 and L3 Gateway Group	Admin Domains > Domains For more information about creating an admin domain, see Creating an Admin Domain, on page 102.	
11	Add devices to route-reflector functional group	Admin Domains > Domains For more information about creating an admin domain, see Creating an Admin Domain, on page 102.	
12	Save the Admin Domain you created	Admin Domains > Domains For more information about creating an admin domain, see Creating an Admin Domain, on page 102.	

Notes Regarding VMware vSphere Distributed Switch

Take care of the following points while you create a vDS:

**Note**

- All the ToRs in the inventory should be part of the vDS.
- 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.

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

Global Settings

This chapter has the following sections:

- [Transaction Settings](#), on page 17
- [Modifying Login Banners](#), on page 18
- [Managing Users](#), on page 18
- [Enabling External Authentication and Authorization](#), on page 19
- [Enabling Accounting and Logging](#), on page 21
- [Creating Authentication Groups](#), on page 22
- [Viewing HA Status](#), on page 23
- [Enabling Auto Southbound Lock](#), on page 24

Transaction Settings

To set up the transactions:

Step 1 Go to **Global Settings > Transaction Settings**.

Step 2 Specify the **Out-of-Sync Commit** behavior to control the Check Sync feature. See [Synchronizing Configuration](#), on page 75 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 3 Enable/disable Device South Bound Lock—Device southbound lock is enabled by default. When VTS has a redundant pair or group, it is possible for a transaction to succeed even when one or more of the redundant members are down, as long as one device is up. When the transaction comes, VTS checks the connectivity to the redundant devices and if it can not reach one of the devices, the admin state of the device will be changed to southbound-locked and the transaction configuration will only be pushed to the active devices. In order for the southbound lock feature to work, you must create a umap and provide the credentials that NSO will use, in the authgroup "vts-default" . This feature currently supports the following redundant groups:

- VPC Pair
- ESI Group

- Static Multi-Homed devices
- DCI
- VTSR

Step 4 Click **Save**.

Modifying Login Banners

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

To modify login banners:

- 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**.
-

Managing 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.

- Step 5** You must add the user to the Authentication group. To do this go to **Inventory > Authentication Group**, and add the user.
- Step 6** Log out and log in again with the new user.

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+ servers and LDAP servers.



Note Cisco VTS does not support Accounting via LDAP servers.

You can add multiple TACACS+ servers and LDAP servers. The authentication servers are chosen from the list of configured servers based on the priority that you set when you configure external authentication.

See the TACACS+ documentation for installing and configuring the TACACS+ server on the IPv4/IPv6 network.

Cisco VTS supports OpenLDAP and MS-ActiveDirectory implementations of LDAP. See the respective documentation for details about installing and configuring the LDAP servers.

For a user logging into VTS to be able to authenticate via TACACS+ server or LDAP server, the VTS admin needs to set up the external authorization servers.

For TACACS+ 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

On LDAP Server, users should be mapped to "Administrator / Operator" role for VTS to authorize the users. This can be done by making the memberUid attribute of Administrator / Operator group, the uid of the user.

See [Setting up Remote Authentication Server](#), on page 20 for details.

Important Notes:

- If the same username is present in the local (Cisco VTS) database and the TACACS+/LDAP server, then the user will be first authenticated using the 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+/LDAP server.
- Cisco VTS users and groups should be consistent across all the participating TACACS+/LDAP servers.
- If the same username is configured in both local and TACACS+/LDAP 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+/LDAP servers.

- If an AAA user is not assigned to any of the Cisco VTS groups in TACACS+/LDAP server, the user authentication will fail.
- AAA users, even AAA admin users, will not be able to disable AAA, but still will be able to add/delete AAA configuration.
- AAA username with special characters are not supported.
- We recommend that you use the *vts-default* authorization group while adding devices into network inventory. This is a system defined authorization group, available in Cisco VTS. If you are not using the *vts-default* authorization group, you need to ensure that you create an auth group which has AAA user added as the VTC Admin User Name.

The servers are contacted for authentication, based on the priority you set while you configure the servers. If a TACACS+/LDAP 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 19](#) section.

-
- Step 1** Go to **Global Settings > Remote Authentication Settings**.
The Remote Authentication Settings page appears.
- Step 2** Set the Global Time Out. This is to set the connection timeout with the external authentication server. The default is 15 seconds. The number of retries for a connection is set to three.
- Step 3** Use the **Enable Protocol** slider to enable the desired protocol. You must add at least one server for the selected protocol. TACACS+ and LDAP are supported.

To enable TACACS+, use the TACACS + slider.

- Click **Add (+)**. The Configure TACACS + popup window appears.
- Enter the IP Address/Host Name, and the port details.
- Enter the secret key in the Key field. This can have 128 characters.
- Enter the secret key in the Key field. This can have 128 characters.
- Click **Logging** toggle button to enable the accounting.

For 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**.

To enable LDAP, use the LDAP slider.

- Click **Add (+)**.
- Enter the IP Address/Hostname. This is the IP address or Hostname of the LDAP server.
- Enter the Root Domain Name. This is the Base DN of the LDAP server. This is a comma-separated list.
- Enter the User Search Base. This is the LDAP directory used to search for the user identity.
- Enter the Group Search Base. This is to determine the organizational unit that contains the groups.

Note Multiple entries for USER SEARCH BASE is not supported. For example,

```
ou=DevUsers ,dc=arun ,dc=net or  
ou=ProdUsers ,dc=arun ,dc=net or  
ou=Users ,ou=DevUsers ,dc=arun ,dc=net"
```

- f) Enter the Group Search Filter. This is to determine the ObjectClass of the group.
- g) Enter the Port of the LDAP server. By default, this is 389.
- h) Enter the User Search Filter.
- i) Enter the Group Membership Attribute
- j) Check the SSL checkbox to enable SSL.
- k) Click **Save**.

Step 4 Specify the priority with which Cisco VTS should contact the external authentication servers.

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).

Following are the fields that can be found in the log:

- Client IP—Client IP from where the request was made
- Server IP—VTS server IP
- User Name—User who performs the transaction
- Message—The model change in the transaction or the REST API url
- Date/Time—The time when the change was made
- Application Name—VTS (static value)
- Operation Type—Derived from the change, could be CREATE, UPDATE or DELETE
- Status—Success or Error (static value)

Setting up Accounting

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



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

- Step 1** Go to **Administration > Remote Authentication Settings**.
The Remote Authentication Settings page appears.
- Step 2** Use the **Enable Protocol** toggle button to enable the desired protocol. You must add at least one configuration instance for the selected protocol. Currently only TACACS+ is supported.
- Step 3** Click **Add (+)**. The Configure TACACS+ popup window appears.
- Step 4** Enter the IP Address/Host Name, and the port details.
- Note** Cisco VTS supports IPv4 and IPv6 addresses.
- Step 5** Enter the secret key in the Key field. This can have 128 characters.
- Step 6** Click **Logging** toggle button to enable the accounting.
- Step 7** Click **Add**.
- Step 8** Click **Save**.
The logs get saved in the local VTS server and TACACS server.

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

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

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

Creating Authentication Groups

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

You can create authentication groups and assign devices you import into Cisco VTS, to these groups. Authentication groups are used to group devices with the same credentials (that is, usernames and passphrases). Once the authentication 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 authentication group can be used. If the credentials are different for different devices, multiple authentication-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 authentication group with a specific device. The applicable authentication group should be used for corresponding device entry in the CSV file.



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

To create an authentication group:

Step 1 Go to **Global Settings > Authentication Group**.

Step 2 Click **Add (+)** icon. The Add Authentication Group popup window appears.

Enter the following details, and click **Save**:

- Authentication Group Name—The authorization group name.
- VTC Admin 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 authentication group gets added to the Groups table.

To edit an authentication group, select the desired Authentication Group Name check box and click the **Edit** icon.

To delete an authentication group, select the desired Authentication Group Name check box and click the **Delete (X)** icon.

Viewing HA Status

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

Go to **Administration > High Availability**.

You can view the role of Policy Plane, VTC and Control Plane, VTSR.

Note You can only view the table here without performing any action.

The VTC table displays the following details:

- Node ID
- Current Role
- Time Stamp
- Configured Role

The VTSR table displays the following details:

- Node ID
 - Current Role
 - Time Stamp
-

Enabling Auto Southbound Lock

Cisco VTS automatically southbound locks all the devices after a period of inactivity.

**Note**

- The auto southbound lock is enabled by default with 15 minute default period of inactivity.
 - This feature is only about southbound lock. You must manually unlock the devices, if required.
 - The period of inactivity is not checked on a device level. Upon hitting the inactivity period timer, all the devices are southbound locked.
-

To enable/disable auto southbound lock:

- Step 1** Go to **Global Settings > Device Settings**. The Device Settings page appears.
- Step 2** Use the **Auto Southbound Lock** toggle button to enable or disable the auto southbound lock feature. It is enabled by default.
- Step 3** Enter the desired period of inactivity in the **Period of inactivity (in minutes)** field. This can be an integer between 15 and 60. The default is 15 minutes.
-



CHAPTER 4

Administering Cisco VTS

This chapter has the following sections:

- [Setting up Site, on page 25](#)
- [Setting Global Route Reflector, on page 26](#)
- [Registering the Virtual Machine Manager using GUI, on page 26](#)
- [Integrating Cisco VTS with Multiple Virtual Machine Managers, on page 28](#)
- [Backing up the Database in non HA Mode, on page 36](#)
- [Restoring the Database in non HA Mode, on page 37](#)
- [Backing up the Database in HA Mode, on page 38](#)
- [Restoring the Database in HA Mode, on page 39](#)
- [Configuring Syslog for Monitoring Logs, on page 41](#)
- [Collect Data Before Contacting Technical Support, on page 44](#)

Setting up Site

To set up the site:

Step 1 Select the Site from the drop-down list.

Step 2 Go to **Administration > Site Settings**.

Step 3 Enter the **DHCP Server IPv4** address. This can be a valid IPv4 address.

Step 4 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 5 Enter the **AnyCast GW Mac**. This is mandatory. Click ? for information about the format.

Step 6 Choose the VTF Mode you want to use. VTF L2 mode means the Hosts in Host Inventory can have vtf-l2 as virtual switch option. The other option is VTF-VTEP mode which means the Hosts in Host Inventory can have vtf-vtep as the virtual switch option.

Note For OpenStack, VTF L2 mode is supported only on OpenStack Newton.

- VTEP
- L2—If you want to use VTF as an L2 switch. This is the default.

Step 7 Specify the **Default Range for Device VLAN Pools**.

Enter:

- Start— Any integer number between 2 and 4094.
- End— Any integer number between 2 and 4094.

Step 8 Specify details regarding **Multiple VNIs to Multicast Address Mapping**.

Step 9 Click **Submit**.

Setting Global Route Reflector

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

To set the global route reflector:

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

Step 2 Use the toggle switch to choose Global.

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

Step 3 Select the device.

Step 4 Click **Save**.

Registering the Virtual Machine Manager using GUI

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

Cisco VTS allows multiple sites to register and install plugin to the same VMM. We recommend that you do not register more than one site with the same VMM.



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

To do this:

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

Step 2 Click the **Add (+)** button.

The Register VMM page is displayed.

Step 3 Enter the VMM Details:

- Name—Name of the VMM.
- Version —Specify the version from the drop-down. If you choose openstack-newton as the Version in the 'Version' drop-down it displays a question 'Do you want VTS to install VMM plugin components?'

If you choose **No**, enter the VMM ID. You can enter the VMM ID present in the file `/etc/neutron/plugins/ml2/ml2_conf.ini` in the controller machine. By default, **Yes** is chosen.

- Mode—Whether the VMM has been registered as Trusted or Untrusted.
- API Endpoint Details—The fields differ based on the VMM you choose.
 - API Endpoint Details for OpenStack
 - API Protocol:IP Address:Port—VMM service endpoint's IPv4/IPv6 address and port. Make sure you use the same IP address format (IPv4/IPv6) for all IP address fields. Mixed mode is not supported.
 - Keystone Protocol:IP Address:Port—Keystone protocol, IP address and port for OpenStack.
 - Openstack Admin Project—Tenant with Administrator privileges in OpenStack. This can be any tenant with Administrator privileges. Any change to this tenant name, username, and passphrase needs to be updated in Cisco VTS for Multi-VMM operations to work properly.
 - Admin User Name—admin user for the admin project in OpenStack.
 - Admin Passphrase—Password of the admin user.
 - API Endpoint Details for vCenter. This is optional.
 - API Protocol:IP Address:Port—VMM service endpoint's IPv4/IPv6 address and port. Make sure you use the same IP address format (IPv4/IPv6) for all IP address fields. Mixed mode is not supported.
 - Datacenter—The name of the datacenter for which Cisco VTS acts as the controller.
 - Admin User Name—Username of the vCenter VMM.
 - Admin Passphrase —Password of the vCenter VMM.

Step 4 Click **Register**.

After the VMM is registered successfully, the Plugin sections opens up.

Step 5 **For OpenStack:**

Note If you choose **No** for the question 'Do you want VTS to install VMM plugin components?' in VMM Details, the radio button mentioned in **a)** is not displayed. It has only the Neutron Server section. The Add Neutron Server popup has the username and password as optional entries. You can choose not to give those. In that case Cisco VTS only saves the IP address. If you enter the Neutron server details you get an option to Save and Validate the plugin installation.

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

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

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

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

For vCenter:

- a) Enter the following in the Plugin details section:

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

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

To delete a VMM, select the check box corresponding to the VMM you need to delete, and click the delete (X) icon. The VMM is deleted after you click **Delete** in the Confirm Delete dialog box.

Uninstalling the OpenStack Plugin

To uninstall the OpenStack plugin from Neutron server:

- Step 1** Go to **Administration > Virtual Machine Manager**.
- Step 2** Select the specific VMM.
- Step 3** Go to Neutron server plugin section, which shows the list of Neutron servers on which you have installed OpenStack plugin.
- Step 4** Check the checkbox next to the neutron server row, and click on “-“ sign next to it. This uninstalls the plugin.

Integrating Cisco VTS with Multiple Virtual Machine Managers

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



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

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

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

The MVMM feature is supported on:

- vCenter/VMware ESXi 6.0 Update 2 and vCenter/VMware ESXi 6.5 Update 1
- Openstack Newton

VMM Registration Modes

When you register a VMM with Cisco VTS, you can specify whether the VMM is a trusted VMM or an untrusted VMM. For information about registering VMMs, see [Registering the Virtual Machine Manager using GUI, on page 26](#).

Trusted VMM

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

In trusted mode:

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

Cisco VTS supports the following variants in trusted mode:

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

- In case two or more VMMs need to share the same tenant, the operators of the VMMs have to co-ordinate on the names before sending the network information to Cisco VTS.

Untrusted VMM

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

In untrusted mode, Cisco VTS:

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

Cisco VTS supports the following variants in the untrusted mode:

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

Workflows in MVMM mode of Operation

To support the above modes, Cisco VTS:

- Enables you to merge the private L2 networks on different VMMs to create a Multi VMM L2 network. The private L2 networks are created by the individual VMMs and the merge operation is controlled by the Cisco VTS administrator. Cisco VTS' involvement is to coalesce two or more network objects in the VTS database into one. After a successful merge operation, all the networks would be tied together by

a unique L2 VNID. This means that the VLAN allocation scheme to VMM private L2 network remains intact. Even if there are workloads belonging to two different VMMs are placed on the same leaf node, there could be two different VLAN allocations, but the same VNI allocation. Traffic between the two workloads will go through VXLAN encap/decap. The normal mode of VNI allocation in Cisco VTS is 'dynamic' (per admin domain) and is assigned per private L2 network.

- Learns L2 networks from trusted VMMs and publishes these to other untrusted VMM under the control of VTS admin. The Cisco VTS GUI is used to create these networks and publish to untrusted VMMs. Cisco VTS can reuse the VNI that was assigned to the originating VMM and push that to the other untrusted VMMs. If there is no originating VMM (and VTS is the originator), then VNI allocation can happen freely in Cisco VTS.
- Creates a router that can interconnect L2 networks across multiple VMMs.

Merge and Publish Operations

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

- **Merge operation**—Cisco VTS learns and auto-creates a Multi VMM L2 network by combining private L2 networks from multiple trusted VMMs. For a successful merge operation, the tenant name, network name, subnet name, subnet CIDR, and underlay multicast address must match. You can select one/multi/all tenants and networks within a source VMM, and then choose a list of VMMs within which the merge would be in effect. Both the source and destination VMMs need to be trusted.



Note We recommend that you ensure that Shared Networks have unique names across all tenants and all VMMs. This is to avoid ambiguity related to network names, which you might encounter during Multi VMM merge operations.

- **Publish operation**—Cisco VTS initiates the creation of a Multi VMM L2 network on untrusted/trusted VMMs. This decides which network (regardless of the source) needs to be published to a list of VMMs. The VMMs can either be trusted or untrusted. Publish operation automatically pushes tenant and network information on the target VMM.



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

**Important**

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

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

Also you need to uncomment the property vni_ranges and update with suitable range values. For example:

```
# Comma-separated list of <vni_min>:<vni_max> tuples enumerating ranges of
# VXLAN VNI IDs that are available for tenant network allocation (list value)
#vni_ranges =
vni_ranges =10:100
```

To make these configuration take effect, you need to restart the neutron-server.

In case of Openstack Newton, these values are, by default, configured in plugin.ini as above.

Deleting Merged Networks

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

Deleting Published Networks

You cannot delete a network or subnet from VTS after a publish operation. You need to delete the publish operation before you change network or subnet from the source VMM or VTS. If you update from source VMM, the target VMM will not get affected. If you update from the VTS GUI, the update will fail.

All operations on published networks can be initiated only from the VTS GUI. If the network was published from VMM 1 to VMM 2, then VMM 1 can remove the network, but the published network will still exist on VMM 2. If VTS published the network to VMM 2, then if VMM 2 deletes the network, Cisco VTS will not allow to delete the network as long as the publish definition exists.

In order to delete a published network/subnet, you have to first unpublish the network, and then perform the delete operation. To unpublish a published network you need to remove the publish definition before you delete the network. To do this go to the source VMM, view the publish definition and deselect the network which you want to unpublish.

**Note**

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

Performing Merge Operation from VMM

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

-
- Step 1** Go to **Administration > Virtual Machine Manager**.
- Step 2** Select the Source VMM and click on the **Merge** icon under the Multi VMM Operations column.. Merge window opens. The Source VMM is the one from which the operation is initiated. It will be selected and highlighted by default.
- Step 3** Click the radio button corresponding to the Target VMM.
The Tenants from Source column lists the tenants that are available. You can use the **Add (+)** button to add a new tenant name. To add a new tenant, enter the tenant name in the text box, and click the **tick** icon. Click **Delete (X)** to delete.
- Note** This will take effect only after the tenant is actually created.
- You can use the filter to view the available tenants or selected tenants. By default, it shows all tenants.
- You may use the select all button to select all tenants. If you use the select all option, you can set the **Include Tenants that will be created in future automatically** toggle switch to **Yes**.
- Step 4** Select the desired tenant(s).
The Networks from Source column lists the Networks available in the source VMM, for that tenant. You can use the **Add (+)** button to add a new network name. To add a new network, enter the network name in the text box, and click the **tick** icon. Click **Delete (X)** to delete.
- Note** This will take effect only after the network is actually created.
- You can use the filter to view the available networks or selected networks. By default, it shows all networks.
- You may use the select all button to select all networks . If you use the select all option, you can set the **Include Networks that will be created in future automatically** toggle switch to **Yes**. This toggle switch will be set to **Yes**, also if you had set **Include Tenants that will be created in future automatically** toggle switch to **Yes**.
- Step 5** Select the desired networks. Click **Save**.
-

Performing Merge Operation from Tenant

To initiate a merge operation from a tenant:

-
- Step 1** Go to **Tenants> Tenant Management**.
- Step 2** Select the VMM from the drop-down. The tenants for the VMM are displayed.
- Step 3** Click the **Merge** icon under the Multi VMM Operations column for the desired tenant.
- Step 4** Click the radio button to select the Target VMM.
- Step 5** Select the Networks from Source to be merged.
You can use the **Add (+)** button to add a new network name. To add a new network, enter the network name in the text box, and click the **tick** icon. Click **Delete (X)** to delete.
- Note** This will take effect only after the network is actually created.
- You can use the filter to view the available networks or selected networks. By default, it shows all networks.
- You may use the select all button to select all network. If you use the select all option, you can set the **Include Networks that will be created in future automatically** toggle switch to **Yes**.

Step 6 Click **Save**.

Performing Merge Operation from Network

To initiate a merge operation from a network:

- Step 1** Go to **Overlay > Network**.
The Overlay / Network window appears.
- Step 2** Select the source from the Select Source drop-down list.
- Step 3** Select the tenant from the Select Tenant drop-down list.
- Step 4** Click the **Merge** icon for the desired network.
- Step 5** Select the target VMM.
- Step 6** Select the network from the Network from Source column.
- Step 7** Click **Save**.
-

Performing Publish Operation from VMM

To publish from VMM:

- Step 1** Go to **Administration > Virtual Machine Manager**.
- Step 2** Select the Source VMM and click the **Publish** icon under the Multi VMM Operations column.
The Publish window opens. The Source VMM is the one from which the operation is initiated. It will be selected and highlighted by default.
- Step 3** Click the radio button corresponding to the Target VMM.
The Tenants from Source column lists the tenants that are available. You can use the **Add (+)** button to add a new tenant name. To add a new tenant, enter the tenant name in the text box, and click the **tick** icon. Click **Delete (X)** to delete.
- Note** This will take effect only after the tenant is actually created.
- You can use the filter to view the available tenants or selected tenants. By default, it shows all tenants.
- Step 4** Select the desired tenant(s).
The Networks from Source column lists the Networks available in the source VMM, for that tenant. You can use the **Add (+)** button to add a new network name. To add a new network, enter the network name in the text box, and click the **tick** icon. Click **Delete (X)** to delete.
- Note** This will take effect only after the network is actually created.
- You can use the filter to view the available networks or selected networks. By default, it shows all networks.
- You may use the select all button to select all networks.
- Step 5** Select the desired networks. Click **Save**.
-

Performing Publish Operation from Tenant

To initiate a publish operation from Tenant:

-
- Step 1** Go to **Tenants > Tenant Management**.
- Step 2** Select the VMM from the drop-down. The tenants for the VMM are displayed.
- Step 3** Click the **Publish** icon under the Multi VMM Operations column for the desired tenant.
- Step 4** Click the radio button to select the Target VMM.
- Step 5** Select the Networks from Source to be merged.
You can use the **Add (+)** New button to add a new network name. To add a new network, enter the network name in the text box, and click the **tick** icon. Click **Delete (X)** to delete.
- Note** This will take effect only after the network is actually created.
- You can use the filter to view the available networks or selected networks. By default, it shows all networks.
- Step 6** Click **Save**.
-

Performing Publish Operation from Network

To perform a publish operation from Network:

-
- Step 1** Go to **Overlay > Network**.
The Overlay / Network window appears.
- Step 2** Select the source from the Select Source drop-down list.
- Step 3** Select the tenant from the Select Tenant drop-down list.
- Step 4** Click the **Publish** icon for the desired network.
- Step 5** Select the target VMM.
- Step 6** Select the Network from Source column.
- Step 7** Click **Save**.
-

Performing Publish Operation from VTS

To perform a publish operation from VTS:

-
- Step 1** Go to **Administration > Virtual Machine Manager**.
- Step 2** Click the **Publish** icon from VTS icon.
The Publish window opens with the source as VTS. It is selected and highlighted by default.
- Step 3** Click the radio button corresponding to the Target VMM.
The Tenants from Source column lists the tenants that are available. You can use the **Add (+)** button to add a new tenant name. To add a new tenant, enter the tenant name in the text box, and click the **tick** icon.

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

Click **Delete (X)** to delete a tenant you do not want to publish from VTS.

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

Step 4 Select the desired tenant(s).

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

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

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

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

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

Note After you register vCenter as a VMM, and, for the first time, perform a publish operation to publish a tenant and multiple networks to this vCenter VMM, the tenant and networks fail to get published to the VMM. The error next to the policy certificate shows an exception related to SSL handshake. Click the **Retry** button to get the tenant and networks published to the VMM.

Backing up the Database in non HA Mode

To backup the database:

Step 1 Login to VTS VM and switch to root environment.

```
admin@VTS-A:~$ sudo su
[sudo] password for admin:
```

Step 2 Source the VTS environment.

```
root@VTS-A:# source /etc/profile.d/ncs.sh
```

Step 3 Verify VTS status.

```
root@VTS-A:# service nso status
```

```
<snip>
Active: active (running) since Wed 2017-08-09 12:08:13 UTC; 12h ago
<snip>
```

Step 4 Stop VTS.

```
root@VTS-A:# service nso stop
```

Verify whether VTS is stopped.

```
root@VTS-A:# service nso status
<snip>
```

```
Active: inactive (dead) since Wed 2017-08-09 12:18:13 UTC; 12s ago
<snip>
```

Step 5 Take backup.

```
root@VTS-A:# ncs-backup --install-dir /opt/nso
INFO Backup /opt/vts/run/nso/backups/ncs-4.3.0.3@2017-08-10T01:05:25.backup.gz created successfully
```

Verify the backup directory.

```
root@VTS-A:# ls -lrt /opt/vts/run/nso/backups
-rw-r--r-- 1 root root 306914477 Aug 10 01:05 ncs-4.3.0.3@2017-08-10T01:05:25.backup.gz
```

Note You must not rename the backup file. If you rename the backup file, restore will fail. We recommend that you make a note of the backup file name to ensure that the correct file is used while you restore. Also, as a best practice, a copy of the backup file may be stored in a location outside of VTS VM to mitigate possible disk failures.

Step 6 Start VTS.

```
root@VTS-A:# service nso start
```

Verify whether VTS is running.

```
root@VTS-A:# service nso status
<snip> Active: active (running) since Thu 2017-08-10 01:06:33 UTC; 4s ago
<snip>
```

Restoring the Database in non HA Mode

To restore the database:

Step 1 Log in to VTS VM and switch to root environment.

```
admin@VTS-A:~$ sudo su
[sudo] password for admin:
```

Step 2 Source the VTS environment.

```
root@VTS-A:# source /etc/profile.d/ncs.sh
```

Step 3 Verify VTS status.

```
root@VTS-A:# service nso status
<snip>
Active: active (running) since Wed 2017-08-09 12:08:13 UTC; 12h ago
<snip>
```

Step 4 Stop VTS.

```
root@VTS-A:# service nso stop
```

Verify whether VTS is stopped.

```
root@VTS-A:# service nso status
<snip>
Active: inactive (dead) since Wed 2017-08-09 12:18:13 UTC; 12s ago
<snip>
```

Step 5 Perform restore. For example:

```
root@VTS-A:#ncs-backup --install-dir /opt/nso --restore
/opt/vts/run/nso/backups/ncs-4.3.0.3@2017-08-10T01:05:25.backup.gz --non-interactive
INFO Restore completed successfully
```

Step 6 Start VTS.

```
root@VTS-A:# service nso start
```

Verify whether VTS is running.

```
root@VTS-A:# service nso status
<snip> Active: active (running) since Thu 2017-08-10 01:06:33 UTC; 4s ago
<snip>
```

Backing up the Database in HA Mode

Perform the following tasks to backup the database, in HA mode:

Perform these on the Master.

Step 1 Login to VTS Master VM and switch to root environment.

```
admin@VTS-A:~$ sudo su
[sudo] password for admin:
```

Step 2 Verify VTS is in Master mode.

```
root@VTS-A: # crm status
<snip>
Master/Slave Set: ms_vtc_ha [vtc_ha] Masters: [ VTS-A ]
Slaves: [ VTS-B ]
<snip>
```

Step 3 Put VTS in maintenance mode.

```
root@VTS-A:# crm configure property maintenance-mode=true
```

Verify whether VTS is in maintenance mode.

```
root@VTS-A:# crm status
<snip>
Master/Slave Set: ms_vtc_ha [vtc_ha] (unmanaged) vtc_ha (ocf::vts:vtc_ha): Started VTS-B (unmanaged)
  vtc_ha (ocf::vts:vtc_ha): Master VTS-A (unmanaged)
<snip>
```

Step 4 Source the VTS environment.

```
root@VTS-A:# source /etc/profile.d/ncs.sh
```

Step 5 Verify VTS status.

```
root@VTS-A:# service nso status
```

```
<snip>
Active: active (running) since Wed 2017-08-09 12:08:13 UTC; 12h ago
<snip>
```


Step 6 Stop VTS.

```
root@VTS-A:# service nso stop
```

Verify whether VTS is stopped.

```
root@VTS-A:# service nso status
<snip>
Active: inactive (dead) since Wed 2017-08-09 12:18:13 UTC; 12s ago
<snip>
```

Step 7 Take backup.

```
root@VTS-A:# ncs-backup --install-dir /opt/nso
INFO Backup /opt/vts/run/nso/backups/ncs-4.3.0.3@2017-08-10T01:05:25.backup.gz created successfully
```

Verify the backup directory.

```
root@VTS-A:# ls -lrt /opt/vts/run/nso/backups
-rw-r--r-- 1 root root 306914477 Aug 10 01:05 ncs-4.3.0.3@2017-08-10T01:05:25.backup.gz
```

Note You must not rename the backup file. If you rename the backup file, restore will fail. We recommend that you make a note of the backup file name to ensure that the correct file is used while you restore. Also, as a best practice, a copy of the backup file may be stored in a location outside of VTS VM to mitigate possible disk failures.

Step 8 Start VTS.

```
root@VTS-A:# service nso start
```

Verify whether VTS is running.

```
root@VTS-A:# service nso status
<snip> Active: active (running) since Thu 2017-08-10 01:06:33 UTC; 4s ago
<snip>
```

Step 9 Take VTS out of maintenance mode.

```
root@VTS-A:# crm configure property maintenance-mode=false
```

Verify whether VTS is out of maintenance mode.

```
root@VTS-A:# crm status
<snip>
Master/Slave Set: ms_vtc_ha [vtc_ha]
Masters: [ VTS-A ]
Slaves: [ VTS-B ]
<snip>
```

Restoring the Database in HA Mode



Note Restore must be done on the Master. If VTC A was the master while you had taken the backup, and at a later point if you had made VTC B the Master, make VTC A the Master and then perform the restore.

Make sure that both VTC master and VTC slave passwords match with the one in the backup file.

To restore the database in HA mode:

Step 1 Log in to VTS VM and switch to root environment.

```
admin@VTS-A:~$ sudo su
[sudo] password for admin:
```

Step 2 Verify VTS is in Master mode.

```
root@VTS-A: # crm status
<snip>
Master/Slave Set: ms_vtc_ha [vtc_ha] Masters: [ VTS-A ]
Slaves: [ VTS-B ]
<snip>
```

Step 3 Put VTS in maintenance mode.

```
root@VTS-A:# crm configure property maintenance-mode=true
```

Verify whether VTS is in maintenance mode.

```
root@VTS-A:# crm status
<snip>
Master/Slave Set: ms_vtc_ha [vtc_ha] (unmanaged)
vtc_ha (ocf::vts:vtc_ha): Started VTS-B (unmanaged) vtc_ha (ocf::vts:vtc_ha): Master VTS-A (unmanaged)
<snip>
```

Step 4 Source the VTS environment.

```
root@VTS-A:# source /etc/profile.d/ncs.sh
```

Step 5 Verify VTS status.

```
root@VTS-A:# service nso status
<snip>
Active: active (running) since Wed 2017-08-09 12:08:13 UTC; 12h ago
<snip>
```

Step 6 Stop VTS.

```
root@VTS-A:# service nso stop
```

Verify whether VTS is stopped.

```
root@VTS-A:# service nso status
<snip>
Active: inactive (dead) since Wed 2017-08-09 12:18:13 UTC; 12s ago
<snip>
```

Step 7 Perform restore. For example:

```
root@VTS-A:#ncs-backup --install-dir /opt/nso --restore
/opt/vts/run/nso/backups/ncs-4.3.0.3@2017-08-10T01:05:25.backup.gz --non-interactive
INFO Restore completed successfully
```

Step 8 Start VTS.

```
root@VTS-A:# service nso start
```

Verify whether VTS is running.

```
root@VTS-A:# service nso status
<snip> Active: active (running) since Thu 2017-08-10 01:06:33 UTC; 4s ago
<snip>
```

Step 9 Take VTS out of maintenance mode.

```
root@VTS-A:# crm configure property maintenance-mode=false
```

Verify whether VTS is out of maintenance mode.

```
root@VTS-A:# crm status
<snip>
Master/Slave Set: ms_vtc_ha [vtc_ha]
Masters: [ VTS-A ]
Slaves: [ VTS-B ]
<snip>
```

Configuring Syslog for Monitoring Logs

From VTC, VTSR, and docker, you can send the logs to rsyslog server and also syslog-ng server. VTSR supports syslog-ng server only on management network. From VTF, it has to be sent to rsyslog server only. You can configure as many rsyslog or syslog-ng server as you require. Cisco VTS supports both TCP and UDP protocols. It also supports and IPv4 / IPv6 addresses for syslog configuration. You can specify multiple syslog servers separated by commas. Make sure you specify the port and protocols also using commas.

- Step 1** Install and configure syslog-ng server on Ubuntu.
- Step 2** Install and configure rsyslog server.
- Step 3** Configure the *ansible all.yaml* file with Syslog server from VTC. For example:

```
#vi /opt/vts/lib/ansible/playbooks/group_vars/all.yaml
```

VPFA_LOG_FILES:

```
CRITICAL: "/var/log/vpfa/vpfa_server_critical.log"
ERROR: "/var/log/vpfa/vpfa_server_errors.log"
WARN: "/var/log/vpfa/vpfa_server_warning.log"
INFO: "/var/log/vpfa/vpfa_server_informational.log"
RSYSLOG_UDP_SERVER_PORT: 514
RSYSLOG_TCP_SERVER_PORT: 515
# Add list items of syslog servers and protocol for each logging level as required
# In the optional PROTOCOL: field use 'TCP' or 'UDP'. Defaults to UDP if not specified
CRITICAL_SERVERS:
- SERVER: "2001:420:10e:2015::202"
PROTOCOL:UDP
- SERVER: "172.23.92.151"
PROTOCOL:UDP
ERROR_SERVERS:
- SERVER: "2001:420:10e:2015::202"
PROTOCOL:UDP
- SERVER: "172.23.92.151"
PROTOCOL:UDP
WARN_SERVERS:
- SERVER: "2001:420:10e:2015::202"
PROTOCOL:UDP
- SERVER: "172.23.92.151"
PROTOCOL:UDP
INFO_SERVERS:
- SERVER: "2001:420:10e:2015::202"
PROTOCOL:UDP
- SERVER: "172.23.92.151"
PROTOCOL:UDP
```

- Step 4** Install the VTSR and complete the registration, then configure syslog from VTC. To do this, copy `/opt/vts/etc/LogConfig.ini.tmpl` to `/opt/vts/etc/LogConfig.ini` and update the new file with the Syslog server host and port, and log level to be set, based on which the corresponding logs from VTC will be sent to the configured external Syslog server. Also the comma separated paths of the log files is monitored for sending the logs to the Syslog.

```
[SyslogSection]
#Provide a comma separated list of syslog server ip, port and protocol
syslog.server=127.0.0.1,2001:0db8:85a3:0000:0000:8a2e:0370:7334
syslog.port=514,514
syslog.protocol=udp,udp
[LogSection]
#Supported log levels EMERGENCY, ALERT, CRITICAL, ERROR, WARNING, NOTICE, INFORMATIONAL, DEBUG
log.level=WARNING
#List of log files to be captured seperated by comma
log.files=/opt/vts/log/nso/ncs-java-vm.log,/opt/vts/log/nso/ncs.log,/opt/vts/log/tomcat/vts_wap.log

[SyslogSection]
#Provide a comma separated list of syslog server ip, port and protocol
syslog.server=2001:420:10e:2015::202,172.23.92.151
syslog.port= 514,515
syslog.protocol= udp,tcp
[LogSection]
#Supported log levels EMERGENCY, ALERT, CRITICAL, ERROR, WARNING, NOTICE, INFORMATIONAL, DEBUG
log.level= INFORMATIONAL
#List of log files to be captured seperated by comma
log.files=/opt/vts/log/nso/ncs-java-vm.log,/opt/vts/log/nso/ncs.log,/opt/vts/log/tomcat/vts_wap.log
```

Note Log levels is set to Warning by default.

- Step 5** As root user, run the python script `ConfigureSyslog.py` which will read the config ini file, and push the necessary configuration on VTC and VTSRs and also automatically start the filebeat and logstash services.

```
# sudo su -
#ConfigureSyslog.py

root@vts14:~# ConfigureSyslog.py
2017-11-14 20:58:59,801 - SyslogConfig - INFO - Start reading configs.
2017-11-14 20:58:59,801 - SyslogConfig - INFO - Syslog_server = 2001:420:10e:2015::202,172.23.92.151
2017-11-14 20:58:59,801 - SyslogConfig - INFO - syslog server = 2001:420:10e:2015::202
2017-11-14 20:58:59,801 - SyslogConfig - INFO - syslog server = 172.23.92.151
2017-11-14 20:58:59,802 - SyslogConfig - INFO - Syslog Servers provided are valid address
2017-11-14 20:58:59,802 - SyslogConfig - INFO - Updating file_beat config
2017-11-14 20:58:59,802 - SyslogConfig - INFO - Created the main filebeat yml file.
2017-11-14 20:58:59,802 - SyslogConfig - INFO - Got the list of files to be monitored for logging.
2017-11-14 20:58:59,802 - SyslogConfig - INFO - Updating file_beat config with input values
2017-11-14 20:58:59,802 - SyslogConfig - INFO - Updating logstash config
2017-11-14 20:58:59,802 - SyslogConfig - INFO - syslog server = 2001:420:10e:2015::202
2017-11-14 20:58:59,802 - SyslogConfig - INFO - syslog server = 172.23.92.151
2017-11-14 20:58:59,802 - SyslogConfig - INFO - Replaced logstash config with input values
2017-11-14 20:58:59,803 - SyslogConfig - INFO - Restarting logstash service
2017-11-14 20:59:06,750 - SyslogConfig - INFO - Restarting filebeat service
2017-11-14 20:59:07,054 - SyslogConfig - INFO - Configuring syslog information
2017-11-14 20:59:07,071 - SyslogConfig - INFO - Configuring syslog information on vtsr01
2017-11-14 20:59:08,151 - SyslogConfig - INFO - Successfully configured syslog server details
```

- Step 6** For HA deployments of VTC, execute steps 4 and 5 on the other node. This ensures that the filebeat and logstash services get started automatically on both the nodes.

- Step 7** Once the configurations are pushed to VTSR and Docker, spawn the VTF from UI.

Example of Config pushed:

```
Configs pushed on VTSR:
logging 172.23.92.151 vrf default port 515 //This is for TCP Port 515
```

```

logging 2001:420:10e:2015::202 vrf default // This is for UDP 515 Port
logging hostnameprefix vtsr01

Configs pushed on Docker:
syslog host-name-prefix vtsr01
syslog host-server vrfs vrf default
ipv6s ipv6 2001:420:10e:2015::202
ipv6-severity-port
!
ipv4s ipv4 172.23.92.151
ipv4-severity-port port 515
!

vtsr-config syslog syslog-servers host-name-prefix vtsr01
vtsr-config syslog syslog-servers syslog-server 172.23.92.151
port 515
severity informational
proto tcp
!
vtsr-config syslog syslog-servers syslog-server 2001:420:10e:2015::202
severity informational
!
vtsr-config vtfs vtf VTF39
mac 00:50:56:88:47:54
ip 42.42.42.39
mode vm-mode
!

```

- Note**
- Only for VTF—To disable rsyslog configuration add the following attribute to the inventory file:
`configure_rsyslog_client=False`
 - There is no uninstall script to cleanup ConfigureSyslog details, or disable option from VTS CLI to clear syslog config. The only way is specify to syslog server as 0.0.0.0 in LogConfig.ini and reconfigure it.

Troubleshooting Syslog Issues

- Step 1** Filebeat configuration files are in `/etc/filebeat/filebeat.yml` and `/etc/filebeat/filebeat_config.yml`. The logs for the filebeat are at the location `/var/log/filebeat`. The log level and files to monitor are populated in the `filebeat_config.yml` file.
- Step 2** If there is a need to start/stop/restart filebeat, then do the following:
- Example:**
- ```
service filebeat start|stop|restart
```
- Step 3** Logstash configuration files are in `/etc/logstash/conf.d`. The syslog configuration is in the file `conf.d/logstash-beatconfig.conf`. Make sure that the syslog info provided in the `ini` file is populated in this logstash conf file. Also the log files for the logstash service are in `/var/log/logstash`.
- Step 4** If there is a need to start / stop / restart logstash then do the following:
- Example:**
- ```
service logstash start|stop
```
- Step 5** If you encounter the below error while running the script `ConfigureSyslog.py`, then you can workaround this by setting the path for python and then running the script again.

Error: File "/opt/vts/lib/python/vtsLogging/ConfigureSyslog.py", line 9, in <module> import ncs

Example:

```
# export PYTHONPATH=/opt/nso/current/src/ncs/pyapi:/opt/vts/lib/python
# ConfigureSyslog.py
```

Collect Data Before Contacting Technical Support

At some point, you might need to contact your technical support representative or Cisco TAC for some additional assistance. This section outlines the steps that you can perform before you contact your next level of support or before you submit the issue to your Product Development team to reduce the amount of time spent resolving the issue.

This process ensures you to identify the root cause of the problem and address them effectively with a little turnaround time in the RCA process.

Table 1: Checklist

Check For..	Description
VTS Release Version	Provide the Software version being used.
Issue	Provide the exact issue.
Issue Type	Identify whether the issue is pertaining to Functional or Performance or Enhancement or Query.
Component	Provide the component that have issues. For example, UI or Install or Upgrade or Templates and so on.
Detailed Scenario	It is recommended to provide the scenarios in detailed steps. Also attach, UI Screenshots and/or CLI screen capture that shows all above steps in the order.
If the issue is isolated to VTS alone when there are some custom development work on top of VTS	Specify Yes or No or Not applicable
Logs pertaining to one of the following issues:	
Show tech	Provide the Show Tech Output as an attachment or provide a download location. For examples, refer the section "Using Show_tech_support-t-a Command"
UI	For UI specific Issues collect UI Screenshots and describe the UI flow (if the UI flow is different from the Detailed scenario).

Check For..	Description
Install/Upgrade	<p>For any Install or Upgrade Issue provide the following:</p> <ul style="list-style-type: none"> • Source Release • Templates (Template Migration issue): Name of the Template and attach the actual template. • Install or Upgrade log files. Refer the Upgrade section.
Template feature	<p>For Template feature issues provide the following:</p> <ul style="list-style-type: none"> • Template Type (Device/L3/L2) • Actual scenario and the template having the issue. Ensure that "show tech" output includes the CDB.
VMM Integration	<p>For VMM Integration issues provide the following:</p> <ul style="list-style-type: none"> • VMM Type. For example, OpenStack OSPD/Red Hat OpenStack/VMWare. • Plug-in logs. Refer the OSPD ML2 Plugin. • Describe the VMM use case in detail.
VTF	<p>For VTF Issues, provide the following::</p> <ul style="list-style-type: none"> • Compute Information (Type of h/w, Model, manufacturer, Compute Spec). • Deployment Type such as OVS or Host Agent and so on. • Types of line cards that are used on Compute. • VTSR logs: <ul style="list-style-type: none"> Login to vtsr docker Logs dir in the vtsr docker: /var/log/vtsr • VPP logs: <ul style="list-style-type: none"> Login to VTF Logs dir: /var/log/vpfa
Performance	<p>For performance issues provide the following:</p> <ul style="list-style-type: none"> • Describe the Deployment scale numbers such as #of TORs, #of computes, # of ports, # of tenants, # of templates as relevant to the issue reported.

Troubleshooting Command Examples and Log Files

Use the following examples to troubleshoot and capture log files:

Using Show_tech_support-t -a

The **Show_tech_Support** CLI command allows administrator to capture all log files and CDB Backup from a given VTC VM. Ensure to attach this CLI output file to any case opened against VTS as this information is very critical to triage issues found in the VTS.

```
root@VTC1:/opt/vts/log/nso# sudo su -
root@VTC1:~# show_tech_support -t -a
2019-01-24 15:58:23,542 - __main__ - WARNING - The execution of show_tech_support may take
several minutes, depending on the state of VTS.
root@VTC1:~# ls
VTS-2.6.2.1-40-2019-01-22--12-21-55.tar.bz2
```



Note If you do not specify '-ta' parameters after show_tech_support CLI, then there will not be a CDB backup file in the .tar file.

Example — Generic Data Collection Data

```
#VTS version
vts_version

cd ~
sudo -i
show_tech_support -t -a
pwd
ls -l
#chmod 777 VTS-2.6.6*.tar.bz2
#Transfer the file
scp user@IP:/<path>/<file>.tar.bz2 .

#Login to VTC and remove the file afterwards
rm <...>
```

Data Collection with CDB

Below is the procedure to Backup NSO on a VTS VM:

```
root@vts291-116:~# /opt/vts/bin/ncs-backup --install-dir /opt/nso
INFO Backup /opt/vts/run/nso/backups/ncs-4.6.1.2@2018-09-08T09:07:44.backup.gz created
successfully
```

Show_tech_support Optional Input Parameters:

- (optional) **parameter --all-device-configurations**—Collects NCS configuration of all managed devices.
- (Optional) **parameter --device-configurations**—Regular expression that defines set of devices whose NCS configuration must be collected.

The selected devices are those whose name contains a substring that matches the regex. To collect all devices use "."



Note Double backslashes are not supported.

- (Optional) **parameter --output-directory**—Directory where the output is generated.
- (Optional) **parameter --text-based-cdb-state**—Includes state from CDB. Intended for internal use only.
- (Optional) **parameter --no-cdb-backup**—Do not include a CDB backup. Intended for internal use only.

Problem-Dependent Data Collection Examples

Service/Template/Device —NSO Logs

```
#NSO Java VM This is where most of the logging for VTS related operations will take place)
/opt/vts/log/nso/ncs-java-vm.log

cd /opt/vts/log/nso/trace
#tail -f <trace file of the device>
#NSO platform:
/opt/vts/log/nso/ncs.log
/opt/vts/log/nso/devel.log

#NED

cd /opt/vts/log/nso/trace
#tail -f <trace file of the device
```

Service/Template/Device—Configuration and Status

```
ncs_cli -u admin
#Check devices name/lock status
show devices list
devices device <name> check-sync
#Check Device CDB:
show full-configuration devices device <device_name> config
```

Template Configuration

```
#Export All Templates:
show running config templates template | save /tmp/XYZ.txt
#Export Specific Template:
show running-config templates template <...> | save /tmp/AB.tx
```

VTS Application

```
#Main log:
/opt/vts/log/tomcat/vts_wap.log
#Log of Northbound REST API calls (HTTP Headers only)
/opt/vts/log/tomcat/localhost.access
```

VTS Application-Additional Logs

```
#Log of SNMP Notifications:
snmp.log
#Audit log shows is users connecting to VTC and would also indicate if a user is connecting
with the wrong password)
audit.log
var/vts/log
/vts-discovery.log (gives information on the LLDP neighbor topology discovery)
/vts-discovery.errors (logs errors of issues encountered during topology discovery)
/setup.log (basic VTC systems information populated during setup)
/hostagent-install.log (logs relevant information about ML3 and ML2 plugin installs performed
by VTC)
```

TACACS

For Tacacs+ server authentication logs, look at `/opt/vts/log/nso/external_authentication.log` `vts-accounting.log`.
For example:

```
root@VTC1:/opt/vts/log/nso# ls
audit.log daemon_err.log devel.log localhost:8080.access ncserr.log.idx ncs-java-vm.log
netconf.log trace vts-accounting.log.lck
ciscoj daemon.log external_authentication.log ncserr.log.1 ncserr.log.siz ncs.log
snmp.log vts-accounting.log webui-browser.log
root@VTC1:/opt/vts/log/nso# vi external_authentication.log
root@VTC1:/opt/vts/log/nso# vi vts-accounting.log
```

OSPD ML2 Plugin

For VTS OSPD ML2 Plugin logs:

```
/opt/vts/log
and look for <Contoller_ansible_logger.log>
```

Upgrade

VTS upgrade logs are located at:

```
/opt/vts/log/vts-upgrade/<upgrade instance>/logs/upgrade.log file.
```

For VTS Upgrade, granular details of upgrade are found in the `ncs-java-vm.logs`:

```
/opt/vts/log/vts-upgrade/<upgrade instance>/logs/<version>/log/nso/ncs-java-vm.log
```

Re-deploy of services during upgrade and post upgrade logs are found in:

```
/opt/vts/log/ncs-java-vm.log
```



CHAPTER 5

Monitoring Cisco VTS

The following sections provide details about Monit and collectd features that enable you to monitor Cisco VTS.



Note Monitoring features (collectd and Monit) are not supported for Data Plane (VTF) when VTF is in VTEP mode, or vCenter (VM mode).

Monit is a process monitoring tool. It collects and displays metrics related to memory consumption, CPU usage, swap information processes, file system, in a dashboard. For each process in each component, the dashboard shows status, uptime, CPU utilization, Memory, and Read/Write bytes on the disk. It is packaged as part of VTC, VTSR, and VTF, and will get installed as part of the respective installations. See [Monitoring Cisco VTS Infrastructure using Monit, on page 49](#) for details.

collectd is a system statistics daemon which collects system and application performance metrics periodically and provides mechanisms to store the values. See <https://collectd.org/> for details about collectd. Cisco VTS installation installs collectd. On a new Cisco VTS installation, collectd plugins are preconfigured to load and run. These plugins have their configurations already saved in VTC. collectd collects various statistics related to VTC and VTF, which includes CPU, memory, number tenants, networks routers etc, based on the plugins that you enable. The metrics can be sent to an external location you specify, in JSON format, which can be used for further processing. See [Monitoring Cisco VTS Infrastructure using collectd, on page 54](#) for details about how collectd is used in Cisco VTS.

Cisco VTS displays the Policy Plane monitoring details in the **Monitor** page, and the Data Plane and Control Plane monitoring details in the **Monitor Site** page.

- [Monitoring Cisco VTS Infrastructure using Monit, on page 49](#)
- [Monitoring Cisco VTS Infrastructure using collectd, on page 54](#)

Monitoring Cisco VTS Infrastructure using Monit

Monit is used to collect status of all the services that are running on the VTS and VTSR VMs as well as compute nodes running VTF. The Policy Plane monitoring details are displayed in a dashboard under **Monitor**. The Data Plane and Control Plane monitoring details are displayed under **Monitor Site**.

The intervals when Monit will collect metrics for each of the Cisco VTS components are:

- VTC—60 seconds

- VTSR—As entered in the `vtsr_template.cfg` file while installing VTSR.
- VTF—30 seconds



Note For VTC and VTF, the intervals are fixed and you cannot change these values.

This means that, when an event happens, the Cisco VTS UI will show the appropriate status only after this interval has passed.

About Monit Username and Passwords

- For Policy Plane (VTC)—Monit is packaged as part of VTC. Monit, when installed, will have a default username and password. The default monit credentials for VTC when installed are:
 - Username—`monit-ro`
 - Password—`monit-ro`

Upon logging into Cisco VTS for the first time after installation, the admin needs to enter this default username and password in the VTC GUI at **Global Settings > Monitoring Settings > Monit Settings** to view the VTC monitoring information in the Monit Dashboard.

- For Control Plane (VTSR)—Monit is packaged as part of VTSR.

The default credentials for VTSR is entered while installing VTSR, by modifying the below two properties in `vtsr_template.cfg`:

```
#VTSR_OPER_USERNAME="monit-ro-oper"
# Password needs an encrypted value
# Example : "openssl passwd -1 -salt <salt-string> <password>"
#VTSR_OPER_PASSWORD="$1$cisco$b88M8bkCN2ZpXgEEc2sG9/"
```



Note See the *Installing VTSR* section in the *Cisco VTS Installation Guide* for details.

Upon logging into Cisco VTS for the first time after installation, the admin needs to enter this default username and password in the VTC GUI at **Administration > Site Monitoring Settings > Monit Settings** to view the VTSR monitoring information in the Monit Dashboard.

- For Data Plane (VTF)—Monit is packaged as part of VTF.

The default monit credentials for VTFs are:

- Username—`monit-ro`
- Password—`monit-ro`

Upon logging into Cisco VTS for the first time after installation, the admin needs to enter this default username and password in the VTC GUI at **Administration > Site Monitoring Settings > Monit Settings** to view the VTF monitoring information in the Monit Dashboard.

- For Data Plane (VTF) deployed via OpenStack Platform Director (OSPD)—When VTF is installed via OSPD, the Monit related properties need to be updated in the `neutron-cisco-vts.yaml` file. See the *Installing Cisco VTS 2.6.6 Components in OpenStack using Red Hat Enterprise Linux OpenStack Director* document for details.

Changing Monit Password Subsequently

To change Monit password for VTC, VTF, or VTSR, the admin needs to run the following script:

```
/opt/vts/bin/update_monit_credentials.sh.
```

- To change password for Policy Plane, run:

```
/opt/vts/bin/update_monit_credentials.sh policy-plane <monit-username> <monit-password>
```

- To change password for Data Plane, run:

```
/opt/vts/bin/update_monit_credentials.sh data-plane <monit-username> <monit-password>  
<site id>
```

- To change password for Control Plane, run:

```
/opt/vts/bin/update_monit_credentials.sh control-plane <monit-username> <monit-password>  
<site id>
```

For VTC, the username cannot be changed. The default is monit-ro.

In HA mode, the password has to be changed on the Master.

To change Monit password for VTF deployed via OSPD, update the following properties under Monit-Configuration section in the neutron-cisco-vts.yaml file. See the *Installing Cisco VTS 2.6.6 Components in OpenStack using Red Hat Enterprise Linux OpenStack Director* document for details.

Metrics Collected using Monit-D

Monit runs in VTC Master, VTC Slave, VTSR Master, VTSR slave, and all the VTFs. The VTS UI Monitoring page displays the monitoring status, on-demand.

Following are the intervals when Monit will collect metrics for each of the planes. This means that, when an event happens, the Cisco VTS UI displays the appropriate status only after this interval has passed.

- VTC—60 seconds.
- VTSR—Based on what is configured in the template.cfg file during installation.
- VTF—30 seconds.

Metrics Collected for VTC

The following metrics are collected:

- Process: The following VTS processes can be monitored.
 - Corosync
 - Pacemaker
 - Filebeat
 - Logstash
 - Collectd
 - Monit
 - Nso

- Ntpd
- Sshd
- Solr
- Nginx
- Nodejs
- Tomcat
- Vtswab

- File System
 - Root
 - Boot

- Network
 - Management
 - Underlay

Metrics Collected for VTSR

The following metrics are collected:

- Process—The following VTSR processes can be monitored.
 - Confd
 - Rc
 - Dl
 - cfg_dl
 - redis
 - stunnel
 - pacemaker
 - corosync
 - logstash
 - monit
 - filebrat-god
 - filebeat

Setting up Monit Credentials

You must set up the credentials for the Policy Plane (VTC), Control Plane (VTSR), and Data Plane (VTF), to enable you to access the metrics collected by Monit via the Cisco VTS UI. See [Monitoring Cisco VTS Infrastructure using Monit, on page 49](#) for details about Monit credentials for VTC, VTSR, and VTF.

Setting up Policy Plane Credentials

To set up the credentials to enable accessing Policy Plane metrics:

Step 1 Go to **Global Settings > Monitoring Settings**.

Step 2 Click the Policy Plane tab.

Step 3 Enter the credentials.

Note The username is monit-ro by default, and cannot be changed. The Password should match the one that was configured during Monit setup.

Setting up Control Plane Credentials

To set up the credentials to enable accessing Control Plane metrics:

Step 1 Go to **Administration > Site Monitoring Settings**.

Step 2 Click the Control Plane tab.

Step 3 Enter the Username and Password that was set during Monit setup.

Setting up Data Plane Credentials

To set up the credentials to enable accessing Data Plane metrics:

Step 1 Go to **Administration > Site Monitoring Settings**.

Step 2 Click the Data Plane tab.

Step 3 Enter the Username and Password that was set during Monit setup.

Viewing Metrics Collected by Monit

The Policy Plane metrics collected by Monit is displayed in Cisco VTS > Monitor UI. The Data Plane and Control Plane metrics are displayed in Cisco VTS > Monitor Sites UI.

Viewing Policy Plane Metrics

To view the metrics collected for Policy Plane (VTC):

Step 1 Go to **Monitor**.

Step 2 Click the Policy Plane tab.

Note Monitoring information is displayed only after you complete the **Monitoring Settings** under **Global Settings**.

Viewing Control Plane Metrics

To view the metrics collected for Control Plane (VTSR):

Step 1 Go to **Monitor**.

Step 2 Click the Control Plane tab.

Note Monitoring information is displayed only after you complete the **Site Monitoring Settings** under **Administration**.

Viewing Data Plane Metrics

To view the metrics collected for Data Plane (VTF):

Step 1 Go to **Monitor**.

Step 2 Click the Data Plane tab.

Note Monitoring information is displayed only after you complete the **Site Monitoring Settings** under **Administration**.

Step 3 Select the VTF IP address from the drop-down list.

Note Monitoring information is displayed only after you complete the **Site Monitoring Settings** under **Administration**. In an HA set up, you must specify the static route to reach the VTF on both Master and Slave. This is to ensure that VTF statistics is displayed even when a Master VTC is switched over to Slave VTC.

Monitoring Cisco VTS Infrastructure using collectd

Cisco VTS embeds collectd to collect metrics and statistics of VTS components. Currently, collectd is embedded as part of VTC and VTFs.

collectd is a system statistics daemon which collects system and application performance metrics periodically and provides mechanisms to store the values.

collectd starts running upon Cisco VTS installation and is configured for a default collection interval of 120 seconds. In an HA setup, collectd runs on both master and slave VTCs. At the configured interval, it will invoke the Input plugins.

The `write_log` output plugin logs the metrics in the local VTC at `/opt/vts/log/collectd/metrics.log`, and `write_http` plugin (when configured) pushes the metrics to an external location.

List of Plugins for VTC

1. Py Custom Plugin—Gets periodically called by `collectd`; calls the UI backend API to get the stats in JSON format.
2. CPU—Inbuilt plugin in `collectd`
3. Memory—Inbuilt plugin in `collectd`
4. Load—Inbuilt plugin in `collectd`
5. Interface—Inbuilt plugin in `collectd`
6. Disk—Inbuilt plugin in `collectd`
7. `log_file`; `write_log`—Used to log the metrics locally.
8. `write_http`—You must configure this plugin if you want to forward the metrics in JSON format to the centralized `collectd` server.



Note Except `write_http`, all plugins are available by default upon installation.

Plugin Configuration for VTC and VTF

`collectd` can be configured for Policy Plane (VTC) from the **Global Settings > Monitoring Settings > Collectd Settings** page. `collectd` configuration for Data Plane can be done at **Administration > Site Monitoring Settings**.

Any change to the Policy Plane `collectd` settings and plugins would take a maximum of three minutes to get reflected in the VTC `collectd` process. Any change to the Data Plane `Collectd` settings and plugins is done immediately but would depend on the number of VTFs to get updated.

Following `collectd` parameters and plugin details can be set up:

1. `collectd Interval`—This is interval for `collect-d` to collect the metrics. This is per plane. The default is 120 seconds.
2. `Enable/Disable collectd`—This toggle switch will help you to enable/disable `collectd` for all the `collectd` plugins within that Plane. For VTC, in an HA setup, it would disable/enable `collectd` process for both master and slave. For VTFs, it would disable/enable `collectd` process for all the VTFs.
3. `Plugin Configuration`—You can configure any `collectd` plugin as required.
 - a. `Plugin Name`—Name of the plugin.

See [collectd Plugin Configuration, on page 209](#) for the default plugin configurations Cisco VTS supports.

Following are the plugins that are supported:

- CPU
- Python
- `log_logstash`

- write_log
- logfile—The log file location: /opt/vts/log/collectd/metrics.log.



Note For VTF, you have to configure the log_file plugin. This is required to write the output to a specific log file. Otherwise it is sent to /var/log/messages.

- Interface
- Memory
- Load
- write_http—You must configure this plugin if you want to forward the metrics in JSON format to the centralized collectd server.

b. Plugin Config—In xml format.

```
LoadPlugin {plugin-name}
  <Plugin {plugin-name}>
    {parameters of the plugin}
  </Plugin>
```

Sample write_http plugin config:

```
<LoadPlugin write_http>
  FlushInterval 10
</LoadPlugin>
<Plugin write_http>
  <Node "example">
    URL "http://10.10.10.10/centralized-collectd"
    Format "JSON"
    BufferSize 10240
  </Node>
</Plugin>
```

FlushInterval ensures that the payloads are sent to the external server at predefined intervals. If the FlushInterval is lesser than collectd-interval, then collectd interval would take precedence, since there is nothing to send when the read-interval has not lapsed.

If the FlushInterval is greater than collect-interval, then FlushInterval will take precedence and payloads will be sent as per the FlushInterval, provided the buffer does not get full.



Note For the FlushInterval to be honored, we recommend that you keep a bigger buffersize (40960). The buffersize can depend on the statistics collected, which depends on scale.

- c. Enable/Disable Plugin—**You can enable or disable the plugins via the Cisco VTS UI.

Plugin Configuration for VTF Deployed via OSPD

For changing the VTF collectd plugin configuration while deploying via OSPD, you need to modify it in the neutron-cisco-vts.yaml file. Any change or addition to the plugins would need a change in this yaml file (under the Collectd Agent Configuration section). See the *Installing Cisco VTS 2.6.6 Components in OpenStack using Red Hat Enterprise Linux OpenStack Director* document for details.

Metrics Collected by collectd

- Default Metrics:
 - CPU
 - Memory
 - Load
 - Disk
 - Interface
 - Python
 - write_log
- VTC Statistics:
 - Total Number of Tenants
 - Total Number of Tenants per VTEP (HW and SW)
 - Total Number of Networks
 - Total Number of Networks per Tenant
 - Total Number of Networks per VTEP (HW and SW)
 - Total Number of Routers
 - Total Number of Routers per Tenant
 - Total Number of Router per VTEP (HW and SW)
 - Total Number of Baremetal per VTEP
 - Total Number of Shared Networks
 - Total Number of Hosts/Servers
 - Total Number of H/W Vteps
 - Total Number of S/W Vteps (VTFs)

Metrics Sent

The following information from the VTS is sent:



Note See [collectd Output JSON Examples, on page 213](#) for output file examples.

1. Master or Slave
2. IP or Hostname of the VTC
3. Stats Category—For example, number of tenants

4. Stats Sub-Category—For example, the VTEP name if we have tenants per VTEP
5. Count—Count of tenants

Setting up collectd Plugins

The **Global Settings > Monitoring Settings > Collectd Settings** page displays the Plugin Types and also shows whether the plugin is enabled or disabled, for the Policy Plane. The **Administration > Site Monitoring Settings > Collectd Settings** page displays the Plugin Types and also shows whether the plugin is enabled or disabled, for the Data Plane.

The Manage Settings and Plugin option allows you to manage the collectd settings and add/remove, enable/disable, and edit plugins, for both Policy Plane and Data Plane. See [Monitoring Cisco VTS Infrastructure using collectd, on page 54](#) for details about usage of collectd in Cisco VTS.

Setting up Policy Plane Plugins

To set up collectd parameters for Policy Plane:

Step 1 Go to **Global Settings > Monitoring Settings > Collectd Settings**.

Step 2 Click the **Policy Plane** tab.

Step 3 Click **Manage Settings and Plugin**.

You can specify the following:

- Collection Interval—The collection interval for metrics collection. By default, this is 120 seconds. This can be between 10 and 1800 seconds.
- Enable Collect D—Use the toggle switch to enable or disable collectd metrics collection. By default, this is set to Yes.
- Plugin Type—Choose from the list of plugins packaged with the collectd server.
- Plugin Config—Enter or edit the configuration for the plugin type you selected.
- Enable Plugin—Use the toggle switch to enable or disable the selected plugin. By default, the selected plugin is enabled.

Use the + button to add plugins. Use the - button to remove the plugin.

Step 4 Click **Save**.

Setting up Data Plane Plugins

To set up collectd parameters for Data Plane:

Step 1 Go to **Administration > Site Monitoring Settings > Collectd Settings**.

Step 2 Click the **Control Plane** tab.

Step 3 Click **Manage Settings and Plugin**.

You can specify the following:

- **Collection Interval**—The collection interval for metrics collection. By default, this is 120 seconds. This can be between 10 and 1800 seconds.
- **Enable Collect D**—Use the toggle switch to enable or disable collectd metrics collection. By default, this is set to Yes.
- **Plugin Type**—Choose from the list of plugins packaged with the collectd server.
- **Plugin Config**—Enter or edit the configuration for the plugin type you selected.
- **Enable Plugin**—Use the toggle switch to enable or disable the selected plugin. By default, the selected plugin is enabled.

Use the + button to add plugins. Use the - button to remove the plugin.

Step 4 Click **Save**.



CHAPTER 6

Managing Inventory

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

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



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

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

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

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



Note Cisco VTS device discovery is performed over secure ports and protocols. You must make sure that the Nexus OS devices are reachable through HTTPS (443). In Nexus 7000 series devices, https is disabled by default. You must make sure it is enabled on port 443.

To enable secure communication for IOS-XR devices over SSH, you need to have the SSH enabled on the devices. Day Zero Configuration for Cisco ASR 9000 has to be updated to support this. (See Day Zero Configuration Examples document for details).

This chapter has the following sections:

- [Importing Inventory using CSV File, on page 62](#)
- [Performing Auto Discovery, on page 65](#)
- [Viewing the Network Topology, on page 73](#)
- [Viewing Network Inventory, on page 74](#)
- [Viewing Host Inventory, on page 77](#)
- [Viewing the VTSR to VTF Mapping, on page 80](#)

- [SR-IOV Support, on page 80](#)
- [Migrating from vPC to ESI, on page 81](#)
- [Redeploying Device Inventory, on page 83](#)
- [Enabling Static Multi Homing, on page 84](#)
- [Administrative State for Devices, on page 85](#)

Importing Inventory using CSV File

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

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



Note

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

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

The CSV file has the following fields:

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

- `interface-name`—Physical port connectivity (interface of the connected device)

**Note**

- Do not use `a/b` for interface names instead, use a letter followed by a number. For example, Excel converts `5/28` to `May-28`.
- Use a text editor like `notepad++`, instead of Excel to append new items to the CSV file.

Prior to CSV import, you can update the existing interface names in `a/b` format by adding a letter next to `a/b`. All dependencies related to an interface name or model data that has an interface name should be changed. This is because Excel does not convert this interface name to a date.

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

**Note**

In a VMware environment, each time you add a leaf, you must create a corresponding VMware vSphere Distributed Switch (vDS). See the [Notes Regarding VMware vSphere Distributed Switch, on page 14](#) section for details.

**Note**

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

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



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

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

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

Step 3 Click **Import CSV** to choose the CSV file. Browse for the CSV file, and click **Open**.

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

- **New**—Shows the new devices, fabric connection, and hosts included in the CSV.
- **Mis Matched**—Shows the mismatch between the uploaded CSV and the existing inventory. You can see the new and existing values for each of the entities, in this view.
- **Existing**—Shows all existing devices, fabric connections, and hosts in the inventory, and also present in the CSV.
- **Missing**—Shows the devices that are in the inventory, but not present in the CSV. Missing devices will be removed from current inventory when you update the inventory.

The following details are displayed for Devices:

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

The following details are displayed for Fabric Connection:

- **Target Device Name**
- **Device Type**
- **Target Device Interface**
- **Target Device IP**
- **Source Device Interface**

- Connection ID

The following details are displayed for Hosts:

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

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

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

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

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

- Network Inventory
- Host Inventory

For reuse, you can use the Export Inventory option to export the current inventory details in the CSV format. You can also export the CSV file :

- By choosing either **Inventory > Import and Discovery** or **Inventory > Network Inventory**.
- Without adding the Connection-Type and Server-id Export is allowed manually.

Note Import fails if Connection-Type and Server-id is Null/Empty.

Performing Auto Discovery

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

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

The auto discovery option has the following prerequisites:

- Link Layer Discovery Protocol (LLDP) has to be enabled on leafs, spine, DCI, and computes. See documentation for the respective devices for details about how to enable LLDP on these devices.
- Enable lldpd on computes. See [Enabling lldpd on Computes, on page 66](#) for details.



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

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



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

- All devices must have a common set of credentials. These credentials will be used during the discovery process. See [Managing Inventory, on page 61](#) for more information. The credentials must be of the appropriate privilege level on the devices.

To perform auto discovery:

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

Step 2 Enter the **Seed Device IP**.

Step 3 Enter the **Seed Device User Name**.

Step 4 Enter the **Seed Device Passphrase**.

Step 5 Click **Discover**.

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

Step 6 Click the desired cell for respective details to be populated in the screen. You may review the details, make changes wherever applicable, and click the **Add to Inventory** button to add the details into the Cisco VTS inventory. See [Working with Discovered Data, on page 68](#) for detailed information about the how to work with the discovered values.

Enabling lldpd on Computes

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



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

Enabling lldpd Using Ansible

To enable lldpd on computes:

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

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

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

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

A sample inventory file is given below:

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

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

Enabling lldpd Manually

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

Step 1 Uninstall lldpad on hosts.

```
yum -y remove lldpad

killall lldpad
```

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

Step 3 Install lldpad.

```
yum -y install lldpd
```

Step 4 Start lldpd daemon process.

```
lldpd
```

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

Step 6 Install the rpm.

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

Working with Discovered Data

Upon completion of discovery, the discovered details about the Devices, Fabric Connections, and Hosts are displayed as a matrix. It displays data in the following buckets. You can click each button, view the details that get displayed in the respective screens, and, wherever Cisco VTS allows edits, change the values. The tables below give detailed information about the discovered values in each bucket for Devices, Fabric Connections, and Hosts, and specifies whether edit option (including Bulk Edit option) is available. Make sure you also review the [Important Notes, on page 73](#) before you update the inventory.

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

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

Values	Discovered	Notes
Device Name	Yes	
Device IP	Yes	Update this with a new value, or retain the discovered data.
Auth Group	No	Select the desired value from the drop-down. Can be edited using Bulk Edit option too.
Device Platform	Yes	Update this with a new value from the drop-down, or retain the discovered data. Can be edited using Bulk Edit option too.
Device Role	No	Select the desired value from the drop-down. Can be edited using Bulk Edit option too.
Group Tag	No	Enter the Group tag value in the text box. Can be edited using Bulk Edit option too.
BGP ASN	No	Enter the ASN value in the text box. Can be edited using Bulk Edit option too.
Underlay Loopback Interface Num	No	Enter the loopback int num in the text box. Can be edited using Bulk Edit option too
Overlay Loopback Interface Num	No	Enter the loopback int num in the text box. Can be edited using Bulk Edit option too.

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



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

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

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



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

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

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



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

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

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

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

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

Value	Discovered	Notes
Device Role	No	Reconciled with existing in inventory. Can be edited after adding to inventory.
BGP ASN	No	Reconciled with existing in inventory. Can be edited after adding to inventory.
Underlay Loopback Interface Num	No	Reconciled with existing in inventory. Can be edited after adding to inventory.
Overlay Loopback Interface Num	No	Reconciled with existing in inventory. Can be edited after adding to inventory.

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

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

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

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

Value	Discovered	Notes
Virtual Switch	No	Reconciled with existing in inventory. Can be edited after adding to inventory.

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

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

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

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

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

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

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

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

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

Important Notes

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

- You must not add UCS 6200 Fabric Interconnects to the inventory even if these Fabric Interconnects are discovered during auto discovery.
- While adding new vCenter hosts into Cisco VTS, which has an existing inventory, you must:
 1. Export the current inventory.
 2. Update the exported inventory CSV file with the new vCenter Hosts.
 3. Reimport the CSV file into Cisco VTS, and update the inventory.
- If, in the CSV file you update existing devices authgroup and import again, in the GUI these will be shown under Mismatch devices. Clicking Update Inventory will update the authgroup of existing devices to the authgroup value you specified in the CSV. This change occurs even if you have a workload attached to the device. The same behavior occurs for BGP-ASN, and Loopback Interface Number also.
- After auto discovery is complete, for New devices, you must add the devices first, then add the fabric connection, and then the hosts.
- If you had changed the name of a TOR, which already exists in inventory, and then do a rediscovery, the TOR whose name is changes will be included in the New Device list, and a mismatch will be shown for Fabric Connection (Target Device Name). If you try to add the discovered fabric connection value to the inventory, it will throw an error. You must first add the new TOR to the inventory, and then add the newly discovered fabric connection.
- If Cisco Nexus 3000 device is used as a Leaf, then in Cisco VTS, the Device Platform needs to be set as Cisco Nexus 9000.
- For the New bucket, first add devices, fabric connections, and then hosts. For Missing bucket, first remove the hosts, then fabric connections, and then devices.
- The discovery process discovers only one connection for Cisco UCS B-Series hosts with multiple connections to the same interface. After discovery, you must manually add the details of the connections that are not discovered, via the Host Inventory page.
- When two ToRs are configured in vPC and no dual-homed host (connected to those ToRs) is in the VTS inventory, VTS does not correctly identify the vPC. You must add the dual-homed host connected to the ToRs in vPC to the VTS inventory, before provisioning a port on a host connected to the ToRs in vPC.
- Different ESI groups/domains must have different ES-id or system MAC. In other words, duplicate ES-id and system MAC are not allowed among ESI groups. This needs to be guaranteed by providing correct Day Zero configurations for ESI on Cisco Nexus 9000 switches.
- The Cisco VTS discovery log file is under `/var/vts/log`. Check for any errors/exceptions in this log file.

Viewing the Network Topology

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

To view the network topology:

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

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

Step 3 Select the tenant for which you need to view the topology, from the **Select Tenant** drop-down list.

The topology is displayed in the Topology window. You can use the following buttons to control the display:

- Select node mode
- Move mode
- Zoom in / Zoom out / Zoom Selection
- Fit Stage
- Full Screen mode

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

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

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

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

Viewing Network Inventory

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

To view the network topology:

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

The following details are displayed:

- Device Name

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

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

- Last Sync Operation

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

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

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

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

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

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

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

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

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

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

Adding Fabric Connection

To add fabric connection:

-
- Step 1** Go to **Inventory > Network Inventory**. The Inventory / Network Inventory window appears with the Network Inventory table displayed.
 - Step 2** Click Fabric Connection tab, then click **Add (+)** icon.
The Add Fabric Connection popup window appears.
 - Step 3** Enter the necessary details and click **Save**.
-

Synchronizing Configuration

You can check if the device configuration is in sync with Cisco VTS database, using the Check Sync option. Once Check Sync is complete, the sync status of the device along with the differences with the device is displayed. Options to Sync From, Sync To, and Reconcile Service are available. See [Important Notes, on page 76](#) section for important information related to Reconcile Service feature.



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

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

Step 2 Click the **Check Sync** link under the Sync column, for the device.

A popup window is displayed with Check Sync Results. The green + indicates additional configuration on the device and the red - indicates the additional configuration in the VTS database.

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

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

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

You can choose to initiate these actions on multiple devices. The requests are placed in a queue and each will be initiated in the order initiated.

If the action succeeds, a success green check icon is displayed in the selected device row. If it fails, a red failure icon is displayed in the selected device row. The tooltip for the critical icon displays which action failed and the reason for failure.

Important Notes

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

- We recommend that you use Out-of-Band reconciliation feature to reconcile configuration that is pushed to the device via ports created from VTS GUI only. Using this feature to reconcile configuration in a VMM integrated VTS setup, where ports are created from the VMM, might cause errors.
- You must ensure the day zero configuration on the device does not include configuration that will be pushed using Cisco VTS services or device templates. That is, device day zero configuration should not include configuration which would conflict with the configuration that VTS would be pushing into the device either via service configuration or device template configuration.

- In certain cases, if a port detach operation fails, you may need to remove any related out-of-band configurations from device, do an out-of-band reconcile operation from the Cisco VTS GUI, and then try the port detach operation again.

Viewing Host Inventory

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

To view host inventory details:

-
- Step 1** Go to **Inventory > Host Inventory**. The Inventory / Host Inventory page appears. The Host Inventory page has two tabs—**Virtual Servers** and **Baremetals**. By default, the page displays Virtual Server details.
- Step 2** To view host details on Virtual Servers, select the VMM from the Select VMM drop-down, and select the device from the Select Device drop-down list. The following details are displayed:
- Host Name
 - IP Address
 - Host Type
 - Associated VMM
 - Virtual Switch
 - Interfaces
 - Installation Status—Shows the installation status.
 - VTF Mode—Displayed on the top right of the table shows the VTF mode you have chosen in the Administration > Site Settings window.
- Step 3** To view host details on Baremetals, select the **Baremetals** tab, then select the device from the Select Device drop-down.
-

Adding a new Host on Virtual Servers

To add a new host:

-
- Step 1** Click the **Add (+)** icon. The Add New Host dialog box appears. It has two tabs—Host Details and Host Interfaces. the Host Details tab is selected by default.
- Step 2** Enter the following host details:
- Host Name—This is mandatory. Only letters numbers, underscore and dashes are allowed. Requires at least one letter or number. The hostname entered here needs to be in FQDN format, that is. <hostname>.<domain>.
 - Host IP Address—This is mandatory.
 - User Name

- Passphrase— User Name and passphrase are mandatory if you choose Non-OSPD VMM name in the VMM Name drop-down of the 'Host Configuration' section in the current popup window.
 - Host Configuration
 - VMM Name—The VMM to which you want to associate the host to. Depending on VMM chosen in the VMM Name section either the VTF Details information is pre-populated or you have to enter the details.
 - Virtual Switch—The following options exist:
 - not defined
 - ovs—If you want to install the VTS host agent on the compute, check the Install VTS agent on save check box.
 - vtf-l2—VTF is used as an L2 switch.
 - vtf-vtep
- Note**
- On selecting multiple virtual hosts for bulk edit, the Virtual Switch drop-down will display 'not-defined' by default. This will be case even if one of the selected host already has Virtual Switch which is different from 'not-defined'. User will have to select the required value in the drop-down and it will be applied to all the hosts selected in bulk edit.
 - The options displayed here depends on what you have specified in the VTF Mode field in Administration > Site Settings and the VMM type.

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

Step 3

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

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

Advanced Configurations Section:

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

If you want to install VTF on the compute select the checkbox 'Install VTF on Save'. Depending on the type of VMM Name chosen in the Host Details tab, either you can 'Save' or 'Save and Validate'.

The VMM can be OSPD/Non-OSPD VMM based on the VMM registration. See [Registering the Virtual Machine Manager using GUI, on page 26](#). For OSPD, the Host will allow for validation of installed plugins (either OVS or VTF). For Non-OSPD, the Host will allow installation of plugins on Host Inventory UI.

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

- Host Interfaces—This is mandatory.
- SR-IOV Enabled—Choose Yes or No from the drop-down to specify whether the interface is SR-IOV enabled.

Note For Host Interface(s) with SRIOV specified as true, it will not install/re-install VTS agent.

- Phys Net— Physnet name associated with the interface. If *SR-IOV Enabled* is Yes, it is the Physnet to be used for SR-IOV. If it is No, the other possibilities are that the port is associated to L2 switch or OVS. In case of OVS, you need to give Physnet intended to be used for OVS.
 - Attached to Device—Choose the device from the drop-down.
 - Device Port—This is mandatory. Choose the device port from the drop-down.
 - Group
- To add more interfaces, use the **Add (+)** icon.

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

Adding a new Host on Baremetal

To add a host:

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

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

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

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

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

- Device Port—This is mandatory. Choose the device port from the drop-down.
- Group

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

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

You cannot edit hosts on which there are workloads associated.

To delete a host from the table, select the Host Name check box corresponding to the device and click the **Delete (X)** icon.

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

Viewing the VTSR to VTF Mapping

Step 1 Go to **Inventory > Virtual Forwarding Groups**. The Inventory / Virtual Forwarding Groups window appears.

The window displays the number of VTFs that are attached to the VTSRs. The table on the right hand side shows the VTFs.

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

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

SR-IOV Support

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

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

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

The following Provider network types are supported:

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

The default tenant network type is VXLAN.

Assigning VLAN Ranges

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

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

See the [Managing Resources, on page 87](#) chapter for details about assigning VLAN ranges.

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

Trunk Port Support

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

Migrating from vPC to ESI

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



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

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

To migrate from vPC to ESI:

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

Step 2 Upgrade VTS to a version which supports ESI.

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

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

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

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

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

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

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

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

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

- Step 8** Redeploy inventory from Cisco VTS only for devices that have new ESI configuration. This is to make sure that Cisco VTS recognizes ESI configuration on Cisco Nexus 9000 series devices. See [Redeploying Device Inventory, on page 83](#) for details.
- Step 9** Remove the peer links between previous vPC peer TORs (**Inventory > Network Inventory > Fabric Connection**).
- Step 10** Add the ESI device group to appropriate functional groups in Admin Domain, and also disable ARP suppression at (**Overlay > Network**).
- Step 11** Upgrade VTSR to the latest image.
- Step 12** Run the Migration script from the path `/opt/cisco/package/vtc/bin/vpc-migration`. For an HA setup, run this on the Active VM.
- For example:

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

Where:

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

Redeploying Device Inventory

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

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

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

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



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

To redeploy device inventory:

-
- Step 1** Go to **Inventory > Network Inventory**, perform a sync-from for the device for which the configuration has changed. See [Synchronizing Configuration, on page 75](#) for more details.
- Step 2** Select the device, click **Redeploy**.
- Note** Redeploy just recalculates the inventory. Existing ports/VMs belonging to old device configurations, would not be updated or redeployed. You might need to delete and recreate the existing ports. We recommend that you use redeploy only if there are no existing ports/router/router interfaces.
- Note** If you delete devices from the inventory and also deleted VTSR with it, when you redeploy or reload the inventory, VTSR will not show up until it is reloaded or restarted. Power on the VTSR and wait for the registration with VTC to complete.
-

Enabling Static Multi Homing

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

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

Enabling Static Multi Homing on Cisco Nexus 7000

To enable static multi homing on Cisco Nexus 7000 devices:

-
- Step 1** Group the interfaces using the **Resources > Devices > Interface Groups** UI.
- Step 2** In Host Inventory, add the same tag for both the interfaces that are connected to the host for which you are enabling static multi homing.
- Step 3** If we want to group N7K devices, we need to use Device Interface Group only. If we want to switch or add couple of devices to interface group and those devices are already part of Admin Domain and have ports then we need to clean up the ports and remove/unchecked devices from Admin Domain, then add to interface group and then add to Admin Domain and create ports.
-

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.

- Step 3** If we want to group N9K devices, we need to use Device Group only. If we want to switch or add couple of devices to device group and those devices are already part of Admin Domain and have ports then we need to clean up the ports and remove/uncheck devices from Admin Domain, then add to device group and then add to Admin Domain and create ports.
-

Administrative State for Devices

You can use the **Admin State** option to change the state of the device(s). A device can have either of the following three admin states:

- **unlocked**: The device can be modified and changes are propagated to the real device.
 - **config-locked**: The device can not be modified. It is still possible to read the configuration of the device and send live-status commands or RPCs.
 - **southbound-locked**: The device can be modified, but changes are not propagated to the real device. Configurations can be created in this state even before the device is available in the network.
 - **locked**: The device can only be read.
-

- Step 1** Go to **Inventory > Network Inventory**. The Inventory / Network Inventory page appears.
- Step 2** Select the device(s) for which you want to change the admin state.
- Step 3** From the **Admin State** drop down, select an admin state (unlocked, locked, config-locked, or southbound-locked).
- Step 4** Click **OK**.
-



CHAPTER 7

Managing Resources

The resources in this chapter refers to VLAN, VNI or EVI, and Multicast-IP pools that need to be managed across all devices in the fabric. Some of the resources are fabric-wide, and apply across all devices in the fabric. Some of them are device scoped. The resources managed across all devices in the fabric are termed 'global'. The resources managed local to a device are termed 'device-local'. Cisco VTS facilitates the creation and management of these resource pools.

Overlay VxLAN networks can utilize the IP multicast capability of underlay fabric switches to handle BUM traffic. A range of multicast IP addresses need to be reserved across the fabric for this purpose. Hence multicast pool is a global resource. Likewise, the VNI allocation needs to be unique within the fabric and is also a global resource.

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 Provider VLAN Range, on page 88](#)
- [Specifying Global VNI Range, on page 89](#)
- [Specifying Global EVI Range, on page 90](#)
- [Specifying VLAN Range, on page 90](#)
- [Specifying Multicast IP Pool, on page 95](#)
- [Resource Pool Use Cases, on page 96](#)

Specifying Global Provider VLAN Range

OpenStack defines the notion of provider VLANs where the connectivity between the computes are backed by external switches carrying those specific VLANs. The expectation is that a given provider VLAN network, the segmentation ID is reserved within all switches in the fabric and dedicated for a specific tenant L2 network.

Cisco VTS can facilitate the provisioning of the provider VLAN in a fabric, by assigning those VLANs to an exclusive pool called the Global Provider VLAN pool. For normal tenant networks, VTS will always allocate VLAN IDs from a dynamic pool (which are device local pools). So, it is likely that for a given tenant L2 network, the VLAN-ID allocated on different switches can vary (They however translate to the same VNI). However, if there is a request from OpenStack admin to provision a provider VLAN network with an explicit segmentation id, and if the segmentation-id falls in the Global Provider VLAN pool range, Cisco VTS will honor the request and consider it to be a provider VLAN network. The VLAN ID specified by OpenStack is essentially 'reserved' across all devices in the fabric for that specific tenant network.

Global Provider VLAN pool has to be mutually exclusive from Cisco Nexus 9000 device VLAN pool and Cisco Nexus 7000 interface VLAN pool.

Global Provider VLAN pool was introduced in Cisco VTS 2.6.1. When upgrading a VTS deployment from prior releases (which did not support this feature), you may want to reserve a range of VLANs for global provider VLAN pool. However, gathering information about the usage of VLAN pools used across all devices can be cumbersome. A script called 'global_provider_vlan_tool.py' is available in Cisco VTS to find suitable 'gaps' in the VLAN range. And if needed, you can reserve specific ranges for Global Provider VLAN.

-
- Step 1** Go to **Resources > Global Provider VLAN Pool**. The Resource / Global Provider VLAN Pool window appears.
 - Step 2** Click the **Add (+)** icon. The Add Global VLAN Pool popup window appears.
 - Step 3** Specify the From and To values. This can be an integer number between 2 and 4094.
 - Step 4** Click **Save**.

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

Global Provider VLAN Tool

The global provider VLAN tool helps you find and free up VLAN range blocks within the resource pools, which can be used later for creation of global VLAN ranges.

The tool prompts the user to enter a comma separated list of range values to be freed up. Upon receiving the input, it re-carve existing device and interface ranges to accommodate the request. After having these ranges freed up, you can choose to create those ranges (or a subset of them) for the global VLAN pool.

The global_provider_vlan_tool.py script is located at /opt/vts/bin.

- Step 1** Run the global_provider_vlan_tool.py script. For example:

```
admin@vtc1:/home/admin# sudo python global_provider_vlan_tool.py
```

```
Available ranges for global vlan pool:
```

```
2-2021
```

```
2023-3009
```

```
3011-3929
```

```
3931-3979
```

```
3981-4095
```

```
Please enter global vlan ranges using common separated ranges (like
1001-1500,1600-1800,2500-3900) :2100-2200,2300-2400,2500-2600
```

```
Shrinking range 8181e1d5-1083-4c5f-90bf-ab2a77c03129 in pool n9k from 1500 - 2204 to 1500 - 2099
```

```
Creating range in pool n9k with values 2201 - 2204
```

```
Shrinking range 2dd0eadc-496e-47bd-ba2e-9498be5937c4 in pool vtf_11.11.11.11 from 2200 - 2214 to 2201
- 2214
```

```
Deleting range 29f37570-c4e1-4971-alab-6f7ad5ec84d9 in pool vts_n7k_ethernet1/5
```

```
Shrinking range f6979637-7699-406d-a59e-6a206872d657 in pool n9k from 2208 - 2999 to 2208 - 2299
```

```
Creating range in pool n9k with values 2401 - 2999
```

```
Shrinking range 07f7083a-eea0-4442-914a-a348703657f6 in pool n9k from 2401 - 2999 to 2401 - 2499
```

```
Creating range in pool n9k with values 2601 - 2999
```

```
Creating range in pool default with values 2100 - 2200
```

```
Creating range in pool default with values 2300 - 2400
```

```
Creating range in pool default with values 2500 - 2600
```

```
Done
```

Step 2 Exit the script.

```
root@VTC1:/opt/vts/bin# exit
exit
```

Specifying Global VNI Range

To specify the global VNI range:

Step 1 Go to **Resources > Global VNI Pool**. The Resources / Global VNI Pool window appears. 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 Global EVI Range

To specify the global EVI range:

- Step 1** Go to **Resources > Global EVI Pool**. The Resources / Global EVI Pool window appears. 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 EVI 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 overlay networks. The default BD range is 1000 to 2000. you can configure this value.

- For computes attached to Cisco Nexus 7000 switch as single homed, Cisco VTS automatically creates a default device interface pool per interface attached to the compute. This default interface pool is of name vts-Device-InterfaceName with a default VLAN range of 2-4094.
- If a compute attached to the Cisco Nexus 7000 is dual homed, Cisco VTS automatically creates a default device interface group pool containing the dual homed interfaces of two switches. This default interface group pool is of name vts-group-`<number>` with a default VLAN range of 0-4096.
- For single homed and non vPC interfaces, Cisco VTS creates a default per interface level VLAN range of 0-4096.
- If computes are attached to FEX, the interface pool with default range is created for the FEX device. This FEX VLAN pool is of name vts-device-`<fexId>` with default range of 0-4096.
- For computes attached in vPC to two different FEX, Cisco VTS automatically creates a group for the two interfaces going to two different FEXs. Overlay network provisioning on this vPC compute uses a common VLAN from the two FEXs ranges.

Cisco VTS only allows grouping of two FEXs to form a logical group. It does not allow a FEX from one logical group to form a grouping with a FEX from a different logical group. For example, if a host compute1 is connected in vPC to two FEXs 101 and 102, these two FEXs will form a logical group. Cisco VTS does not allow having a host compute2 connected in vPC to FEX 102 and FEX 103. This is because the same VLAN across multiple FEXs for a given network would be difficult to maintain.

- The default VLAN pool of a device gets deleted when the device is added to a Device Group. This is because it will be using the default VLAN pool of the Device Group once it is part of the group. However, when the device is removed from the Device Group, Cisco VTS does not add back the original VLAN pool to it. The Cisco VTS admin has to add the VLAN pool back to the device, manually.



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

See the following sections for details:

- [Specifying Device VLAN Range, on page 91](#)
- [Specifying Group VLAN Range, on page 92](#)
- [Specifying Interface VLAN Range, on page 92](#)
- [Creating Interface Groups, on page 93](#)

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 popup 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 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 popup 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 greyed 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.

- 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 greyed out. Editing is enabled for ports that are connected to computes.

- d) Click **Define Ranges** to define VLAN ranges. The Define Ranges and Group Details popup window appears.

You can use the Restricted range toggle button to restrict a range. You can add ranges using the Add (+) button.

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

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

Auto Select/Auto Delete functionality

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

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

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

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

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

When you deselect one of the peer ports/devices that is part of a system defined group (port channel in this case) or SMH group, then the corresponding peer port/device also gets auto de-selected/deleted and gets deleted from the interface group (both Physical and FEX interface group).

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

Specifying Multicast IP Pool

You can specify the number of overlay networks that can be mapped to a single multicast address. Choose Enter VNI (Network Count) 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.

Important Note

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

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

Resource Pool Use Cases

This section provides information about the VTS behavior related to resource pool allocation use cases.

Table 2: Terms and Description

Terms	Description
Restricted Range	A range of VLAN or VNI from which VTS is restricted from allocating. Typically this sort of restriction is done to accommodate the following use cases: - Deployments where the fabric controller (VTC) will honor the VLAN/VNI pool allocation done by the VMM for tenant networks. - Deployment where the fabric controller (VTC) will honor VLAN/VNI segment id allocation done by VMM for realizing Provider VLAN networks
Dynamic Range	A range of VLAN or VNI from which VTS would allocate for ports. Dynamic and Restricted ranges could overlap.
Out of Range	Any VLAN IDs that have not been allocated to either Dynamic Range or Restricted Range.
Provider VLAN/VNI	A VLAN/VNI id explicitly requested by the OpenStack Administrator to be assigned to a network. These are provider networks and are realized using hardware switches outside of compute.
Static VNI	A VNI ID explicitly is requested by the VTC administrator to be assigned to a network.
Static VLAN	A VLAN ID explicitly requested by the VTC administrator to be assigned to a port. Note For more information about Static VLAN options and the behavior of each options see, Static VLAN Options, on page 97

Table 3: Cisco VTS Behavior for Various Use Cases

Use Case	VTS Behavior		
	In Restricted Range ¹	In Dynamic Range	Out of Range
Provider VLAN segment ID used for SR-IOV links from Openstack	VTS will honor this if port VLAN is available	VTS will reject the request	VTS will reject the request
VLAN segment ID from Openstack	VTS will honor this if port VLAN is available	VTS will ignore VLAN from Openstack. Will allocate VLAN from dynamic range	VTS will ignore VLAN from Openstack. Will allocate VLAN from dynamic range
VxLAN segment ID from Openstack	VTS will honor this if VNI is available	VTS will ignore VNI from Openstack. Will allocate VNI from dynamic range	VTS will ignore VLAN from Openstack. Will allocate VNI from dynamic range
Static VLAN segment ID on port (VTC admin)	VTS will honor this if VLAN is available	VTS will honor this if VLAN is available	VTS will honor this if legacy-mode ² is OFF VTS will reject if legacy-mode is ON

Use Case	VTS Behavior		
	In Restricted Range ¹	In Dynamic Range	Out of Range
Static VxLAN segment ID (VTC admin)	VTS will honor this if VxLAN is available	VTS will honor this if VxLAN is available	VTS will honor this if legacy-mode flag is OFF VTS will reject if legacy-mode flag is ON
Multi-VMM VLAN provider network (merge)	VLANs needs to match on both VMMs. Else Reject	VTS will ignore VLAN from VMMs. Will allocate VLAN from dynamic range	VTS will ignore VLAN from VMMs. Will allocate VLAN from dynamic range
Multi-VMM VLAN provider network (publish)	VTS preserves the source VMM's VNI in published network. VLAN is not published.		
Multi-VMM VxLAN provider network (merge)	VNI needs to match on both VMMs. Else Reject	VTS will ignore VNI from VMMs. Will allocate VNI from dynamic range	VTS will ignore VNI from VMMs. Will allocate VNI from dynamic range
Multi-VMM VxLAN provider network (publish)	VTS preserves the source VMM's VNI in published network.		

- ¹ Restricted range is meant to be a global setting in VTS. But is currently realized as a device specific resource pool. Hence to achieve a global behavior, the same restricted range needs to be configured on all devices.
- ² Legacy mode—If enabled, VTS rejects if VLAN allocation is not in VTS managed pool.

Static VLAN Options

Following are the static VLAN options and the behavior of each option in legacy modes.

Request Type	VTS Version	Pool and Range Options	Static VLAN Behavior
Standard Static VLAN Request from VTS GUI or API	Before VTS 2.3.1	Only one range type is available for both dynamic and static allocations.	On BM ports attach operation, specify a VLAN. Note Ensure to add this VLAN as part of a predefined range else, the request fails.

Request Type	VTS Version	Pool and Range Options	Static VLAN Behavior
Standard Static VLAN Request from VTS GUI or API	VTS 2.3.1-> Present	<p>Define ranges as restricted. You can carve out restricted ranges from inside the dynamic ranges. Dynamic allocated (often called VTS allocated) will not touch restricted ranges.</p> <p>Note The term 'restricted' means ranges are restricted from dynamic allocation, not restricted from the user requesting it.</p>	<p>On BM ports attach operation, you can request a VLAN. This VLAN is available from a dynamic range, a restricted range, or no range at all.</p> <p>Note It is not recommended to create a range of 1 just to allocate a single VLAN outside of any range.</p>
Static VLAN Request with Range Check from VTS GUI or API	VTS 2.5.0 > Present	Same as above	Before 2.3.1, if you want to view the behavior of VTS failure with all requests that are not in a range, set a flag to re-enable this legacy option.

Request Type	VTS Version	Pool and Range Options	Static VLAN Behavior
Honoring Provider VLANs from Openstack	VTS 2.5.1 > 2.6.0	<p>Same as above.</p> <p>Note Create a restricted range matching the range of provider VLANs that are sent from Openstack in honor for VTS to honor it. A restricted range matching the provider range must be created on every device that provider networks create ports on.</p>	<p>If a port attach comes from Openstack and the network segmentation id VLAN is not in the restricted range, VTS allocates its own dynamic VLAN.</p> <p>If it is an SRIOV port attach operation, which must have a known provider VLAN, VTS fails the operation.</p> <p>For VTS GUI static VLAN requests, there is no change.</p>
Honoring Provider VLANs from Openstack	VTS 2.6.1 > Present	<p>Same standard options. You need not create a restricted range on every device. Instead, create a single range in the Global Provider VLAN Pool. Multiple ranges can be given if needed. The ranges though cannot overlap any of the ranges of the site's device pools.</p>	<p>Same as above.</p> <p>VTS does not check the restricted range for nonlegacy networks (networks that are created before 2.6.1). VTS will instead check the Global Provider VLAN Pool.</p> <p>For VTS GUI static VLANs requests, there is no change. Static VLANs are provided from the global provider pool ranges or the device pool ranges.</p>



CHAPTER 8

Creating and Managing Admin Domains

This chapter has the following sections:

- [Admin Domain Overview, on page 101](#)
- [Viewing Admin Domain, on page 102](#)
- [Creating an Admin Domain, on page 102](#)
- [Creating DCI Interconnect Profiles, on page 105](#)

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
- Route Reflector Functional Group

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 65](#) and [Managing Inventory, on page 61](#) sections for details.
- Reviewed the Supported Platforms section in the *Cisco VTS Installation Guide* which provide information about the platforms that Cisco VTS support, and their roles.

Step 1 Go to **Admin Domains > Domains**.

The Admin Domains / Admins window appears.

Step 2 Click **Create (+)**.

The Create New Admin Domain popup window appears.

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

Step 4 Click **Create**.

The Admin Domain canvas appears.

You can see the following functional groups on the left-hand side of the canvas:

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

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

Functional Group	Property
DCI	<p>Specify:</p> <ul style="list-style-type: none"> • Whether it is a New or Shared DCI. • The Redundancy / Availability settings: <ul style="list-style-type: none"> • Enable/Disable Redundancy using the toggle switch. • ICCP—VXLAN/fabric ICCP group number. Valid range is 1 to 4294967295. MPLS/core ICCP group number. Valid range is 1 to 4294967295 • ESI—Ethernet Segment ID for NVE overlay. Valid entry is a nine octet string. Each octet can contain one or two numbers in the range 0 to F. <p>Click Stitching Profile and choose the required profile.</p>
DC GW	<p>Specify:</p> <ul style="list-style-type: none"> • Whether it is a New or Shared DC GW. • The Control Protocol—BGP EVPN.
L3 GW	<p>Specify:</p> <ul style="list-style-type: none"> • Whether it is a New and Shared L3 GW. • The Control Protocol—BGP EVPN. • The Replication Mode—Multicast or Ingress. This is the data plane replication mode that will be used for VXLAN data plane traffic. The admin domain can contain devices that support common replication mode. <p>Note</p> <ul style="list-style-type: none"> • Cisco Nexus 5600 and Cisco Nexus 7000 supports Multicast replication mode only. • VTF supports Ingress mode only. • Cisco Nexus 9000 supports both modes. • Distribution Mode—Decentralized. <p>Note</p> <p>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.</p> <p>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.</p>

Functional Group	Property
L2 GW	Specify: <ul style="list-style-type: none"> • Whether it is a New and Shared L2 GW. • The Control Protocol—BGP EVPN. • The Replication Mode—Multicast or Ingress. • The Distribution Mode—Decentralized.

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.

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

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 popup box. Click **Yes** in the confirmation box to remove the link. See [Extending Layer 2 Network Across Data Centers, on page 162](#) for details.

While creating admin domain with large number of devices per L2GWgroup or L3GWGroup, user must add devices in smaller batches. Note that locally we have tried batch of 40 devices and the admin domain creation gets completed within 11 minutes.

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.

Note When we add DCI and DCGW without connectivity to L3GW & L2GW and save Admin Domain, upon deletion of DCI/DCGW from UI does not delete from NCS_CLI.

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 9

Managing Templates and Device Objects

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. 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 the border leaf. You must have an external network, and the External network should be set as the Router gateway. If you do not have the external gateway set, the template will be attached to the router, but configuration will not be pushed. After you have your external network as the router gateway, it will push the configuration.

- **L2 Service Extension templates**—An L2 Service Extension template allows you to extend Cisco VTS Layer 2 service configuration.

The creation of L2 Service Extension templates is done by authoring configuration that are specific to a device type. You can modify the L2 configuration that is attached to the network. The

configuration is for the L2 construct which can be applied to the virtual interface. Following L2 service configuration templates are supported:

- L2 QoS—Supported on Cisco Nexus 9000 series. See [Important Notes—L2 QoS Template, on page 119](#).
 - VPLS—Supported on Cisco ASR 9000 series. The L2 Extension templates for Cisco ASR 9000 series devices include the VPLS configuration under the L2VPN and EVPN containers. See [Important Notes—VPLS Template, on page 119](#).
- **Underlay template**—Underlay templates enables you to configure the Day Zero configuration on underlay devices via the VTS UI.



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

Cisco VTS requires you to preview the template configuration before you attach, detach, or edit L2 Service Extension, L3 Service Extension, and Device Templates. See [Previewing Template Configuration, on page 124](#) for details.

The Device Object feature enables you to create a parameterized device object of a configuration that is frequently used and apply the device object per device (For example, Port-channel). While creating Device Objects, instead of creating the values as is the case with Device Templates, you create parametrized variables, so that the same Device Object can be applied to multiple devices. While a device template is very specific to a device and tied to a device, a Device Object can be applied to different devices with different values.

Device objects feature leverages the device template concepts and exposes a set of targeted configs as device objects that can be easily attached to devices. For example, for Ethernet device object type, ethernet specific configuration is exposed in the device object. You do not need to go through the . entire model and search for various sub-trees in the device template creation flow.

The following sections provide more details about working with templates and device objects:

- [Creating Route Templates, on page 109](#)
- [Creating L3 Extension Templates, on page 111](#)
- [Editing Templates, on page 113](#)
- [Copying a Template, on page 114](#)

- [Deleting Templates, on page 114](#)
- [Importing and Exporting Templates, on page 115](#)
- [Attaching Templates to Routers, on page 116](#)
- [Creating L2 Extension Templates, on page 118](#)
- [Creating Underlay Templates, on page 122](#)
- [Previewing Template Configuration, on page 124](#)
- [Searching Template Content, on page 127](#)
- [Creating Device Objects, on page 129](#)
- [Editing Device Objects, on page 129](#)
- [Deleting Device Objects, on page 130](#)
- [Associating Device Objects to Devices, on page 130](#)
- [Editing Device Object Instances, on page 131](#)
- [Deleting Device Object Instances, on page 131](#)
- [Searching Device Object, on page 132](#)
- [Device Objects Notes and Caveats, on page 133](#)

Creating Route Templates

To create templates:

-
- Step 1** Choose **Templates > Overlay Template Management**. The Templates / Overlay Template Management page appears.
 - Step 2** Click **Add (+)**. 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 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 111](#) for details.
 - Step 9** Add Route Targets. See [Adding Route Targets, on page 109](#) for details.
 - Step 10** Click **Save**. The template is saved and listed in the Template Management page.
-

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 **Add (+)**. The Add Route Target(s) popup 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 External 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 **Add (+)**. The Add Route Target(s) popup 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# gets pushed to devices. In addition, the RT seed text box gets enabled again.

Creating L3 Extension Templates

To create L3 Extension templates:

- Step 1** Choose to **Templates > Overlay Template Management**. The Templates / Overlay Template Management page appears.
- Step 2** Click **Add (+)**. 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 Nexus 5000 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 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.

Note Currently, the *remote-as* (neighbor > IP > *remote-as*) attribute is not available in the L3 Service Extension templates for both Cisco Nexus 7000 series and Cisco Nexus 9000 series devices. This is now replaced by *inner-remote-as* (neighbor > IP > *inner-remote-as*) attribute.

The Author Template page provides two types of views:

- The Editable Tree view—This is the default view.
- The Read-only Config Preview—Lets you to view a complete summary of the configuration that will be pushed on the device. From the read only view, you can copy the configuration and paste it into any other editor, to view the configuration.

You can toggle between the views using the toggle button on the top right of the config pane.

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 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 112](#) for details.

Step 10 Click **Save Template**.

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

Supported System Variables

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

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

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

Supported variables in Cisco Nexus 9000 Series for L2 Service Extension:

- VLAN ID

Supported variables in Cisco ASR 9000 Series for L3 Service Extension:

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

Supported variables in Cisco ASR 9000 Series for L2 Service Extension:

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

Editing Templates

You can modify a template that you have created.

Step 1 Choose **Templates > Device Template Management** or **Templates > Overlay Template Management**. The Templates / Device Template Management or the Templates / Overlay 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 Click the **Edit Configuration** button to edit a template.

Step 3 Make the desired changes, then click **Review** to view the configuration summary for the device. It opens the Review modal and also processes dry-run/preview config call in the backend simultaneously and you can see spinning wheel on modal. Once done, it displays the 'Review' modal window. Click on the accordion icon/Template name/device name row will display the configs differences, if any.

Note The Review button is enabled on device/overlay template pages only if template has at least one device attached to it.

Step 4 Click **Save Template**.

Note If user wants to edit last route target of that combination of direction and type, then user needs to add the new route target of same combination. Go back to edit mode and delete the old pair of route template and save the template.

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 Choose **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 Choose **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 **Delete (X)**.

The Delete popup appears.

Step 3 Click **Delete (X)** to delete the template.

Importing and Exporting Templates

You can import and export Cisco VTS device templates and overlay templates.

If you have master templates created on a Cisco VTS lab instance, these can be exported to an external server. Modifications pertaining to the Cisco VTS production instance can be done on the server, and then the templates can be imported on the production instance. This helps you to avoid creating similar templates on multiple production instances.



Note Templates that are imported into Cisco VTS need to be in the valid JSON format. You might encounter errors if the format is incorrect.

The following sections give details about importing and exporting templates.

- [Importing Templates, on page 115](#)
- [Exporting Templates, on page 115](#)

Importing Templates

You can import a device templates or overlay templates into Cisco VTS. You can import a single file containing single/multiple device templates at a time. You cannot import templates with names that conflict with existing templates in Cisco VTS .

To import templates:

-
- Step 1** Go to **Templates > Overlay Template Management** or **Templates > Device Template Management**, based on your requirement.
 - Step 2** Click the Import icon.
 - Step 3** Locate the template JSON file and click open.
The template file gets imported into Cisco VTS, and is visible in the template list.
-

Exporting Templates

You can export a device templates or overlay templates from Cisco VTS to an external server. You can export a single or template in bulk.

To export templates:

-
- Step 1** Go to **Templates > Overlay Template Management** or **Templates > Device Template Management**, based on your requirement.
 - Step 2** Select the templates you want to export.
 - Step 3** Click the Export icon.

Step 4 Save the template file at the desired location.

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 116](#)
- [Attaching Template while Editing a Router , on page 117](#)



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 151](#) for details about attaching port.

Attaching Templates while Adding Routers

You can attach L3 Extension templates or Route template while you add a router.

- Step 1** Click **Overlay > Router**. The Overlay / Router page is displayed.
- Step 2** Click **Add (+)** icon. The Overlay / Network / Add Router page is displayed.
- Step 3** Click the **Templates** tab. The table displays the L3 extension templates and Route templates you have created for all supported device platforms. The drop-down list lets you display the desired template type for the desired device platform.
- Step 4** Select the Template type from **Showing** drop-down.
Template Type can be **L3 Extension** or **Route**.
- Step 5** Select the template(s) and click the **Attach** icon to attach the device(s).

Note You can attach only one Route template to devices at a time. If a Route template is already attached and you want to attach a new Route template, then you need to detach the existing Route template. You can attach multiple L3 extension templates to devices.

The Attach Devices to Templates(s) page appears with the following details:

- Template Name
- Device Platform Name
- Device Name
- IP Address
- Role
- Group Tag

- BGP-ASN
- Admin State

Step 6 Selected templates are listed in the Selected Templates pane. Select the desired template from the list, and then select the device(s) you want to attach.

You may choose to apply the selected templates in bulk to all of the devices, or choose specific devices on which you may want to apply the template.

You can remove the templates that you do not want to attach by clicking the **Remove** icon.

Step 7 For L3 Extension templates, you can review or skip the review and attach the configuration before you proceed. Select the device(s), then click **Review**. The review summary popup displays a hyperlink enabled Preview Config option. Click Preview Config link to view config changes in each device.

Note This feature is not supported for Route templates.

Step 8 Click **Skip Review: Attach** to attach the configuration directly to the device.

Step 9 Click **Save** on the Overlay / Network / Add Router page.

Devices added to the templates will be attached to the network.

Note You will see *Pending Save* in **Template Status** for the changes done, until you do a final save for the devices to be added to the templates.

Attaching Template while Editing a Router

You can attach L3 Extension templates or Route template while you edit a router.

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

Step 2 Select the Router you want to attach the template to, then click **Edit** icon. The Overlay / Network / Edit Router page is displayed.

Step 3 Click the **Templates** tab. The table displays the L3 extension templates and Route templates you have created for all supported device platforms. The drop-down list lets you display the desired template type for the desired device platform.

Step 4 Select the Template type from **Showing** drop-down.

Template Type can be **L3 Extension** or **Route**.

Step 5 Select the template(s) and click the **Attach** icon to attach the device(s).

Note You can attach only one Route template to devices at a time. If a Route template is already attached and you want to attach a new Route template, then you need to detach the existing Route template. You can attach multiple L3 extension templates to devices.

The Attach Devices to Templates(s) page appears with the following details:

- Template Name
- Device Platform Name
- Device Name

- IP Address
- Role
- Group Tag
- BGP-ASN
- Admin State

Step 6 Selected templates are listed in the Selected Templates pane. Select the desired template from the list, and then select the device(s) you want to attach.

You may choose to apply the selected templates in bulk to all of the devices, or choose specific devices on which you may want to apply the template.

You can remove the templates that you do not want to attach by clicking the **Remove** icon.

Step 7 For L3 Extension templates, you can review or skip the review and attach the configuration before you proceed. Select the device(s), then click **Review**. The review summary popup displays a hyperlink enabled Preview Config option. Click Preview Config link to view config changes in each device.

Note This feature is not supported for Route templates.

Step 8 Click **Skip Review: Attach** to attach the configuration directly to the device.

Step 9 Click **Save** on the Overlay / Network / Add Router page.

Devices added to the templates gets attached to the network.

Note You will see *Pending Save* in **Template Status** for the changes done, until you do a final save for the devices to be added to the templates.

Creating L2 Extension Templates

To create L2 Extension templates:

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

Step 2 Click **Add (+)**. The Create Template page appears.

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

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

Step 5 Choose L2 Extension as the template type.

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

- Cisco Nexus 9000 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 or configuration items which can take multiple instances, the Add Configurations pane appears in the authoring pane.

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

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

Step 10 Click **Add Instance** to add an instance. New Instance gets added on the left pane.

For configuration items which can take multiple instances, the Add Instance button appears in the authoring pane. Click **Add (+)** after you add the configuration for the instance.

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

The Author Template page provides two types of views:

- The Editable Tree view—This is the default view.
- The Read-only Config Preview—Lets you to view a complete summary of the configuration that will be pushed on the device. From the read only view, you can copy the configuration and paste it into any other editor, to view the configuration.

You can toggle between the views using the toggle button on the top right of the config pane.

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

Important Notes—L2 QoS Template

Review the section below before you create QoS template:

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

Important Notes—VPLS Template

Review the section below before you create a VPLS template:

- Ensure that the necessary Day Zero configuration is complete on the DCI. See *Cisco VTS Day Zero Configuration Examples* document for details.

- Create the Admin Domain. Make sure you link the L2GW and DCGW, which is required for L2VNI. See [Creating an Admin Domain, on page 102](#) for details.
- Create the DCI Interconnect profile with the specific details. See [Creating DCI Interconnect Profiles, on page 105](#) for details.
- Associate the DCI Interconnect profile to the DCI in the Admin Domain under DCI Interconnect profile tab.
- Create the L2 Extended Network with an EVI ID assigned to the Network under L2VPN tab. See [Creating a Network using Cisco VTS GUI, on page 151](#) for details.
- Make sure the subnets and port attach is done for the Network.
- You can create a template with just PW (pseudo wire) configuration, or just Access VFI configuration, or can have both configurations in one template. Make sure that you enter correct PW and Access VFI under L2VPN, in the DC that is not managed by VTS.
- The following are supported as system variables in VPLS templates:
 - Bridge Domain Name
 - Bridge Group Name
 - Interface NVE ID
 - EVI ID

Attaching Devices to L2 Extension Templates

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

-
- Step 1** Click **Overlay > Network**. The Overlay / Network page is displayed.
- Step 2** Select the network for which you want to attach the L2 extension template, then click **Edit** icon.
- Step 3** Click the **Templates** tab. The table displays a list of all L2 extension templates you have created for all supported device platforms. The drop-down list lets you choose the desired device platform.
- Step 4** Select the desired template(s), then click the **Attach** icon.
- The Attach Devices to Templates(s) page appears with the following details:
- Device Platform Name
 - IP Address
 - Role
 - Group Tag
 - BGP-ASN
 - Admin State
- Step 5** Selected templates are listed in the Selected Templates pane. Select the desired template from the list, and then select the device(s) you want to attach.

You may choose to apply the selected templates in bulk to all of the devices, or choose specific devices on which you may want to apply the template.

You can remove the templates that you do not want to attach by clicking the **Remove** icon.

Step 6 For L2 Extension templates, you can review or skip the review and attach the configuration before you proceed. Select the device(s), then click **Review**. The review summary popup displays a hyperlink enabled Preview Config option. Click Preview Config link to view config changes in each device.

Step 7 Click **Skip Review: Attach** to attach the configuration directly to the device.

Step 8 Click **Save Template**.

Devices added to the templates will be attached to the network.

Note You will see *Pending Save* in **Template Status** for the changes done, until you do a final save for the devices to be added to the templates.

Detaching Devices from L2 Extension Templates

You can detach a device from an L2 Extension template only to an existing network. To detach a device from an L2 Extension template:

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

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

Step 3 Click the **Templates** tab. The table displays a list of all L2 extension templates you have created for all supported device platforms. The drop-down list lets you choose the desired device platform.

Step 4 Select the desired template(s), then click the **Detach** icon.

The Detach Devices from Templates(s) page appears with the following details:

- Device Platform Name
- IP Address
- Role
- Group Tag
- BGP-ASN
- Admin State

Step 5 Selected templates are listed in the Selected Templates pane. Select the desired template from the list, and then select the device(s) you want to attach.

You may choose to apply the selected templates in bulk to all of the devices, or choose specific devices on which you may want to apply the template.

You can remove the templates that you do not want to detach by clicking the **Remove** icon.

Step 6 For L2 Extension templates, you can review or skip the review and attach the configuration before you proceed. Select the device(s), then click **Review**. The review summary popup displays a hyperlink enabled Preview Config option. Click Preview Config link to view config changes in each device.

Step 7 Click **Skip Review: Detach** to detach the configuration directly from the device.

Step 8 Click **Save Template**.

Devices added to the templates will now be detached.

Note You will see *Pending Save* in **Template Status** for the changes done, until you do a final save for the devices to be added to the templates.

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 **Add (+)**. 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.
- The Author Template page provides two types of views:
- The Editable Tree view—This is the default view.
 - The Read-only Config Preview—Lets you to view a complete summary of the configuration that will be pushed on the device. From the read only view, you can copy the configuration and paste it into any other editor, to view the configuration.
- You can toggle between the views using the toggle button on the top right of the config pane.
- 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 Underlay Template to Devices

To attach an Underlay template to a device:

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

Step 2 Click **Device Templates**.

Note You can filter templates by Device Platforms.

Step 3 Select the template(s) and click the **Attach** icon to attach the device(s).

The Attach Devices to Templates(s) page appears with the following details:

- Template Name
- Device Platform
- Device Name
- IP Address
- Role
- Group Tag
- BGP-ASN
- Admin State

Selected templates are listed in the Selected Templates pane. Select the desired template from the list, and then select the device(s) you want to attach.

You may choose to apply the selected templates in bulk to all of the devices, or choose specific devices on which you may want to apply the template.

You can remove the templates that you do not want to attach by clicking the **Remove** icon.

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

An attach icon is seen adjacent to the devices which have templates already attached.

Step 5 For L2 Extension templates, you can review or skip the review and attach the configuration before you proceed. Select the device(s), then click **Review**. The review summary popup displays a hyperlink enabled Preview Config option. Click Preview Config link to view config changes in each device.

Step 6 Click **Skip Review: Attach** to attach the configuration directly to the device.

The devices are successfully attached to the underlay template.

Detaching Underlay Template from Devices

To detach an Underlay template from a device:

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

Step 2 Click **Device Templates**.

Note You can filter templates by Device Platforms.

Step 3 Select the template(s) and click the **Detach** icon to attach the device(s).

The Detach Devices from Templates(s) page appears with the following details:

- Template Name
- Device Platform
- Device Name
- IP Address
- Role
- Group Tag
- BGP-ASN
- Admin State

Selected templates are listed in the Selected Templates pane. Select the desired template from the list, and then select the device(s) you want to detach.

Note You can detach a template only if the Template Status is *In use*.

You may choose to apply the selected templates in bulk to all of the devices, or choose specific devices on which you may want to apply the template.

You can remove the templates that you do not want to attach by clicking the **Remove** icon.

Step 4 Check **Devices(s)** to detach the devices.

Step 5 For L2 Extension templates, you can review or skip the review and attach the configuration before you proceed. Select the device(s), then click **Review**. The review summary popup displays a hyperlink enabled Preview Config option. Click Preview Config link to view config changes in each device.

Step 6 Click **Skip Review: Detach** to detach the configuration directly from the device.

The devices are successfully detached from the underlay template.

Previewing Template Configuration

While you attach, detach, or edit an L2 Service Extension or L3 Service Extension, Cisco VTS requires you to preview the template configuration before you proceed with the desired action.

The Review option displays a preview of the configuration summary for each of the devices. The Preview Config link displays the Preview Config Summary popup, which displays the configuration difference that will be pushed on each device. It also shows the timestamp at which the preview was run. The Change Summary gives a count of the number of configuration additions and deletions.

The configuration that is added to the device is shown in Green. The configuration that is removed is shown in Red. A + and - sign shows the additions and deletions respectively. You can click the > icon to get a preview of the configuration.

Preview Config Summary shows the configuration difference for all devices, including devices to which the template was previously attached. Devices to which the template has already been attached will be preselected and grayed out.

The Red out of sync icon indicates that the device is not in sync with the VTS database. You will not be able to attach, detach, or edit a template to the device, if the device is out of sync. You can perform the attach, detach, and edit operations only on those devices which are in sync with the VTS database (as indicated by the Green icon). However, you can still view the configuration difference with the current VTS database, using the drop-down arrow.

You either need to go back and deselect the out of sync device, or you need to go to Network Inventory and sync the device to VTS. Out of sync devices can be synced in Network Inventory

**Note**

You cannot deselect a device to which the template was previously attached, if the device is currently out of sync with VTS. This is because the UI does not allow you to deselect a device to which the template was already attached. You need to go to Network Inventory and sync the device with VTS before you proceed.

To review the template configuration while you edit a template, at least one device should be attached to the template before you edit the template. Otherwise, the option to review is grayed out.

When you detach a template from a device, the Preview Config Summary shows the configuration difference, which includes the configuration that will be removed from the device upon template detachment, and the additional configuration that Cisco VTS adds to the device.

See [Preview Template Configuration Examples, on page 126](#) to see examples of preview config summary for an attach, detach, and edit operation.

If the template application creates no change in the device configuration, the preview displays No difference.

For L2 Service Extension template, while you attach the template, the configuration summary shows the port configuration and the device configuration. If you have not created ports on a device during network creation, the summary will be shown as No difference for that device. This because the template configuration does not get pushed to the device unless you do a port attach.

For L3 Service Extension templates, the configuration summary shows the interface configuration and device configuration. If an interface is not attached to the Router, the configuration summary will be shown as No difference.

Under certain conditions, for instance, when the commands in the template are invalid, or if the feature is not supported by the device, the summary displays errors.


Preview Template Configuration Examples

The following images show examples of Read-only Config Preview and Preview Config during attach, detach, and edit operations, of an L2 Service Extension template, with ports on two devices, where one device is out of sync.

Figure 1: Read-only Config Preview

```
{
  config {
    vlan {
      vlan-list [%v:VLAN_ID] {
        id [%v:VLAN_ID]
        name DemoVlan
        shutdown false
        state active
      }
    }
    dot1Q {
      tag {
        native true
      }
    }
  }
}
```

Figure 2: Preview Config—Attach Template

 Attach Devices to Template(s)


Select atleast one device for each of the selected template.

Selected Template(s) (1)

DocSample_template1 (2)

Template Name: DocSample_template1
Device Platform: N9K

Showing Show All Devices

<input type="checkbox"/>	Name	IP Address	Role	Group Tag	BGP-ASN	Admin State
<input type="checkbox"/>	spine1	172.23.209.94	Spine-rr		1500	Unlocked
<input checked="" type="checkbox"/>	tor2 	172.23.209.91	Leaf		1500	Unlocked
<input checked="" type="checkbox"/>	tor3	172.23.209.92	Leaf		1500	Unlocked

Showing 1 to 3 of 3 entries Previous 1 Next

Cancel
Review
Skip Review: Attach

Figure 3: Preview Config—Edit Template

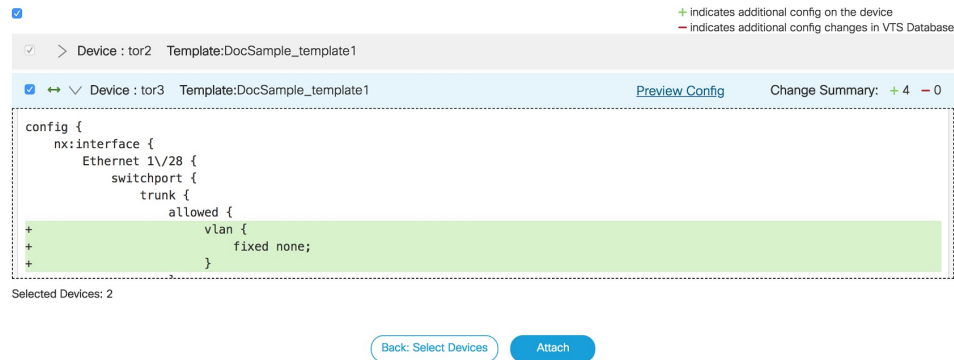
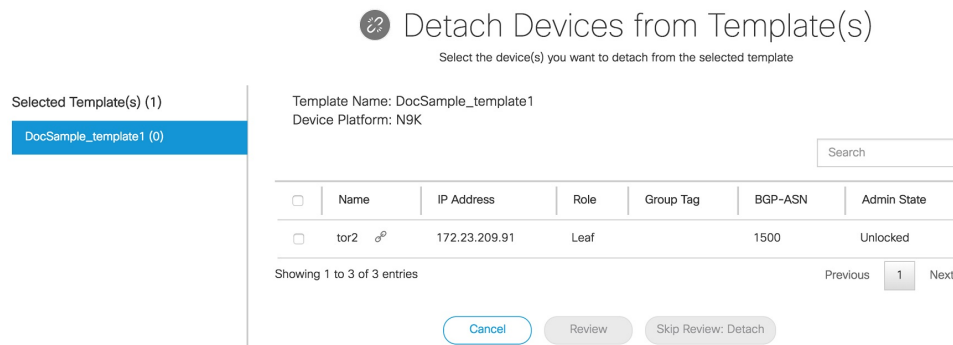


Figure 4: Preview Config—Detach Template



Searching Template Content

VTS Templates provide full text search in Overlay Template Management page and Device Template Management page.

Template search enables search on content and configuration across all the templates using solr engine. From VTS UI > Overlay Templates Management and/or Device Templates Management pages, you could enter a "term" in the search box and the Search Results page shows all the templates wherein the term is found either in template name or description or template config.

To perform a template search:

- Step 1** Choose **Templates > Device Template Management** or **Templates > Overlay Template Management**. The Templates / Device Template Management or Templates / Overlay Template Management page appears with a table showing a list of templates.
- Step 2** Enter the search text based on your requirement in the **Search text box** to search across the templates, and press **Enter**. The / Search Results page appears with the instances of the searched term highlighted.
- Step 3** You can refine your search results based on the filters on the left hand side of the search results page. The following filters are available:
 - Category—Filter based on template category, that is, Overlay Templates or Device Template.

Note If the Category is selected, then template type further refines based on categories (that is, it is dependent). If the Category is not selected, then template type refines based on the type of template.

- Template Type—Within a category, you can refine the search results based on the template type.
- Overlay—Filter based on the overlay association—Whether associated to Routers, Networks, or Devices.

Step 4 Click the template name to view the summary page which lists the number of occurrence of the search term in that template.

Step 5 Click the < or > button adjacent to search term in the summary page to navigate to the instance in the search result. The summary page lists the number of occurrence of the search term in that template config (that is, in config path and config value).

Step 6 Click **Export** or **Edit** if you want to export or edit the template.

Using Search Hints

You can search template content using search hints. You can use the following search hints.

Table 4: Using Search Hints

Search Hint	Description	Example
Text	Used for normal text-based search.	Type bgp to search for all strings that match with the term bgp
" "	Used to search complete phrases and IP addresses.	"service extension", "L3_TEMPLATE_N9K", "10.1.1.0/26", "2001:0db8:85a3:0000:0000:8a2e:0370:7334", "mtu:1500"
AND	Used to search when two terms exist. Note You must type AND in uppercase.	vrf AND bgp
OR	Used to search when either of the term exists. Note You must type OR in uppercase.	vrf OR bgp
NOT	Used to search when you need to exclude terms that are mentioned after NOT. Note You must type NOT in uppercase.	vrf NOT bgp or (vrf AND bgp) NOT ipv6
*	Used for wildcard search.	*—to display all search results swi*—to search for any word that starts with 'swi' virt*tch—to search for any word that starts with 'virt' and ends with 'tch'
?	Used for single character wildcard search.	te?t—to match test, text, and so on.

Creating Device Objects

To create a device object:



Note VTS 2.6.4 ships with four sample device objects that are read-only. Duplicate these device objects before use.

- Step 1** Choose **Device Objects > Manage**.
- Step 2** Click **+**. The Create Device Object popup is displayed.
- Step 3** Enter the **Name**. This is mandatory.
- Step 4** Enter a Description.
- Step 5** Choose the Device Platform. This is mandatory. Currently, Cisco VTS supports Cisco Nexus 7000 and Cisco Nexus 9000 series devices.
- Step 6** Choose the Device Object Type. The supported types are Ethernet and Port Channel.
- Step 7** Click **Add Configuration**. The Author Device Objects page is displayed.

After you pick a particular device object type, port channel for example, only those elements related to port channel get populated in the tree. You cannot delete these elements.

If you want to add more configurations under the device object, you can use the cheese icon to add these.

Under each configuration field, you can choose to enter static values or enter '\$value' to get it parameterized. If you choose a particular value to be parametrized you can enter the actual values while you associate the device objects to the device. See [Associating Device Objects to Devices, on page 130](#) for details.

All the mandatory config fields are given '\$values' by default and values for any other config fields like shutdown, mtu etc can be provided by the user.

Note Channel-group Id under Ethernet configuration is an exception. It is populated with a default '\$value' to match with the port-channel Name field.

- Step 8** Click **Save Device Object**. The Device Object gets listed in the Device Objects page. To view a Summary of the device object, click on the Name link. The Summary page shows Device Object Details and the Associations.

Editing Device Objects

To edit a device object:

- Step 1** Choose **Device Objects > Manage**.
- Step 2** Check the check box adjacent to the device object you need to edit. The Edit Device Object window appears.
- Step 3** Edit the Description, if required.
- Step 4** Click **Edit Configuration**. The Edit Device Objects page is displayed.

Make the necessary changes. You need to review the messages that might get displayed based on your changes.

Step 5 Click **Save Device Object**.

Note If device object is already associated with devices, then warning icon appears next to device object in Manage page listing. Also when the Device Object name is clicked, warning icon appears next to the associated devices. The warning icon indicates that the device object instance needs to be edited to update values based on device object edit.

Deleting Device Objects

To delete device objects:

Before you begin

Before deleting the device object, make sure no devices are attached/associated to the device object.

Step 1 Choose **Device Objects > Manage**.

Step 2 Check the check box adjacent to the device object you need to delete.

Step 3 The **X** icon is enabled.

Step 4 Click **X** icon to delete the selected device object.

Associating Device Objects to Devices

To associate device objects to devices:

Step 1 Choose **Device Objects > Associate**. The Associate page provides you options to create or delete device objects instances on devices.

Step 2 Select the Device Object you want to attach a device to and click **Add Instance**. The Add Instance page is displayed.

Step 3 Select a device.

After the device is selected, Device Object Instances are displayed. All fields are not mandatory to be filled. User can skip any '\$value' they added during the device-object creation. Any non-substituted values will be skipped on save. The device object config which is static (not parameterized) will appear read-only in this page.

If the same device object is to be attached to a different device, then the procedure has to be repeated.

Summary button is always displayed on Associate page irrespective of devices associated or not. The **Edit** button is used to exit the summary view and continue with the Add Instance flow.

The **Review** button is enabled only if instance name (base path param name – port-channel name and Ethernet interface name) is filled in for the device object.

Note In the case of Port-channel, you can create as many instances as required. In such cases of Device Objects, the Add Instance button is available. To create a new Port-channel Instance, you may click Add Instance button, and the configuration details/values of port-channel config are displayed.

The config which is of type list or collection (e.g. Ethernet) under port channel device object is shown with repetition capability wherein multiple values can be given.

Step 4 Click **Review**. A review dialog page displays all devices having instance(s) associated with the device objects with an optional Preview Config link on each listed device. The Configuration Summary displays all the configuration that will be pushed to the device. If you want to make any changes, click **Back to: Add** for any additional information.

Step 5 Click **Save** to push the configuration directly to the device.

Editing Device Object Instances

To edit a device object instance:

Step 1 Choose **Device Objects > Associate**. The Associate page provides you options to create or edit instance of device objects.

Step 2 Select the Device Object you want to attach to a device and click **Edit Instance**. The Edit Instance page is displayed.

Only one device and all instances tied to that device can be edited when you initiate this flow.

Step 3 Select a device which was entered earlier.

If you select Edit Instances in the Summary page and initiate this flow, the device is already selected and the instances tied to the device are displayed.

The fields display the values that were set in each of the fields. You can edit any of the fields or add additional instances or delete instances.

Step 4 Click **Review**. A review dialog page displays all devices having instance(s) associated with the device objects with an optional Preview Config link on each listed device. The Configuration Summary displays all the configuration that will be pushed to the device. If you want to make any changes, click **Back to: Add** for any additional information.

Step 5 Click **Save** to apply the changes on device.

Note Since the UI takes time to reflect the actual position of the popover, click on the (i) icon or eye icon to open the popover information.

Deleting Device Object Instances

To delete device object instances:

Step 1 Go to **Device Objects > Associate**. The Associate page provides you options to create or edit instance of device objects.

Step 2 Select the Device Object from where you want to delete the instances.

Step 3 Click **Edit Instance**. The Edit Instance page is displayed.

The fields display the instances associated to the device.

Step 4 Click **Delete Instance**.

All the instances associated to the selected device can be deleted but only one instance can be deleted at a time when you initiate this flow.

Searching Device Object

VTS 2.6.6 provides full-text search capability for Device Objects. Full-text search enables search on content and configuration across all the Device Objects using solr engine. The Search function is specific to the Device Objects in each site.

In the **Device Objects** page, you can enter a term in the **Search** field. The search result displays all the Device Objects and Device Object Instances that contain the search term in **Name** or **Description** or **Device Object config**.

Follow these steps to perform a Device Object search:

Step 1 Choose **Device Objects > Manage or Association** tab. A list of available Device Objects is displayed.

Step 2 Enter the search text, based on your requirement in the **Search all Device Objects** field, to search across the device objects, and then press **Enter**.

The search result displayed splits each Device Object into the two different type of items, namely, Device Object Definition and Device Object Instance. If the search term matches with the configs of the Device Object, the corresponding keypaths are listed after each item.

Note If the device object is associated with multiple devices, each device object instance per device is shown in the search results. Device Object Instance name has the following format:

<DeviceObjectName_InstanceType(and)InstanceName_DeviceName>

It displays the device name, device object name along with other metadata of the device.

Step 3 You can refine your search results using the filters on the left pane of the search results page. The following filters are available:

- **Categories:** Filters the Device object Instances based on device object definition or device object Instance.

Note You can refine a category based on the selection.

- **Device Object Type:** Filters the Device object based on device object type.
- **Device Object Instance Type:** Filters the Device object Instances based on device object instance type.

Step 4 Click the **Device Object Name** or **Device Object Instance** name to view the **Summary** page. The Summary page has the following two tabs:

- Device Object Details
- Association Information

The **Device Object Details** tab displays the number of occurrence of the search term in that entity.

Step 5 Click the < or > button adjacent to search term on the **Device Object Details** tab to navigate to the matching configuration in the search result.

The **Summary** page lists the number of occurrence of the search term in that device object config, namely, in the **config path** and **config value**.

Step 6 Click **Export** or **Edit** if you want to export or edit the Device Object definition.

Device Objects Notes and Caveats

Following table lists some of the important notes and Caveats for Device Objects.

Table 5: Important Notes and Caveats

Item	Description
Configuration	Using “switch port trunk allowed vlan none” config as part of device template or device object wipes the vlan id configs at later points (redeploy, reconcile and upgrade of services) on device with port attach on same interface. Hence, configure and use it with expected impact in all functional flows. For example, Port-attach and then attaching device template will wipe off the vlan id of port attach config.
Sample Device Objects	Four sample device objects (Port-channel DO, Ethernet DO each for N9k and N7k platform types respectively) for read-only from UI are available. If you edit these Device objects from the Manage/Authoring and Associate page, following valid error will be displayed on the UI. Sample Device Objects cannot be updated. Copy Sample Device Object for further usage
Configuration	N7K configurations under "isis" container config should be configured as day0 on device. NED needs to provide the sub-level configs support under "isis" container config. RT#34389.

Item	Description
Sub-config under a config with composite-key	<p>If a normal config (non-composite keys config) contains two \$PARAMS of list type in keypath, combination of their (m * n) values given to lists is used to create that many substituted config keypaths to be pushed to device.</p> <p>In case of config with composite-key that is having a list subconfig, combination cannot be done for the same and should have 1-1 mapping.</p> <p>For example, Here is the config for “VR and Address family” and addresses within that composite key.</p> <pre>config/nx:interface/Ethernet{\$ETHERNET-NAME}/vrrpv3 {\$ETHERNET_VRRPV3-VR " \$ETHERNET_VRRPV3-ADDRESS-FAMILY} "</pre> <p>It has sub config of primary addresses (which are of list type)</p> <pre>config/nx:interface/Ethernet{\$ETHERNET-NAME}/vrrpv3 {\$ETHERNET_VRRPV3-VR \$ETHERNET_VRRPV3-ADDRESS-FAMILY}/address/ primary-list{\$ADDRESS_PRIMARY-LIST-ADDRESS}</pre> <p>If user provides values for \$ETHERNET_VRRPV3-VR as 100, 200 and for \$ETHERNET_VRRPV3-ADDRESS-FAMILY as ipv4, ipv6 and for ADDRESS_PRIMARY-LIST-ADDRESS as 10.10.10.10, 2001:db8:abcd:0012::0/64</p> <p>In this case, you cannot do a combination of primary addresses to the composite keys</p> <p>That is, composite key {100, ipv4} having both the addresses 10.10.10.10 and 2001:db8:abcd:0012::0/64 -- This will not work as address family is ipv4. Hence 1-1 mapping of the list values under composite-key needs to be done so it should look like.</p> <pre>{100,ip4} having just 10.10.10.10 {200,ipv6} having just 2001:db8:abcd:0012::0/64</pre> <p>Currently this creates a limitation that the user cannot supply multiple primary addresses to address family.</p>

Item	Description
<p>Configuration</p>	<p>Device specific behavior: Wherein cleanup is not working properly for vlan ids for Ethernet interface configs with channel-group id referenced in it.</p> <p>The “switchport trunk allowed vlan <id>” will be pushed to both port-channel and Ethernet interface referencing it. But, while detaching the config “switchport trunk allowed vlan <id>” was not removed from Ethernet interface. You need to enable the “switchport” config in device object definition under port-channel container for the cleanup to work properly on device.</p> <p>Note Device does not reject the config push “switchport trunk allowed vlan <id>” without enable “switchport” config.</p>
<p>Port-channel Ethernet supported config</p>	<p>Under Port-channel device object, the only ethernet config allowed and supported from VTS UI and API is:</p> <pre> keypath-value "config/nx:interface/Ethernet{\$ETHERNET-NAME}" { value {"Ethernet\": {"name\": \"\$ETHERNET-NAME\", \"shutdown\": \"\$SHUT1\"}}"; } keypath-value "config/nx:interface/Ethernet{\$ETHERNET-NAME}/channel-group" { value {"channel-group\": {"id\": \"\$PORT-CHANNEL-NAME\"}}"; } keypath-value "config/nx:interface/Ethernet{\$ETHERNET-NAME}/enable" { value {"enable\": {"switchport\": \"true\"}}"; } </pre>

Item	Description
CSCvm29427	<p>The Edit instance page shows address with comma separated values for each of the addresses in instance.</p> <p>Build#13 262 throttle build</p> <p>Issue: Edit instance page shows address with comma separated values for each of the addresses in instance.</p> <p>Steps:</p> <ol style="list-style-type: none"> 1. Create Ethernet DO type with vrrpv3 and address config under it. 2. Create 2 instances for above DO on Add instance page and add 2 sets of vrrpv3 configs and attach it to device. 3. On Edit instance page, we can see address with comma separated values for each of the addresses in instance. <p>Workaround: Under vrrpv3 Choose address repeater list and then update the address with individual values again in Edit instance page.</p>
CSCvm42414 Configuration (Preview config behavior of device object in a particular scenario)	<p>Deletion of DO instance does not work if config is updated after config push.</p> <p>For port-channel device object with ethernet interface associated to it, if “enable -> switchport” config is not set to “true” then vlan id will be pushed only to “port-channel”. Once we update the enable -> switchport” config to “true” then the vlan id is pushed to associated ethernet interface accordingly.</p> <p>But, when deleting the above port-channel instance device throws invalid command error. Hence, this is expected N9K platform behavior with cleanup of configs as per the sequence of configs pushed.</p> <p>Work around: Manually disassociate the channel-id from Ethernet interface and then try updating “enable -> switchport” config to default param value, delete the port-channel from UI/API.</p>



CHAPTER 10

Managing Tenants

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

This chapter has the following sections:

- [Viewing Tenant Details, on page 137](#)
- [Adding Tenants, on page 137](#)
- [Editing Tenants, on page 138](#)

Viewing Tenant Details

Go to **Tenants > Tenant Management**. The Tenant Management page lists all the available tenants.

By default, the tenants under VTS are displayed. You can choose individual VMMs from the drop-down to display the tenants under these.

The page displays the following:

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

Adding Tenants

To add tenants:

Step 1 Go to **Tenants > Tenant Management**. The Tenant Management page appears.

- Step 2** Click **Add (+)** icon.
- Step 3** Enter the following:
- Tenant Name
 - Description
- Step 4** Click **Save**.
- Step 5** To add a Zone, click the **Add (+)** icon. The Add Zone popup appears.
- Step 6** Enter the zone name, and click **OK**.
- Step 7** 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** Click **Save**.
- Step 5** To enable or disable network extension, use the **Extend all networks** toggle switch. By default, Extend all networks is **Yes**.
- Step 6** Modify the zone details.
- Step 7** Click **Save**.
-



CHAPTER 11

Deploying Security Groups

In a cloud-enabled Data Center, security enforcement is no longer just network-centric (such as network addresses and VLAN attributes). Security has to be enforced specific to application requirements, who the tenant is, and which tier of the application is being protected.

Isolation is the basis for any network security strategy. Isolation has been accomplished in traditional environments through the manual configuration of ACLs or firewall rules on physical devices. In case of VTS enabled overlay networking, tenant isolation and network isolation are enforced by default. Overlay network isolation is achieved by the associated encapsulation mechanism on the VXLAN data plane. If an attack is started by an application workload inside a virtual network, the physical infrastructure of a cloud is completely protected by this isolation.

Segmentation adds security controls to smaller groups of workloads. In an overlay/virtual network, the ACL services are required to be provisioned near the application workloads. ACLs enable to place restrictions on a selective basis to restrict the communication between VMs. The ACLs can be realized on the hardware and software VTEPs.

Security group is a named collection of network access rules that are used to limit the types of traffic that have access to instances. Security rules define access rules within a security group. Security groups consists of security rules on the underlying hardware. They minimize the risk of data leak and protect the datacenter deployments through a proactive stance.

Security Policies are instantiated on VTS Policy plane based on OpenStack Security Groups APIs.

OpenStack Security Group is a named collection of network access rules that are used to limit the types of traffic that have access to instances. When launching an instance, administrator can assign one or more security groups to it. If not assigned, new instances are automatically assigned to the default security group. See OpenStack documentation for more details about OpenStack Security Groups.

The associated rules in each security group control the traffic to instances in the group. Any incoming traffic that is not matched by a rule is denied access by default. Rules can be added to or removed or modified for the default and any other security group. Rules are automatically enforced as soon as it is created or modified.

Admin can modify the rules in a security group to allow access to instances through different ports and protocols. For example, admin can modify rules to allow access to instances through SSH, to ping instances, or to allow UDP traffic; for example, for a DNS server running on an instance by specifying the following parameters for rules:

- **Source of traffic**—Enable traffic to instances from either IP addresses inside the cloud from other group members or from all IP addresses.
- **Protocol**—Choose TCP for SSH, ICMP for pings, or UDP.

- **Destination port on virtual machine**—Define a port range. To open a single port only, enter the same value twice. ICMP does not support ports; instead, you enter values to define the codes and types of ICMP traffic to be allowed.

When the OpenStack security is passed through the VTS ML2 plugin, VTS programs these policies as the ACLs in the underlying forwarding elements.

OpenStack security groups are realized using:

- ACLs on VTFs.
- OVS and Linux IP tables on compute nodes.
- ACLs on TORs for Bare metal and Virtual workloads.

This chapter has the following sections:

- [Security Group - Feature Scope, on page 140](#)
- [Support for Reflexive ACLs, on page 142](#)
- [Creating Security Groups from Cisco VTS GUI, on page 142](#)
- [Attaching Security Group to Baremetal Port, on page 143](#)
- [Detaching Security Group from Baremetal Port, on page 144](#)
- [Attaching Security Groups to OVS, VTF, and SR-IOV Ports, on page 144](#)
- [Detaching Security Groups from OVS, VTF, and SR-IOV Ports, on page 144](#)
- [Security Group - Examples, on page 144](#)

Security Group - Feature Scope

Following are the Port types supported in Security Group (SG):

Table 6: Port Types Supported

Port Types	Details
VTF Ports	<ul style="list-style-type: none"> • No support for remote security group. • All other OpenStack Security Group functionality can be fully realized on VTF Ports.
OVS Ports	Fully Supported
Baremetal Ports and SRIOV Ports	<ul style="list-style-type: none"> • No support for remote security group. • Reflexive ACLs are not supported. • Security Group Rules applied to traffic ingressing SRIOV port may not get enforced when the traffic is L2 traffic coming from a ToR different from destination ToR. • ACLs on Cisco Nexus 9000 series device cannot block Intra-compute SRIOV traffic. This is device platform issue. • No support for Cisco Nexus 7000 series device.

Table 7: Feature Supported - Detailed Table

Security Group Features	OVS	VPP	SR-IOV on Cisco Nexus 9000	BM on Cisco Nexus 9000
Default SG without Remote SG	Yes	Yes	Yes	Yes
Default SG with Remote SG	Yes	The default SG will be ignored.	The default SG will be ignored.	NA
Custom SG without Remote SG	Yes	Yes	Yes	Yes
Custom SG with Remote SG	Yes	The remote-sg rule will be ignored.	The remote-sg rule will be ignored.	NA
Reflexive Policies	Yes	Yes	No	No
Implicit DHCP allow	Yes	Yes	Yes	Yes
Routed Traffic	Egress/Ingress	Egress/Ingress	Egress/Ingress	Egress/Ingress
Bridged Traffic	Egress/Ingress	Egress/Ingress	<ul style="list-style-type: none"> Egress Only for Inter-Compute. None for Intra Compute (Traffic does not come in TOR). 	Egress Only for Inter-Compute

**Note**

- OpenStack, by default, associates all VMs with their respective Tenant (or Project) 'default' sg. As OpenStack does not support SG for SRIOV Ports, 'default' sg associated with SRIOV ports gets ignored and all traffic will be allowed to passthrough. Same is the case with VTF Ports, as in prior releases VTS did not support SG for VTF ports. From Cisco VTS 2.6.0, the intent of these SGs—'default' or not, will start getting fully realized by Cisco VTS for SRIOV and VTF ports, provided these rules do not contain remote-sg rules. 'remote-sg' rules are not support for non OVS Ports—VTF, SRIOV and Baremetal. If a given SG happens to have a remote-sg rule then please refer to this section for details about expected behavior.
- See Cisco VTS syslog for error details.
- For Reflexive policies, reverse ACLs/Security rules need to be configured explicitly. There will not be any error logs.
- Cisco VTS does not allow you to create rules with remote SG.



Important Review the Security Groups feature specific information in the *Limitations and Restrictions* section of the *Cisco VTS Release Notes* before you create or attach security groups.

Support for Reflexive ACLs

This feature allows the ACLs configured on VTF Ports to be of reflexive nature. Reflexive ACL takes a packet flow, gets session information, and creates dynamic ACL entry in access-list in reverse direction. This entry gets automatically removed either after the session completes or times out. This dynamic insertion of rules rids the user of the need to explicitly program rules to allow reverse direction traffic.

Prior to this feature support, ingress rules corresponding to each egress rule to allow reverse traffic (and vice versa) had to be explicitly added.

With Reflexive ACLs feature, VTF behavior for Security Groups configured through OpenStack is brought to parity with OVS. For OVS, reflexive is always turned on. If you desire to turn this feature off for VTF ports then set the flag `vtf-sg-reflexive-acl-enabled` in `global-settings` to false. This setting applies only to the VTF ports and not OVS.

Creating Security Groups from Cisco VTS GUI

To create security groups in Cisco VTS:



Note These security groups can be attached only to Baremetal Ports from Cisco VTS.

-
- Step 1** Go to **Tenants > Security Groups**. The Tenants / Security Groups window appears.
- Step 2** Select **VTS** and **Tenant** as the source from the drop-down list.
- Note** If you have created a Security Group from OpenStack, it will show under a different source (OpenStack) and it will not show under VTS. You cannot add or edit or delete a Security Group from Cisco VTS after creating it under OpenStack.
- Step 3** Click **Add (+)** icon to create a new Security Group. The Tenants / Security Groups / Create New Security Group window appears.
- Step 4** Enter the Security Group name. The name requires at least one alphabet or number. Characters "and" are not allowed. The character limit is 255.
- Step 5** Select the Tenant from the drop-down list, if you want to change the tenant.
- Step 6** In the Description field, enter a description for the Security Group. The character limit is 255.
- Step 7** Click **Create**. The Tenants / Security Groups / <new security group name> window appears in which you can see the Security Group details with two default rules that gets added to the new Security Group created. You may remove the default rules if you wish to. To do this, check the check boxes and click **Delete (x)** icon.
- Step 8** Click **Add (+)** icon to create a new rule for the Security Group.

Note Rules you create here cannot be edited. Rules can only be added or removed.

Specify the following Parameters:

- Direction
- IP Protocol
- Port/ Port Range
- IP Protocol Number
- Remote CIDR

Step 9 Click **Save**.

The rules created get saved to the VTS database.

Note You may click on the Security Group name link in the table to review the details.

Attaching Security Group to Baremetal Port

To attach a Security Group to a Baremetal Port:

Before you begin

Create a network before you do a port attach. See [Creating a Network using Cisco VTS GUI, on page 151](#) for details.

Step 1 After you enter the details for attaching a port, click the **Next: Attach Security Groups** button to attach the Security Groups. The Attach Security Group window appears.

Step 2 Specify the Baremetal IPv4 / IPv6 address. You may use CIDR notation.

Step 3 Check the check box corresponding to the Security Groups you want to attach from the Available Security Group(s) table.

Note You can attach different Security Groups from different source by selecting it from **Source** pull down list.

The selected Security Group gets added to the Attached Security Group(s) table. Click the **Expand (>)** icon if you want to expand the Security Groups to view its rules.

Step 4 Click **Review** icon. The Review window appears for you to review the details.

Step 5 Click **Done** icon. The Overlay / Network / Fabric Host Networks / Edit Tenant Network window appears.

Click the link **View / Edit** icon in the Baremetal Ports table if you want to view or edit the Security Group attached.

Step 6 Click **Review** to review the details of the Security Group you have attached.

Step 7 Click **Done** .

Step 8 Click **Save**.

Detaching Security Group from Baremetal Port

To detach a Security Group from a Baremetal Port:

-
- Step 1** Go to **Overlay > Network**. The Overlay / Network window appears.
- Step 2** Click **Fabric Host Networks** tab and select the Network.
- Step 3** Click **Edit** button.
- Step 4** Click on the **Ports** tab.
- You can see both Baremetals Ports and Virtual Machines Ports on the left hand side panel.
- Step 5** Click the **View / Edit** link in the Baremetal Ports table. The Review window appears.
- Step 6** Click **Edit** on the Attached Security Group(s) pane. The Attach Security Group window appears.
- Step 7** Uncheck the check box corresponding to the Security Group you want to detach from the port. The Security Group moves from the Attached Security Group(s) pane to the Available Security Groups pane.
- Step 8** Click **Review**. The Review window appears. Make sure that the Security Group you wanted to detach is not listed in the Attached Security Group(s) pane.
- Step 9** Click **Done**.
- Here you can see the number of the Security Groups that are currently attached, after you have detached the Security Group(s).
- Step 10** Click **Save** to save the changes.
-

Attaching Security Groups to OVS, VTF, and SR-IOV Ports

To attach Security Groups from OpenStack to OVS, VTF, and SR-IOV, see the OpenStack Horizon documentation.



Note Ensure that you do not select any remote-sg rules while you attach security groups to VTF and SR-IOV ports.

Detaching Security Groups from OVS, VTF, and SR-IOV Ports

Detach Security Groups from OVS, VTF, and SR-IOV from OpenStack. For more information, see the OpenStack Horizon documentation.

Security Group - Examples

This section provides examples of Security Group use cases.

Creating Security Group to Restrict Access to a Given Application

Associating SRIOV port with Security Group



CHAPTER 12

Provisioning Overlay Networks

This chapter has the following sections:



Note

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

- [Provisioning Overlay Networks Using Cisco VTS, on page 147](#)
- [Creating Overlays, on page 148](#)
- [Creating Network using VMware, on page 150](#)
- [Creating Subnetwork using VMware, on page 150](#)
- [Creating Routers using VMware, on page 150](#)
- [Attaching Network to Router, on page 150](#)
- [Attaching a Virtual Machine to Network, on page 151](#)
- [Creating a Network using Cisco VTS GUI, on page 151](#)
- [Creating Router using Cisco VTS GUI, on page 152](#)
- [Port Extensions Support , on page 155](#)
- [Assigning BVI Interface IP Address, on page 161](#)
- [Extending Layer 2 Network Across Data Centers, on page 162](#)
- [Enabling Global Route Leaking Service, on page 162](#)
- [Enabling L3VPN to EVPN Route Stitching, on page 164](#)
- [Adding Static Routes, on page 164](#)
- [OpenStack Allowed Address Pairs Support, on page 167](#)

Provisioning Overlay Networks Using Cisco VTS

Cisco VTS 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. Cisco VTS supports hardware and software VTEPs to segment the data center network.

Cisco VTS 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 Cisco VTS environment.

Cisco VTS implements the highly scalable MP-BGP with the standards-based EVPN address family as the overlay control plane to:

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

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

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

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

- Prepare the physical environment to be managed by Cisco VTS to build virtual overlays. See the *Prerequisites* section in the *Cisco VTS Installation Guide* for details.
- Discover the network topology in the data center. See the [Managing Inventory, on page 61](#) chapter for details.
- Define Admin Domains. See the [Creating and Managing Admin Domains, on page 101](#) chapter for details.

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

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

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

Creating Overlays

As part of overlay provisioning, you may need to:

- Create Tenant
- Create Network
- Create Subnet
- Create Router

- Create VM

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

**Note**

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

Using OpenStack

**Note**

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

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

Using VMware

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

- [Attaching Network to Router, on page 150](#)
- [Creating Network using VMware, on page 150](#)
- [Creating Subnetwork using VMware, on page 150](#)
- [Creating Routers using VMware, on page 150](#)
- [Attaching a Virtual Machine to Network, on page 151](#)

**Note**

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

Using Cisco VTS GUI

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

- [Creating a Network using Cisco VTS GUI, on page 151](#)
- [Creating Router using Cisco VTS GUI, on page 152](#)

Creating Network using VMware

To create a network:

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

Creating Subnetwork using VMware

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

To create subnetworks:

-
- Step 1** Select one of the vDS switches you had created, then click the **Manage** tab.
 - Step 2** Select the **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 the **Create Subnet** button to create subnetwork.
 - Step 5** Click the **Refresh** button to see the subnetwork.
-

Creating Routers using VMware

To create routers:

-
- 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 the **Tenant Name** and enter the **Router Name**.
 - Step 4** Click the **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 the **VTS Router** tab and click the **Router Name** where network has to be added.
The Router Details dialog box 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 VTS plugin. This will create portgroup for the network.
- Step 2** Create the VM in vCenter and attach the created portgroup to the VM.
This will attach the VM to the network created via Cisco VTS.
-

Creating a Network using Cisco VTS GUI

To create a network from the Cisco VTS GUI:



Important

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

-
- Step 1** Go to **Overlay > Network**. The Overlay / Network window appears.
- Step 2** Click Fabric Host Networks or External Networks, based on your need.
- Step 3** Click **Add (+)** icon.
- Note** For External Networks you need to specify the Name, Tenant, and Zone.
- Step 4** Enter the network name. This is mandatory.
- Step 5** Select the Tenant for which you to create the network.
- Step 6** Select the Zone.
- Step 7** If the network is not external, enter the Static VNI. This can be an integer between 4096 and 65535.

- Step 8** Enter the **MTU** value for the port. This can be an integer between 576 and 9216. The entered MTU value is pushed to BDI/SVI interfaces during port attach. The default value of MTU is empty.
- Step 9** Specify whether the network is a shared network.
- Step 10** Specify whether the network has to be Extended. If you select Yes, VPN Service becomes available for use.
- For the L2 Extended Network, click the **L2VPN** tab and enter the EVI number. This can be an integer between 1 and 65534. Select the Load Balance Per EVI check box to introduce the load balance CLI in the device. See [Extending Layer 2 Network Across Data Centers, on page 162](#).
- Step 11** Specify whether ARP Suppression needs to be enabled.
- Step 12** Click **Save**.
- You can also add a subnet, and add port.

Creating a Subnetwork

To create a subnetwork:

- Step 1** Click **Add (+)** in the Subnet pane of the Add Network page.
- Step 2** Enter the subnet name. Only IPv4/IPv6 addresses, alphabets, space, numbers, and special characters /, - and _ are allowed.
- Step 3** Enter the IP details. You can enter an IPv4 or IPv6 address. You must ensure that the network address and the gateway IP are in sync.
- You can create subnets with /31 prefix. /31 subnet masks are used for point to point links. The gateway IP for a subnet with /31 prefix should be within the two allowed IP addresses. For Example:
- For subnet 10.20.30.0/31, the allowed IPs are 10.20.30.0 and 10.20.30.1
 - For subnet 10.20.30.5/31, the allowed IPs are 10.20.30.5 and 10.20.30.4
- Note** When you have a network with /31 subnet, then you cannot make it as an external network and vice versa.
- Step 4** Click **OK**.
- The table displays the Subnet Name, Network Address, Gateway IP, and the IP Version (whether IPv4 or IPv6).

Creating Router using Cisco VTS GUI

To create a router using Cisco VTS GUI:



Important

If you attach VTS subnets (Baremetal) to a router from Cisco VTS GUI, and then attach the OpenStack subnets to the same router from the Cisco VTS GUI, all subsequent operations on these subnets need to be done from Cisco VTS. You can specify route-map name for both address families ipv4/ ipv6 under router bgp. Only unicast needs to be supported in address family.

- 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 116](#) for details.
- Step 8** Enter a VRF name. This is optional. If this is left empty, when the **Save** button is clicked, a default VRF name gets automatically generated.
- The custom VRF name accepts up to 24 characters.
 - If there is no input for custom VRF name, a default VRF name gets generated in form of `<tenant-name>-<router-name>`. Both tenant-name and router-name accept up to 15 characters.
 - If Cisco ASR 9000 series router is configured as DCI in the domain, and you have not given a custom VRF name, then you must ensure that the default VRF name does not exceed 27 characters. Otherwise, the configuration will fail.
 - If VTSR is configured, and you have not given a custom VRF name, then you must ensure that the default VRF name does not exceed 24 characters. Otherwise, the configuration will fail.
 - If configuration fails because the default VRF name exceeds the limit, you can choose to use custom VRF name instead.
 - If the configuration fails because the default VRF name exceeds 27 characters, an error message appears on the Network > Port Attach screen, which indicates invalid input for “bridge-domain” configured on Cisco ASR 9000 series router.
 - For VTSR configuration, a similar error is displayed if the default VRF name exceeds 24 characters.
 - VRF name change from VTS GUI is not supported for VTSR. Cisco VTS does not allow changing the name of a router if it connects to a port on a V node. (A V node is compute node where there is a VTF present, and the workload is behind a VTF where the VXLAN Tunnel originates.)
 - If you modify the VRF name after saving the router, the Router Gateway IP address gets removed. You can reconfigure it back after saving the VRF name change.
- Step 9** Select the router gateway from the **Router Gateway** drop-down list. When you select External GW from drop-down list, two additional fields for Router Gateway IPv4 and Router Gateway IPv6 get displayed. These are optional.

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

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

```

router bgp 65539
vrf t1-rohi-rout
  rd 2.2.2.11:10009
  address-family ipv4 unicast
    label mode per-vrf
    maximum-paths ebgp 2
    maximum-paths ibgp 2
    network 0.0.0.0/0
  aggregate-address 3.2.3.0/24 summary-only
  redistribute connected
exit
address-family ipv6 unicast
  label mode per-vrf
  redistribute connected

```

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

Step 10 Choose the values in the **Route Map IPv4** and **Route Map IPv6** fields to capture/fetch route-map ipv4 and ipv6 information. This allows a custom route-map to be specified during redistribution in BGP.

When static route is configured, VTS will automatically add redistribute static under BGP with a specific route-map (For example by name : vts-static-route-map-ip). Static route configured under first class object allows a custom route-map to be specified per BGP peer for policy control during redistribution in BGP.

Step 11 If the router is used to add shared networks from different tenants as interfaces, set the **Provider Router** toggle switch to **Yes**.

Step 12 Choose a **Maximum Path** from the drop-down list and its values. Specify the value for max-paths and the options are none/eibgp/ibgp/mixed and no. of max. paths.

Example: Sample Configuration :

```

Sample Configuration:
router bgp <ASN>
vrf <VRF_NAME>
address-family <ipv4/v6> unicast
export-gateway-ip
maximum-paths mixed <1-64>
Where VALUE can be:
eibgp Configure multipath for both EBGp and IBGP paths
ibgp Configure multipath for IBGP paths
mixed Configure multipath for local and remote paths

```

The number of max paths is limited to 1-64. For N7K, the number of max paths is limited to 1-32. If maximum-paths value is mixed, then export-gateway-ip will be configured automatically.

Note For N9k (Release : 9.2(1)), Setting **eibgp** value in maximum-paths will fail to configure because, the VRF (Router Interface) is by default created with label mode “per-vrf”, and the error message is shown as “Cannot configure EIBGP multipath alongwith per-vrf label mode”.

View the label mode using the cli command “show bgp l3vpn detail vrf <vrf-name>” on device. Ensure that you disable the external/internal Border Gateway Protocol (BGP) multipath feature if it is enabled before you configure the per-VRF label allocation mode.

At the time of creating a Router, if the router interface is associated and the maximum-paths value is **eibgp**, then error is displayed in case of per-vrf. However, in the Router window you can still view the value "eibgp". This is because Service model and device model/device have two different api calls.

When you Edit the Router details, if maximum-paths value is modified from <any value> to eibgp and vrf is with per-vrf label mode, then device rejects the request. Router UI will reset to none and in the device the maximum-paths value is cleared. This is because, In N9K OS, the earlier value is modified with a new value.

- Step 13** To add a Log-neighbor-changes for each router, use the **Log-neighbor-changes** toggle switch. By default the toggle switch is ON.
- Step 14** Enter the **MTU** value for the port. This can be an integer between 576 and 9216. The entered MTU value is pushed to BDI/SVI interfaces during port attach (router interface should be added). The default value of MTU is 9100 and you can change it between the given range.
- Step 15** Use the **IPv4 IP Redirect** and **IPv6 IP Redirect** toggle switches to enable/disable the IPv4 and IPv6 redirects at the interface level, respectively. By default, the toggle switch is OFF.
- Step 16** Click **Add (+)** icon. The Add Interface dialog box appears.
- Step 17** Choose the subnet from the drop-down list, and click **OK**.
- Step 18** Click **Save** in the Add Router window to save the router and its interface.
-

Port Extensions Support

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

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

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

Creating a Port Extension

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

- Step 1** Go to **Overlay > Port Extensions**. The Overlay / Port Extensions window appears.
- Step 2** Select a tenant from the Tenant drop-down. The following details are displayed.
- Port Extension Name—The Port Extensions that have been created for the tenant.
 - Description
 - Type
 - Zone
 - Attached Ports
- Step 3** Click **Add (+)** icon to create a new Port Extension.

- a. In the Add Port Extension pane, enter the following details:
 - Port Extension Name—This is a mandatory field.
 - Description
 - Type—Type it BGP.
 - Tenant—The tenant under which the Port Extension is being created.
 - Zone—Select any zone from the drop-down.

- b. BGP Profile Information—BGP Profile Information (VRF Config) contains details that will influence the type and nature of the BGP peering sessions initiated by the TOR (connected to overlay ports) towards VM instances. Enter the following details:
 - Route Reflector Mode—Applicable in iBGP scenarios. Select **None** for eBGP. The following options are available:
 - None
 - Client

- c. In the Neighbor List pane, you can either add or edit the neighbor list. Click **Add (+)**.

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

The Add Neighbor page appears with the following details:

1. BGP Neighbor Information:

- Neighbor Id—Neighbor IP to which the BGP session needs to be established. This is a mandatory field, there is a format.
- Description
- Neighbor ASN—Applies only in eBGP case.
- Local ASN—Applies only in eBGP case.

Note If you enter a value for Neighbor ASN, then Local ASN value can not be the same as that of Neighbor ASN.
- Local Source Loopback Number
- Passphrase—Password to establish the BGP session with the neighbor.
- Suppress 4-byte ASN—Suppress 4-byte AS Capability.
- BFD—Bidirectional Fast Detection for the neighbor.
- eBGP-Multihop—Number of hops the eBGP peer is away. For directly connected peers, leave this field empty.

Note All eBGP fields need to be removed before moving the neighbor session from eBGP type to iBGP. Due to platform limitation, to switch from eBGP to iBGP with an attach Port Extension, you need to follow the platform flow with the following steps:

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

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

2. Address Family List:

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

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

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

- Type—Route Map or Prefix List.
- Name
- Direction—filter-in or filter-out.

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

- Step 4** Click **Add** to add the neighbor details.
- Step 5** Click **Save**. The Overlay / Port Extensions page appears where you can see that the Port Extension is created successfully.
-

Editing a Port Extension

You can modify a Port Extension that you have created.

- Step 1** Go to **Overlay > Port Extension**. The Overlay / Port Extensions window appears.
- Step 2** Select the check box corresponding to the Port Extension you need to edit, and click the **Edit** icon.
- Step 3** Make the desired changes in the attributes.
- Note** You can edit the Port Extension name only for the ones that are not attached.
You can make only four entries to the Address Family List.
- Step 4** Click **Add (+)** icon to add any number of Route Filter Lists based on your requirement. Click the **Remove (-)** icon to remove any Route Filter List.
- Step 5** Click **Add** to add the neighbor details.
- Step 6** Click **Update**.
- Step 7** Click **Save**. The Overlay / Port Extensions page appears where you can see that the Port Extension is updated successfully.
-

Deleting a Port Extension

You can delete a Port Extension that you have created.

- Step 1** Go to **Overlay > Port Extension**. The Overlay / Port Extensions window appears.
- Step 2** Select the check box corresponding to the Port Extension you need to delete, and click the **Delete (x)** icon.
- Note** You will not be able to delete a Port Extension that is attached to any ports. You need to detach the Port Extensions from those ports and then delete the Port Extension.
- Step 3** Click **Yes** to delete the selected Port Extension that does not have any ports attached.
-

Attaching Port Extension to Baremetal Ports

You can attach a single port extension to one or multiple Baremetal ports. Select either Zone or VRF or Network or Device filters to view list of Baremetal ports. To reduce delays in fetching ports for association with Port extensions and to view sub-set of ports you can use either Network or Device filters.

Before you begin

You need to have a Port Extension created already.

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

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

The table shows the following details:

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

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

Step 4 Select the **Zone**.

Step 5 Choose a VRF from the **VRF** drop-down list.

Note Port extensions can be attached to a port only after the VRF has been created. It cannot be attached to a port with no VRF. At the time of adding a port, Port Extension option will not be available.

Step 6 Choose a network from the **Network** drop-down list for a VRF. For the selected network, list of corresponding devices are listed.

Step 7 Choose a device from the **Device** drop-down list for a network. For a selected device, view list of matching ports on that device.

The list of available ports gets displayed for the selected VRF, network and device.

Step 8 Select a port (s) that you want to attach from the displayed list. For the selected port(s), view the port extension association list to associate with port(s).

Step 9 Select the Port Extension to attach to from the **Port Extension** drop-down list.

Step 10 Click **Attach**. Baremetal ports page is displayed with the attached Port Extension displayed as a link.

Detaching Port Extension from Baremetal Ports

To detach a Port Extension from Baremetal Ports:

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

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

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

- Step 4** Select the **Zone**.
- Step 5** Select a VRF from the **VRF** drop-down.
- Step 6** Select the ports that you want to detach from the displayed list.
- Step 7** Click **Detach**.
- Step 8** Click **Yes** to confirm.

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

Attaching Port Extension to Virtual Machine Ports

You can attach a single port extension to one or multiple Virtual Machine (VM) ports. Select either Zone or VRF or Network or Device filters to view list of Baremetal ports. To reduce delays in fetching ports for association with Port extensions and to view sub-set of ports you can use either Network or Device filters.

Before you begin

You need to have a Port Extension created already.

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

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

The table shows the following details:

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

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

Step 4 Select the **Zone**.

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

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

Note Port extensions can be attached to a port only after the VRF has been created. It cannot be attached to a port with no VRF. At the time of adding a port, Port Extension option is not available.

Step 6 Choose a network from the **Network** drop-down list for a VRF. For the selected network, list of corresponding devices are listed.

- Step 7** Choose a device from the **Device** drop-down list for a network. For a selected device, view list of matching ports on that device.
- Step 8** Select the ports that you want to attach from the displayed list.
- Step 9** Select the Port Extension to attach to from the **Port Extension** drop-down.
- Step 10** Click **Attach**. Virtual Machine ports page is displayed with the attached Port Extension displayed as a link.
-

Detaching Port Extension from Virtual Machine Ports

To detach a Port Extension from Virtual Machine Ports:

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

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

Assigning BVI Interface IP Address

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

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

- Step 10** Assign the **Router Gateway IP address** for the selected external network for BVI interface and click **Save**.
- Step 11** Verify whether the configuration is pushed to DCI and the IP address is assigned to BVI interface.

Extending Layer 2 Network Across Data Centers

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

To complete the L2VNI extension workflow:



Note

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

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

Enabling Global Route Leaking Service

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



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

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

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

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

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

See [Synchronizing Configuration, on page 75](#) 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 105](#).

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

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

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

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

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

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

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

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

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

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

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

Step 11 Click **Add**. The Global Route Leaking service gets added. You can click on the name to get a summary of the global route leaking service you created.

Step 12 Click **Save**. Once the service is attached to the router, all the networks for the router will be leaked outside. To disassociate the service you need to select the **Detach** button and save the edit.

Enabling L3VPN to EVPN Route Stitching

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



Note As a prerequisite, you must create an external network and extend to L3. You must then attach the router interfaces to the external network. See [Creating a Network using Cisco VTS GUI, on page 151](#) and [Creating Router using Cisco VTS GUI, on page 152](#) sections for details.

Step 1 Configure BGP VPNv4/v6 neighbor using Device Templates. A single template can be used for all the neighbors, or you can have a template each for each neighbor. Create the template at **Templates > Device Template Management**. Attach the template to the DCI. See [Managing Templates and Device Objects, on page 107](#) chapter for details.

Step 2 Create an External Route Stitching Template. Choose the routes which you want to leak between your core and EVPN, or vice versa. Create the template at **Templates > Overlay Template Management** (use the Fabric External RT option). Attach the template to the DCI.

Adding Static Routes

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

Add the following types of static routes:

- Internal Static Route—See [Adding Fabric Static Routes, on page 165](#) for details.
- External Static Route—See [Adding External Static Routes, on page 165](#) for details.
- Port Static Route—See [Adding Port Static Routes, on page 166](#) for details.

If the static route is marked external, then it gets applied to the VRF on Border Leaf/DCI.

This release supports Static Route - BFD feature which enables association of static routes with a static Bidirectional Forwarding Detection (BFD) configuration.

During upgrade, all the routes in the route templates would be converted to static route objects and pushed to all leaf devices. Post-upgrade, if you need to modify these, you can remove these static route objects.

Adding Fabric Static Routes

Static Routes can be configured directly on the router, and it will be pushed to all the nodes that have vrf for that router. You can track an object in a static route and name a static route (name of next hop) for first class static routes.

-
- Step 1** Go to **Overlay > Router**. Select the router you need to edit, and click the **Edit** icon.
 - Step 2** Click **Static Route**, then select **Fabric Static Route** from the Static Route Scope pane.
 - Step 3** Click **Add (+)** in the Internal Static Route pane, to add an internal static route. The New Internal Static Route popup is displayed.
 - Step 4** Enter the **Destination Prefix**—Destination IP address and the Subnet Mask.
 - Step 5** Enter the **Static Route Tag**. This is the BGP tag which is used for redistribution. This is optional.
 - Step 6** Enter the name in the **Static Route Name** field.
 - Step 7** Enter the **Next Hop Details**—IP Address and Subnet Mask.
 - Step 8** Specify the **VRF**, if a router need to send traffic to a specific destination, via a specific next hop, and if that next hop is reachable only via a certain VRF.
 - Step 9** Enter the track ID in the **Track** field.
 - Step 10** Specify the **Weight/Preference**, when you add multiple next hops.
 - Note** If you want the packet to a destination to be dropped, you can configure a null interface, by selecting the **Drop traffic that matches Destination Prefix** radio button.
 - Step 11** Click **Save**. The static route will be pushed to all the nodes.

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

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

Adding External Static Routes

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



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

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

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

Adding Port Static Routes

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

Port Scoped Static Route and Fabric static Route will be supported for virtual VTEP only when VTSR is configured as VTEP. VTSR will not support multiple next hops static routes. You can also set VRF for port-scoped static route.



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

When static route is configured, VTS will automatically add redistribute static under BGP with a specific route-map (For example by name : vts-static-route-map-ip) . Static route configured under first class object allows a custom route-map to be specified per BGP peer for policy control during redistribution in BGP.

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

- Static Route Tag—This is an optional parameter.
- Static Route Name — Enter a name for the static route.
- Next Hop Details
 - Port ID—The ID of the port (VM) which you want to reach. You can select the subnet and choose from the ports on the interfaces.
 - Next Hop Prefix
 - Weight—Specify the preference here. This is optional.
 - BFD—Enable Bidirectional Fault Detection. This is supported only on Cisco Nexus 9000 Series devices. Port Scope Static Route supports BFD on both VTF (V-side) and (P-side).
- Enter the **VRF** name and the **Track** object ID for static routes.

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

When you create a router and populate Route Map Ipv4 and Route Map Ipv6 information, you can view the configuration changes and the outer payload that includes the route-map value, under the device model.

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

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

```
vtsr01# show running-config vtsr-config ip-routes
vtsr-config ip-routes ip-route admin-rtr-1 217.217.217.217/32
scope port
paths path 101.1.1.67
  next-hop-vrf admin-rtr-1
  bfd true
  interface vhost-f8d76e3b-8943-4cde-90f1-df8141274aad
!
```

OpenStack Allowed Address Pairs Support

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

Important Notes:

- VPP adds default allowed address pair:



Note This is specific to VTF.

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



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



CHAPTER 13

Viewing Overlay Details

You can view details about the network at the device, VM, baremetal, network, and router levels. This chapter has the following sections:

- [Viewing Device Details, on page 169](#)
- [Viewing Virtual Machine Details, on page 170](#)
- [Viewing Baremetal Port Details, on page 170](#)
- [Viewing Network Details, on page 171](#)
- [Viewing Router Details, on page 172](#)

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 Port ID
 - Status
 - Network Name
 - Binding Host
 - SRIOV Enabled
 - VLAN
 - Connected Host
 - Security Groups—Click the **View** link in the table to view the details of the Security Group(s) attached to the Virtual Machine.
-

Viewing Baremetal Port Details

To view the Baremetal Port details:

Go to **Overlay > Baremetal**. The Overlay / Baremetal window appears.

The following details are displayed:

- Baremetal Port ID
- Status
- Network Name
- Binding Host
- VLAN Number

- Security Groups—Click the **View** link to view the details about the Security Group(s) attached to the Baremetal Port.
-

Viewing Network Details

To view details about the network:

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

Step 2 Choose the source from the Source drop-down. Then select the tenant from the Tenant drop-down. The following details are displayed for the tenant you select:

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

You can filter to display all networks or only shared networks using the Select Network drop-down.

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

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

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

Viewing Router Details

To view details about the routers in the network:

-
- Step 1** Go to **Overlay > Router**. The Overlay / Router window appears.
- Step 2** Select the source from the Select Source drop-down, then select the tenant from the Select Tenant drop-down..

The following details are displayed:

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

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

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



APPENDIX **A**

Service Extension Templates-Supported Configuration Examples

This appendix has the following sections.

- [Supported Service Extension Template Configuration Examples for Cisco Nexus 7000 Series Switches, on page 173](#)
- [Supported Service Extension Template Configuration Examples for Cisco Nexus 9000 Series Switches, on page 176](#)
- [Supported Service Extension Template Configuration Examples for Cisco ASR 9000 Series Routers , on page 179](#)

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

This section provides service extension configuration examples for Cisco 7000 series switches.

Interface Configuration

```
interface port-channel 6.1001

    description towards PE
    encapsulation dot1q 1001
    vrf member <demo_name>
    bfd interval 50 min_rx 50 multiplier 3
    no ip redirects
    ip address 10.10.10.10/24
    no ipv6 redirects
    no shutdown

interface port-channel 6.1041

    description towards PE
    encapsulation dot1q 1041
    vrf member <demo_name>
    bfd interval 50 min_rx 50 multiplier 3
    no ip redirects
    ipv6 address 10:10:10:10:10:10:10:10/64
    no ipv6 redirects
    no shutdown
```

```

interface vlan 2471
  no shut
  vrf member <demo_name>
  ip address 10.10.10.10/24
  vrrpv3 1 address-family ipv4
    address 10.10.10.10

```



Note This configuration works for IPv4 and IPv6.

```

interface vlan 2475
  no shut
  vrf member <demo_name>
  ipv6 address 10:10:10:10:10:10:10:10/64
  vrrpv3 1 address-family ipv6
    address fe80::1 primary
    address 10:10:10:10:10:10:10:10

```



Note This configuration works for IPv4 and IPv6.

```

interface vlan 2400
  no shut
  vrf member <demo_name>
  ip address 10.10.10.10/24
  ipv6 address xxxx::x/64 ! dual stack or create a difference interface

```



Note This configuration works for IPv4 and IPv6.

VRF Configuration

```

vrf context <demo_name>

    ip route 10.32.10.0/24 10.42.10.4
    ip route 10.52.10.0/24 10.42.10.4
    ipv6 route 2001:db8::5/128 2001:db8:0:1:2a0:a502:0:19da

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

```



Note This works configuration for IPv4 and IPv6.

```

vrf context <demo_name>
  vni 50001
  ip route static bfd Vlan1050 68.50.50.50
  ip route 210.0.0.1/32 Vlan1050 68.50.50.50

```




Note This works configuration for IPv4 and IPv6. User should enter the vlan value other than 1000-2000, which is reserved for bridge domain. If the range is configured differently, make sure the values within that range is not used.

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



Note This configuration works for IPv4 and IPv6.

Router BGP Configuration

```
router bgp 65537
vrf <demo_name>
  local-as 65539
  address-family ipv4 unicast
    network 10.32.10.0/24 route-map <demo_name>_LOCAL_COMMUNITIES
    network 10.52.10.0/24 route-map <demo_name>_LOCAL_COMMUNITIES
    advertise l2vpn evpn
  neighbor 10.23.65.0 remote-as 65541
  bfd
  password 3 XXXX
  description towards PE
  address-family ipv4 unicast
    send-community
    route-map <demo_name>_ROUTE_POLICY in
    route-map <demo_name>_LOCAL_ROUTE_POLICY out

router bgp 65539
Vrf <demo_name>
  router-id 192.168.0.25
  address-family ipv4 unicast
    network 150.0.0.1/32 route-map ONLY_FABRIC
    advertise l2vpn evpn
  redistribute direct route-map vts-subnet-policy augmentation and deviation
  redistribute static route-map staticMap
  maximum-paths 32
  maximum-paths ibgp 32
  address-family ipv6 unicast
    advertise l2vpn evpn
  redistribute direct route-map vts-subnet-policy
  maximum-paths 32
  maximum-paths ibgp 32
  neighbor 68.50.50.50
  bfd
  remote-as 65538
  address-family ipv4 unicast
    send-community
    send-community extended
  neighbor 210.0.0.1
  bfd
  remote-as 65538
  update-source loopback150
  ebgp-multihop 255
  address-family ipv4 unicast
```

```

send-community
send-commuqnity extended

```



Note This configuration works for IPv4 only.

ICMP v6 Configuration

```

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

```



Note This configuration works for IPv6 only.

Interface Loopback Configuration

```

interface loopback1
vrf member <demo_name>

```



Note This configuration is done in L3 Service Extension.

BFD Multi-hop Configuration

```

router bgp 100
vrf abcs
neighbor <IPv4 or v6 address>
bfd multihop interval 900 min_rx 900 multiplier 4

```



Note This configuration works for IPv4 and IPv6.

Supported Service Extension Template Configuration Examples for Cisco Nexus 9000 Series Switches

This section provides service extension configuration examples for Cisco 9000 series switches.

Interface Configuration

```

int vlan 2471
no shut
vrf member <demo_name>
ip address 10.10.10.10/24
vrrpv3 1 address-family ipv4
address 10.10.10.10

```



Note This configuration works for IPv4 and IPv6.

```
int vlan 2475
  no shut
  vrf member <demo_name>
  ipv6 address 10:10:10:10:10:10:10:10/64
  vrrpv3 1 address-family ipv6
    address fe80::1 primary
    address 10:10:10:10:10:10:10:10
```



Note This configuration works for IPv4 and IPv6.

```
interface vlan 2400
  no shut
  vrf member <demo_name>
  ip address 10.10.10.10/24
  ipv6 address xxxx::x/64 ! dual stack or create a difference interface
```

VRF Configuration

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



Note This works configuration for IPv4 and IPv6.

```
vrf context <demo_name>
  vni 50001
  ip route static bfd Vlan1050 68.50.50.50
  ip route 210.0.0.1/32 Vlan1050 68.50.50.50
```



Note This works configuration for IPv4 and IPv6.

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



Note This configuration works for IPv4 and IPv6.

Router BGP Configuration

```

router bgp 65539
Vrf <demo_name>
  router-id 192.168.0.25
  address-family ipv4 unicast
    network 150.0.0.1/32 route-map ONLY_FABRIC
    advertise l2vpn evpn
  redistribute direct route-map vts-subnet-policy
  redistribute static route-map staticMap
  maximum-paths 32
  maximum-paths ibgp 32
  address-family ipv6 unicast
    advertise l2vpn evpn
  redistribute direct route-map vts-subnet-policy
  maximum-paths 32
  maximum-paths ibgp 32
  neighbor 68.50.50.50

  bfd
  remote-as 65538
  address-family ipv4 unicast
    send-community
    send-community extended
  neighbor 210.0.0.1
  bfd
  remote-as 65538
  update-source loopback150
  ebgp-multihop 255
  address-family ipv4 unicast
    send-community
    send-commuqnity extended

```



Note This configuration works for IPv4 only.

Interface Loopback Configuration

```

interface loopback1
  vrf member <demo_name>

```



Note This configuration is done in L3 Service Extension.

BFD Multi-hop Configuration

```

router bgp 100
  vrf abcs
  neighbor <IPv4 or v6 address>
    bfd multihop interval 900 min_rx 900 multiplier 4

```



Note This configuration works for IPv4 and IPv6.

Supported Service Extension Template Configuration Examples for Cisco ASR 9000 Series Routers

This section provides service extension configuration examples 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 <demo_name>
!
```

Router Static Configuration

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

Router BGP Configuration

```
router bgp 65540
  bgp router-id 49.1.1.1
  address-family ipv4 unicast
    maximum-paths ebgp 2
    maximum-paths ibgp 2
  !
  neighbor-group ng1
    remote-as 65539
    password encrypted 07982c55db2b9985d3391f02e639db9c
    update-source Loopback0
    address-family ipv4 unicast
      next-hop-self
  !
  vrf <demo_name>
    rd auto
    bgp router-id 49.1.1.1
    address-family ipv4 unicast
  !
  neighbor 13.1.1.8
    remote-as 65539
    address-family ipv4 unicast
    route-policy vts-route-policy in
    default-originate
```

```

!
!
!
!

```

VRF Configuration

```

vrf <demo_name>
  address-family ipv4 unicast
!
!

```

Interface/{any}-subinterface Configuration

```

interface GigabitEthernet0/0/0/1.1
  vrf <demo_name>
  ipv4 address 10.10.10.10
  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 <demo_name>
!
!

```

Interface NVE Configuration

```

interface nve1
  description desc123
  vrf <demo_name>
  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 Examples

This appendix provides examples of supported underlay template configuration.

- [Supported Underlay Configuration Examples, on page 181](#)

Supported Underlay Configuration Examples

Configuration Area	Sample Configuration	Device Type	Device Role
<p>Note Prerequisite: You should enable the following configuration for the device configuration to work.</p>	<pre>feature telnet feature nxapi feature ospf feature bgp feature pim feature udld feature interface-vlan feature vn-segment-vlan-based feature hsrp feature lacp feature vpc feature lldp feature nv overlay feature pbr feature sla sender feature sla responder feature vrrpv3 feature bfd</pre>		
UNDERLAY IGP ROUTING OSPF routing process	<pre>router ospf 10</pre>	N9K	Leaf
	<pre>router ospf 10</pre>	ASRK	DCI
OSPF Area	<pre>interface ethernet 1/5 ip ospf router 10 area 0.0.0.0</pre>	N9K	Leaf
	<pre>interface ethernet 1/5 ip ospf router 10 area 0.0.0.0</pre>	N9K	Leaf
	<pre>router ospf 10 area 0</pre>	ASRK	DCI

OSPF router-id	router ospf 10 router-id 10.218.20.15	N9K	Leaf
	router ospf 10 router-id 10.218.20.15	ASPK	DCI
OSPF auto-cost reference	router ospf 10 ! auto-cost reference-bandwidth 800000	ASPK	DCI
OSPF Network type	interface ethernet1/5 ip ospf network point-to-point	N9K	Leaf
	interface vlan10 ip ospf network point-to-point	N9K	Leaf
	router ospf 10 area 0 interface GigabitEthernet0/0/1/3 network point-to-point	ASPK	DCI
OSPF Authentication	interface Ethernet1/5 ip ospf authentication message-digest	N9K	Leaf
	interface Ethernet1/5 ip ospf message-digest-key 1 md5 0 xxx	N9K	Leaf
	router ospf 10 area 0 interface <Fabric Interface> authentication message-digest message-digest-key 1 md5 encrypted 202cb962ac59075b964b07152d234b70	ASPK	DCI
OSPF Passive-interface	interface loopback3 ip router ospf 100 area 0.0.0.0	N9K	Leaf
	router ospf 10 area 0 interface Loopback10 passive enable	ASPK	DCI
OSPF Convergence	router ospf 10 timers lsa arrival 15 timers throttle lsa 0 20 5000 timers throttle spf 50 100 5000	N9K	Leaf
	router ospf 10 timers throttle lsa all 0 20 5000 timers throttle spf 50 100 5000 timers lsa min-arrival 15	ASPK	DCI
OSPF BFD (per-link)	feature bfd router ospf 10 bfd	N9K	Leaf
	interface Ethernet1/5 no ip redirects	N9K	Leaf

	<pre>router ospf 10 bfd minimum-interval 150 bfd multiplier 3 area 0 interface TenGigE0/0/2/1 bfd fast-detect</pre>	AS9K	DCI
	<pre>interface vlan 10 no bfd echo</pre>	N9K	Leaf
Multicast Routing	<pre>feature pim</pre>	N9K	Leaf
	<pre>interface loopback1 ip address 10.10.10.10/24 ip router ospf 10 area 0.0.0.0 ip pim sparse-mode</pre>	N9K	Spine
	<pre>ip pim rp-address 10.218.20.250 group-list 239.255.0.0/16 override</pre>	N9K	Spine
	<pre>ip pim anycast-rp 10.218.20.250 10.218.20.249 ip pim anycast-rp 10.218.20.250 10.218.20.248</pre>	N9K	Spine
	<pre>feature pim</pre>	N9K	Leaf
	<pre>ip pim rp-address 10.218.20.250 group-list 239.255.0.0/16 override</pre>	N9K	Leaf
	<pre>interface Vlan10 ip pim sparse-mode</pre>	N9K	Leaf
	<pre>interface loopback0 ip pim sparse-mode</pre>	N9K	Leaf
	<pre>interface Ethernet2/1 ip pim sparse-mode</pre>	N9K	Leaf
	<pre>interface Ethernet2/2 ip pim sparse-mode</pre>	N9K	Leaf
L2 Technologies	<pre>interface Ethernet 1/10 switchport mode trunk</pre>	N9K	Leaf
	<pre>interface Ethernet 1/10 switchport trunk allowed vlan none</pre>	N9K	Leaf
	<pre>interface Ethernet 1/10 spanning-tree port type edge trunk</pre>	N9K	Leaf
	<pre>interface Ethernet 1/10 spanning-tree bpduguard enable</pre>	N9K	Leaf
	<pre>interface Ethernet 1/10 spanning-tree bpdufilter enable</pre>	N9K	Leaf
	<pre>interface Ethernet 1/10 storm-control broadcast level 20.0</pre>	N9K	Leaf
	<pre>interface Ethernet 1/10 storm-control multicast level 30.0</pre>	N9K	Leaf
	<pre>interface Ethernet 1/10 storm-control unicast level 50.0</pre>	N9K	Leaf

	interface Ethernet 1/10 storm-control action shutdown	N9K	Spine
vPC Role and Priority	vpc domain 1 role priority 100	N9K	Leaf
	vpc domain 1 role priority 200	N9K	Leaf
vPC Peer Keep-alive Link	vrf context management	N9K	Leaf
	interface mgmt 0 vrf member management	N9K	Leaf
	interface mgmt 0 ip address 10.10.10.10/24 no shutdown	N9K	Leaf
	vpc domain 1 peer-keepalive destination 172.20.118.20	N9K	Leaf
vPC Peer-Link	interface Ethernet 1/1 spanning-tree port type network channel-group 1 mode active no shutdown	N9K	Leaf
	interface Ethernet 1/2 spanning-tree port type network channel-group 1 mode active no shutdown	N9K	Leaf
	interface port-channell switchport switchport mode trunk spanning-tree port type network vpc peer-link	N9K	Leaf
vPC Port	interface Ethernet 2/9 channel-group mode active id 51	N9K	Leaf
	interface port-channel 51 switchport	N9K	Leaf
	interface port-channel 51 switchport vpc 51	N9K	Spine
vPC Peer-Switch Option	vpc domain 1 peer-switch	N9K	Leaf
vPC ARP Synchronization	vpc domain 1 ip arp synchronize	N9K	Leaf
vPC in VXLAN environment adjustment	vpc domain 10 peer-switch system-priority 100 (could not find this option) peer-keepalive destination 172.20.118.120 delay restore 200 peer-gateway ip arp synchronize	N9K	Leaf
	interface port-channell description vPC peer-link	N9K	Spine

	interface port-channel1 description vPC switchport mode trunk	N9K	Leaf
	interface port-channel 1 description vPC switchport mode trunk	N9K	
	interface port-channel 1 description vPC spanning-tree port type network	N9K	
	interface port-channel 1 vpc peer-link	N9K	
	interface port-channel 10 switchport trunk allowed vlan	N9K	
	interface port-channel 10 spanning-tree port type edge trunk	N9K	
	interface port-channel 10 spanning-tree bpdupfilter enable	N9K	
	interface port-channel 10 spanning-tree bpduguard enable	N9K	
	interface port-channel 10 vpc 10	N9K	Spine
	interface Ethernet 1/10 switchport trunk allowed vlan none	N9K	Spine
	interface Ethernet 1/10 spanning-tree port type edge trunk	N9K	Leaf
	interface Ethernet 1/10 spanning-tree bpduguard enable	N9K	Leaf
	interface Ethernet 1/10 spanning-tree bpdupfilter enable	N9K	Leaf
	interface Ethernet 1/10 channel-group 10 mode active	N9K	Leaf
	interface loopback 0 ip address 10.10.10.10/24	N9K	Leaf
	interface loopback 0 ip address 10.10.10.10/24 secondary	N9K	Leaf
	interface loopback 0 ip router ospf 100 area 0.0.0.0	N9K	Leaf
	interface loopback 0 ip pim sparse-mode	N9K	Leaf
	interface Vlan 10 ip address 10.10.10.10/24	N9K	Leaf
	interface Vlan 10 description Underlay vPC Backup link no shutdown no bfd echo	N9K	Spine
	interface Vlan 10 ip ospf network point-to-point	N9K	Leaf

	interface Vlan 10 ip router ospf 100 area 0.0.0.0	N9K	Leaf
	interface Vlan10 ip pim sparse-mode	N9K	Leaf
STP	interface Ethernet 1/10 switchport mode trunk	N9K	Leaf
	interface Ethernet 1/10 switchport mode trunk allowed vlan 10	N9K	Leaf
	interface Ethernet 1/10 spanning-tree port type edge trunk	N9K	Leaf
	interface Ethernet 1/10 spanning-tree bpduguard enable	N9K	Leaf
	nx:interface Ethernet 1/10 spanning-tree bpdufilter enable	N9K	Leaf
	interface Ethernet 1/10 no shutdown	N9K	
	interface port-channel 10 switchport mode trunk	N9K	Leaf
	interface port-channel 10 switchport mode trunk trunk allowed vlan ids 1	N9K	Leaf
	interface port-channel 10 spanning-tree port type edge	N9K	Leaf
	interface port-channel 10 spanning-tree bpduguard enable	N9K	Leaf
	interface port-channel 10 spanning-tree bpdufilter enable	N9K	Leaf
	interface port-channel 10 no shutdown	N9K	Leaf
	interface port-channel 10 vpc port-channel-number 10	N9K	Leaf
	interface Ethernet 1/10 switchport mode trunk	N9K	Leaf
	interface Ethernet 1/10 switchport mode trunk allowed vlan 10	N9K	Leaf
	interface Ethernet 1/10 spanning-tree port type edge trunk	N9K	Leaf
	interface Ethernet 1/10 spanning-tre guard root	N9K	Leaf
	interface Ethernet 1/10 no shutdown	N9K	Leaf
	interface ethernet <xxxx> description <leaf/Spine Fabric> ip address	N9K	Leaf

	interface Ethernet 1/10 description leaf mtu 9216	N9K	Leaf
	interface Vlan 1 description <attachment/border facing intf>	N9K	Leaf
	interface Vlan 1 description ip address <addr>	N9K	Leaf
	interface Vlan 1 description ip address mtu 1500	N9K	Leaf
	interface <To Spine> mtu 9214	N9K	Leaf
	interface <To Border Leaf> mtu 1518	N9K	Leaf
	interface GigabitEthernet0/0/1/5 mtu 9214	ASR9K	DCI
Nexus 9500 QoS	system qos service-policy type queuing output default-out-policy	N9K	Leaf
	policy-map type network-qos Jumbo-nq-policy class type network-qos c-nq3	N9K	Leaf
	policy-map type network-qos Jumbo-nq-policy class type network-qos c-nq3 match qos-group 3	N9K	Leaf
	policy-map type network-qos Jumbo-nq-policy class type network-qos c-nq3 mtu 9216	N9K	Leaf
	class type network-qos c-nq3 match qos-group 3 mtu 9216 class type network-qos c-nq2 match qos-group 2 mtu 9216 class type network-qos c-nq1 match qos-group 1 mtu 9216 class type network-qos c-nq-default match qos-group 0 mtu 9216	N9K	Leaf
	system qos service-policy type network-qos Jumbo-nq-policy	N9K	Leaf
QoS Hardware resources configuration	"hardware access-list tcam region racl 0 hardware access-list tcam region e-racl 0 hardware access-list tcam region span 0 hardware access-list tcam region vqos 256 hardware access-list tcam region e-qos 256 hardware access-list tcam region arp-ether 256"	N9K	Leaf
N 9500 QoS	system qos service-policy type queuing output default-out-policy	N9K	Leaf

	system qos service-policy type network-qos Jumbo-nq-policy	N9K	Leaf
N 9500 QoS Queuing policy	policy-map type queuing default-out-policy class type queuing c-out-q3 priority level 1 class type queuing c-out-q2 bandwidth remaining percent 0 class type queuing c-out-q1 bandwidth remaining percent 0 class type queuing c-out-q-default bandwidth remaining percent 100	N9K	Leaf
	System qos Service-policy type queuing out default-out-policy	N9K	Leaf
N 9500 QoS Queuing Policy	policy-map type queuing default-out-policy class type queuing c-out-q3 priority level 1 class type queuing c-out-q2 bandwidth remaining percent 0 class type queuing c-out-q1 bandwidth remaining percent 0 class type queuing c-out-q-default bandwidth remaining percent 100	N9K	Leaf
	System qos Service-policy type queuing out default-out-policy	N9K	Leaf
Network Management Ethernet (Mgmt0)	interface mgmt0 ip address 10.10.10.10/24	N9K	Leaf
	vrf context management ip route 0.0.0.0/0 10.218.23.254	N9K	Leaf
Configuring Hostname on Nexus 9000	hostname nw_lf_cnx9_001.4lgebz_o01_s01	N9K	Leaf
Time Zone and day-light saving	clock timezone EET 2 0 clock summer-time EEST 4 Sunday March 02:00 4 Sunday October 03:00 60	N9K	Leaf
DNS	ip domain-name <cust_name> no ip domain-lookup	N9K	Leaf
SNMP	snmp-server contact <contact_name> snmp-server location <location_name>	N9K	Leaf
	snmp-server host 85.29.26.36 traps version 2c <SNMP_Community_1> snmp-server host 85.29.56.136 traps version 2c <SNMP_Community_1> snmp-server host 85.29.60.191 traps version 2c <SNMP_Community_1> snmp-server host 85.29.60.235 traps version 2c <SNMP_Community_1> snmp-server host 213.74.189.232 traps version 2c <SNMP_Community_1> snmp-server host 213.74.189.233 traps version 2c <SNMP_Community_1>	N9K	Leaf

	<pre>snmp-server host 85.29.26.36 use-vrf management snmp-server host 85.29.56.136 use-vrf management snmp-server host 85.29.60.191 use-vrf management snmp-server host 85.29.60.235 use-vrf management snmp-server host 213.74.189.232 use-vrf management snmp-server host 213.74.189.233 use-vrf management</pre>	N9K	Leaf
	<pre>snmp-server source-interface trap mgmt0</pre>	N9K	Leaf
	<pre>snmp-server community <community> group network-admin</pre>	N9K	Leaf
	<pre>15 permit ip host 213.74.197.43 any ... 390 permit ip host 176.43.250.25 any</pre>	N9K	Leaf
LLDP on Nexus 9000	<pre>feature lldp</pre>	N9K	Leaf
Network Security Disable IP Redirects	<pre>interface Ethernet slot#/port# no ip redirects no ipv6 redirects</pre>	N9K	Leaf
Device Access Security	<pre>NX-OS(config)#no ssh server enable NX-OS(config)#ssh key {dsa [force] rsa [bits [force]]} NX-OS(config)#ssh server enable NX-OS#show ssh key ***** rsa Keys generated:Fri Apr 10 20:13:21 2010 <clipped> !</pre>	N9K	Leaf
AAA-N	<pre>NX-OS(config)#feature tacacs+ NX-OS(config)#tacacs-server host {ipv4-address ipv6-address host-name} NX-OS(config)#tacacs-server key [0 7] key-value NX-OS(config)#aaa group server tacacs+ group-name server {ipv4-address ipv6-address host-name} deadtime minutes use-vrf <demo_name> NX-OS(config)#tacacs-server timeout seconds NX-OS(config)#tacacs-server host {ipv4-address ipv6-address host-name} port tcp-port NX-OS(config)#tacacs-server deadtime minutes</pre>	N9K	Leaf
	<pre>feature tacacs+ aaa group server tacacs+ TacacsGroup use-vrf management server 10.35.175.1 aaa authentication login console group TacacsGroup aaa authentication login default group TacacsGroup aaa authentication login error-enable ! tacacs-server host 10.35.175.1 key <shared-key> port 49 tacacs-server directed-request ip tacacs source-interface mgmt 0 ! ! Device Login Authorisation with AAA !</pre>	N9K	Leaf

	<pre> aaa authorization config-commands default group TacacsGroup local aaa authorization commands default group TacacsGroup local ! ! Device Login Accounting with AAA ! aaa accounting default group TacacsGroup ! ! Local User Configuration ! username admin Pword <Pword> role network-admin </pre>	N9K	Leaf
Device Hardening 3.9.9.4 COPP policy and class maps	<pre> policy-map type control-plane copp-system-p-policy-strict class copp-system-p-class-l3uc-data set cos 1 police cir 250 pps bc 32 packets conform transmit violate drop </pre>	N9K	Leaf
	<pre> class copp-system-p-class-critical set cos 7 police cir 19000 pps bc 128 packets conform transmit violate drop </pre>	N9K	Leaf
	<pre> class copp-system-p-class-important set cos 6 police cir 3000 pps bc 128 packets conform transmit violate drop </pre>	N9K	Leaf
	<pre> class copp-system-p-class-multicast-router set cos 6 police cir 3000 pps bc 128 packets conform transmit violate drop </pre>	N9K	Leaf
	<pre> class copp-system-p-class-management set cos 2 police cir 3000 pps bc 32 packets conform transmit violate drop </pre>	N9K	Leaf
	<pre> class copp-system-p-class-multicast-host set cos 1 police cir 2000 pps bc 128 packets conform transmit violate drop </pre>	N9K	Leaf
	<pre> class copp-system-p-class-l3mc-data set cos 1 police cir 3000 pps bc 32 packets conform transmit violate drop </pre>	N9K	Leaf
	<pre> class copp-system-p-class-normal set cos 1 police cir 1500 pps bc 32 packets conform transmit violate drop </pre>	N9K	Leaf
	<pre> class copp-system-p-class-ndp set cos 6 police cir 1500 pps bc 32 packets conform transmit violate drop </pre>	N9K	Leaf
	<pre> class copp-system-p-class-normal-dhcp set cos 1 police cir 300 pps bc 32 packets conform transmit violate drop </pre>	N9K	Leaf

	<pre>class copp-system-p-class-normal-dhcp-relay-response set cos 1 police cir 400 pps bc 64 packets conform transmit violate drop</pre>	N9K	Leaf
	<pre>class copp-system-p-class-normal-igmp set cos 3 police cir 6000 pps bc 64 packets conform transmit violate drop</pre>	N9K	Leaf
	<pre>class copp-system-p-class-redirect set cos 1 police cir 1500 pps bc 32 packets conform transmit violate drop</pre>	N9K	Leaf
	<pre>class copp-system-p-class-exception set cos 1 police cir 50 pps bc 32 packets conform transmit violate drop</pre>	N9K	Leaf
	<pre>class copp-system-p-class-exception-diag set cos 1 police cir 50 pps bc 32 packets conform transmit violate drop</pre>	N9K	Leaf
	<pre>class copp-system-p-class-monitoring set cos 1 police cir 300 pps bc 128 packets conform transmit violate drop</pre>	N9K	Leaf
	<pre>class copp-system-p-class-12-unpoliced set cos 7 police cir 20000 pps bc 8192 packets conform transmit violate drop</pre>	N9K	Leaf
	<pre>class copp-system-p-class-undesirable set cos 0 police cir 15 pps bc 32 packets conform transmit violate drop</pre>	N9K	Leaf
	<pre>class copp-system-p-class-fcoe set cos 6 police cir 1500 pps bc 128 packets conform transmit violate drop</pre>	N9K	Leaf
	<pre>class copp-system-p-class-nat-flow set cos 7 police cir 100 pps bc 64 packets conform transmit violate drop</pre>	N9K	Leaf
	<pre>class copp-system-p-class-12-default set cos 0 police cir 50 pps bc 32 packets conform transmit violate drop</pre>	N9K	Leaf
	<pre>class class-default set cos 0 police cir 50 pps bc 32 packets conform transmit violate drop</pre>	N9K	Leaf

	N9k-ST-Leaf-01# sh copp status Last Config Operation: None Last Config Operation Timestamp: None Last Config Operation Status: None Policy-map attached to the control-plane: copp-system-p-policy-strict	N9K	Leaf
	N9k-ST-Leaf-01# sh copp profile ? dense Display dense profile lenient Display lenient profile moderate Display moderate profile strict Display strict profile	N9K	Leaf
BFD	feature bfd bfd interval 50 min_rx 50 multiplier 3	N9K	Leaf
	router ospf UNDERLAY bfd	N9K	Leaf
	router bgp 65539 vrf <demo_name> address-family ipv4 unicast	N9K	Leaf
	router bgp 65539 vrf <demo_name> local-as 65539	N9K	Leaf
	router bgp 65539 vrf <demo_name> neighbor 10.23.65.0 remote-as 65541 bfd	N9K	Leaf
OSPF Routing Process	feature ospf ! router ospf UNDERLAY log-adjacency-changes detail bfd	N9K	Leaf
OSPF Router ID	router ospf UNDERLAY log-adjacency-changes detail bfd router-id <loopback17-ip-address>	N9K	Leaf
Enabling OSPF on interfaces	router ospf UNDERLAY passive-interface default	N9K	Leaf
	continue from the above... interface Ethernet1/5 ip router ospf UNDERLAY area 0.0.0.1 ip ospf bfd ip ospf network point-to-point no ip ospf passive-interface	N9K	Leaf
	interface loopback<id> ip router ospf UNDERLAY area 0.0.0.1	N9K	Leaf
OSPF Authentication	interface eth <slot>/<port> ip ospf authentication message-digest ip ospf message-digest-key <key-id> md5 0 <clear-text-key>	N9K	Leaf
OSPF Reference-Bandwidth	router ospf UNDERLAY auto-cost reference bandwidth 100Gbps	N9K	Leaf

Underlay OSPF Configuration on Leaf Underlay OSPF Configuration on Spine	interface loopback17 ip router ospf UNDERLAY area 0.0.0.1	N9K	Leaf
	interface eth<slot>/<port> ip router ospf UNDERLAY area 0.0.0.1 ip ospf network point-to-point no ip ospf passive-interface ip ospf bfd ip ospf authentication message-digest ip ospf message-digest-key <key-id> md5 0 <clear-text-key>	N9K	Leaf / Spine
Enabling Multicast Routing - PIM	feature pim	N9K	Leaf
	ip pim long-neighbor-changes	N9K	Spine
	interface ethernet 1/10 ip pim sparse-mode	N9K	Spine
	interface ethernet 1/10 ip pim bfd-instance	N9K	Spine
	interface loopback<id> ip pim sparse-mode	N9K	Leaf
Mapping Layer 2 VNI VXLAN segment to ASM group	interface nve<id> member vni <L2-VNID> mcast-group 239.239.0.1 member vni <L2-VNID> mcast-group 239.239.0.2	N9K	Leaf
PIM Anycast RP (RFC 4610)	interface loopback18 ip pim sparse-mode	N9K	Leaf
	interface loopback17 ip pim sparse-mode	N9K	Leaf
	ip pim rp-address <loopback18> group-list 239.239.0.0/16	N9K	Leaf
Multicast configuration for Leaf	ip pim rp-address <anycast-loopback> group-list 239.239.0.0/16	N9K	Leaf
	feature pim ip pim log-neighbor-changes	N9K	Leaf
	interface loopback17 ip pim sparse-mode	N9K	Leaf
	interface ethernet<slot>/<port> ip pim sparse-mode ip pim bfd-instance	N9K	Leaf
	interface nve1 member vni <L2-VNID> mcast-group 239.64.64.1 member vni <L2-VNID> mcast-group 239.64.64.2	N9K	Leaf

	<code>ip pim rp-address <anycast-loopback> group-list 239.239.0.0/16</code>	N9K	Leaf
Multicast configuration for Spine	<code>feature pim ip pim log-neighbor-changes</code>	N9K	Spine
	<code>interface ethernet 1/10 ip pim sparse-mode</code>	N9K	Spine
	<code>interface ethernet 1/10 ip pim bfd-instance</code>	N9K	Spine
	<code>interface loopback17 ip pim sparse-mode</code>	N9K	Spine
	<code>interface loopback18 ip pim sparse-mode</code>	N9K	Spine
	<code>ip pim rp-address <loopback18> group-list 239.239.0.0/16</code>	N9K	Spine
	<code>ip pim anycast-rp <loopback18> <loopback17></code>	N9K	Spine
Service Extensions for OSPF routing	<code>vlan 17 vn-segment 10019</code>	N9K	
	<code>interface Vlan17 mtu 9216 vrf member <demo_name> ip ospf cost 10 ip ospf passive-interface ip router ospf 1 area 0.0.0.0</code>	N9K	
Service Extensions for Static routing	<code>vrf context <demo_name> ip route 0.0.0.0/0 Vlan1605 11.0.23.30</code>	N9K	
Service Extension for default route injection on N9K BL/redistribute mode.	<code>router bgp 65542 vrf <demo_name> address-family ipv4 unicast network 0.0.0.0/0</code>	N9K	
route-map	<code>route-map RM-IN-S2 permit 10 match tag 1000 route-map RM-IN-S3 permit 10 match tag 1000</code>	N9K	
	<code>route-map RM-S-to-O permit 10 match tag 131 132 133 139 134 135 set metric-type type-1</code>	N9K	
vrf context <demo_name>	<code>vrf context <demo_name> ip route 9.59.207.0/24 Vlan1603 11.0.34.30 name <test_name> tag 1000 50</code>	N9K	
	<code>vrf context <demo_name> ip route 9.59.207.0/24 Ethernet1/46.2 11.0.40.142 name <test_name> tag 1000 10</code>	N9K	
	<code>vrf context <demo_name> ip route 10.0.0.0/12 Vlan1603 11.0.34.30 tag 1000 50</code>	N9K	

	<pre>vrf context <demo_name> ip route 10.0.0.0/12 Ethernet1/46.2 11.0.40.142 tag 1000 10</pre>	N9K	
	<pre>vrf context <demo_name> ip route 10.2.52.0/24 Vlan6 10.2.42.3 tag 1000</pre>	N9K	
	<pre>vrf context <demo_name> ip route 192.168.0.0/16 Vlan1603 11.0.34.30 name <test_name> tag 1000 rd auto address-family ipv4 unicast route-target both auto route-target both auto evpn</pre>	N9K	
vrf context <demo_name>	<pre>vrf context <demo_name> ip route 10.2.0.0/19 Vlan1607 11.0.34.14 tag 131 50</pre>	N9K	
	<pre>vrf context <demo_name> ip route 10.2.0.0/19 Ethernet1/45.1 11.0.40.145 tag 131 10</pre>	N9K	
	<pre>vrf context <demo_name> ip route 10.2.96.0/19 Vlan3203 11.0.39.14 tag 134</pre>	N9K	
interface Vlan1601	<pre>interface Vlan1601 no shutdown vrf member <demo_name> no ip redirects ip address 10.10.10.10/24 no ipv6 redirects hsrp version 2 hsrp 1601 preempt priority 110 ip 11.0.34.33</pre>	N9K	
interface Vlanxx	<pre>interface Vlan1602 no shutdown vrf member <demo_name> no ip redirects no ipv6 redirects ip ospf cost 10 ip ospf passive-interface ip router ospf 100 area 0.0.0.0</pre>	N9K	
interface Ethernet (IPv4 and IPv6)	<pre>interface ex/y mac aaa.bbbb.cccc vrf member <demo_name> ip address x.x.x.x/31 ipv6 address x:x:x::x ip policy route-map TO_VPER_OR_FW ipv6 policy route-map TO_VPER_OR_FW_v6 no shut</pre>	N7K	
interface Ethernet1/46.1	<pre>interface Ethernet1/36.1 mtu 1500</pre>	N9K	

	<pre>interface Ethernet1/36.1 encapsulation dot1q 1602 mac-address 0000.0000.2222 vrf member <demo_name> no ip redirects ip address 10.10.10.10/24</pre>	N9K	
interface Ethernet1/47.1	<pre>interface Ethernet1/37.1 mtu 1500 encapsulation dot1q 1608 vrf member <demo_name> no ip redirects ip address 10.10.10.10/24 ip ospf dead-interval 20 ip ospf hello-interval 5 ip ospf network point-to-point ip router ospf 100 area 0.0.0.0</pre>	N9K	
router ospf 1	<pre>router ospf 1 vrf <demo_name> router-id 55.2.32.5 vrf <demo_name> router-id 55.2.32.5 vrf <demo_name> router-id 55.2.32.5 redistribute static route-map RM-S-to-O</pre>	N9K	
router bgp 65543	<pre>router bgp 65543 vrf <demo_name> address-family ipv4 unicast advertise l2vpn evpn redistribute direct route-map vts-subnet-policy redistribute static route-map RM-IN-S2</pre>	N9K	
	nv overlay evpn	N9K	
	clock protocol ntp vdc 1	N9K	
role name nsdcheck	<pre>role name nsdcheck rule 4 permit command show * rule 3 permit command terminal length * rule 2 permit command ping * rule 1 permit read</pre>	N9K	
	<pre>role name devcheck rule 8 permit command tac-pac * rule 7 permit command dir * rule 6 permit command ssh * rule 5 permit command traceroute * rule 4 permit command ping *</pre>	N9K	
	<pre>role name devopera rule 1 permit read-write</pre>	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	<pre>username user password 5 \$1\$lDuqR.60\$eNzZ5I22WxJT58gdEm88N0 role network-operator</pre>	N9K	

	username vtsadmin password 5 \$5\$MmpswImI\$vbZhp/52dNjHY5KWj4yBvmiDvuOZZ9gd2vo2oZc61b4 role network-admin	N9K	
	username nsdcheck password 5 \$5\$dpIXMjZs\$jDIZVf6grMulyq79vTts2mcgPlt0QWp5z3tDnw3N5W8 role nsdcheck	N9K	
snmp-server	snmp-server source-interface trap loopback1	N9K	
	snmp-server user user network-operator auth md5 0x3eaa4221f6bbf8722cbdea7ea6bf2f11 priv 0x3eaa4221f6bbf8722cbdea7ea6bf2f11 localizedkey	N9K	
	snmp-server host 55.6.8.1 traps version 2c COMMUNITY1 snmp-server host 55.6.8.1 use-vrf default	N9K	
	snmp-server enable traps bgp snmp-server enable traps ospf snmp-server enable traps callhome event-notify snmp-server enable traps callhome smtp-send-fail snmp-server enable traps cfs state-change-notif snmp-server enable traps lldp lldpRemTablesChange snmp-server enable traps aaa server-state-change snmp-server enable traps hsrp state-change snmp-server enable traps feature-control FeatureOpStatusChange snmp-server enable traps sysmgr cseFailSwCoreNotifyExtended snmp-server enable traps config ccmCLIRunningConfigChanged snmp-server enable traps snmp authentication snmp-server enable traps link cisco-xcvr-mon-status-chg snmp-server enable traps vtp notifs snmp-server enable traps vtp vlancreate snmp-server enable traps vtp vlandelete snmp-server enable traps bridge newroot snmp-server enable traps bridge topologychange snmp-server enable traps stpx inconsistency snmp-server enable traps stpx root-inconsistency snmp-server enable traps system Clock-change-notification snmp-server enable traps feature-control ciscoFeatOpStatusChange	N9K	
	snmp-server community COMMUNITY1 group network-operator	N9K	
ntp	ntp source-interface loopback0 ntp logging	N9K	
ip pim	ip pim ssm range 232.0.0.0/8	N9K	
spanning-tree	spanning-tree pathcost method long spanning-tree mst 1 priority 4096 spanning-tree mst configuration name CFG01 revision 1 instance 1 vlan 1-4094	N9K	
hardware	hardware access-list tcam region qos 0	N9K	

vpc domain	vpc domain 151 peer-keepalive destination 55.2.34.2 source 55.2.34.1 vrf default	N9K	
	vpc domain 151 auto-recovery	N9K	
interface vlan	interface Vlan1602 no shutdown vrf member <demo_name> no ip redirects fabric forwarding mode anycast-gateway	N9K	
interface port-channel	interface port-channel101 no switchport mtu 9216 no ip redirects ip address 10.10.10.10/24 ip ospf cost 10 ip ospf dead-interval 20 ip ospf hello-interval 5 ip ospf network point-to-point ip router ospf 1 area 0.0.0.0 ip pim sparse-mode	N9K	
interface Ethernet	interface Ethernet1/45 no switchport mtu 9216 mac-address 0000.0000.1111	N9K	
	interface Ethernet1/47 no switchport mtu 9216 udld enable	N9K	
	interface Ethernet2/5 switchport mode trunk switchport trunk allowed vlan 2-4094 channel-group 21 mode active	N9K	
interface mgmt0	interface mgmt0 no lldp transmit no lldp receive	N9K	
clock timezone	clock timezone PRC 8 0	N9K	
ip route	ip route 0.0.0.0/0 Ethernet1/46.452 55.6.34.198 tag 1000 10 ip route 0.0.0.0/0 Vlan3903 55.6.40.14 tag 1000 50	N9K	
router ospf	router ospf 1 redistribute static route-map RM-S-to-0	N9K	
router bgp	router bgp 65543 router-id 55.2.32.5 address-family ipv4 unicast address-family l2vpn evpn neighbor 55.2.32.1 remote-as 65543 update-source loopback1 address-family ipv4 unicast address-family l2vpn evpn send-community extended	N9K	

<p>router bgp (IPv4 only)</p>	<pre>router bgp 65539 router-id 192.168.0.25 log-neighbor-changes address-family ipv4 unicast maximum-paths 32 maximum-paths ibgp 32 address-family ipv6 unicast maximum-paths 32 maximum-paths ibgp 32 address-family l2vpn evpn neighbor 192.168.0.3 remote-as 65539 password 3 2b7cf4643b66b222 update-source loopback17 address-family l2vpn evpn send-community send-community extended</pre>	<p>N9K and N7K</p>	
<p>Event manager config (IPv4 and IPv6)</p>	<pre>event manager applet TRACK-PING-FOR-BGP-DOWN event track 1 state down action 1.0 syslog msg CANNOT PING FW. GOING TO SHUTDOWN BGP PEER action 2.0 cli config term action 3.0 cli router bgp 65539 action 4.0 cli vrf <demo_name> action 5.0 cli neighbor 175.175.175.175 action 6.0 cli shutdown event manager applet TRACK-PING-FOR-BGP-UP event track 1 state up action 1.0 syslog msg CAN PING FW. GOING TO NO SHUTDOWN BGP PEER action 2.0 cli config term action 3.0 cli router bgp 65539 action 4.0 cli vrf <demo_name> action 5.0 cli neighbor 175.175.175.175 action 6.0 cli no shutdown</pre>	<p>N9K</p>	
<p>IP sla config (IPv4 only for N9K) (IPv4 and IPv6 for N7K)</p>	<pre>!On BL-1 Track the local VPER-1 ip sla 1 icmp-echo 69.83.32.36 source-interface vlan 2400 vrf <demo_name> □ forward reference to VRF threshold 500 timeout 500 frequency 1 ! Start the SLAs ip sla schedule 1 life forever start-time now ! Setup a track object for sla 1 track 1 ip sla 1 reachability delay up 180 down 3 ! Set up a track open that returns a DOWN only if both objects 1 and 2 are down. track 111 list boolean or object 1</pre>	<p>N9K and N7K</p>	
<p>Track config (IPv4 and IPv6)</p>	<pre>track 10 ip route 0.0.0.0/0 reachability vrf member <demo_name></pre>	<p>N9K and N7K</p>	

Interface port channel (IPv4 and IPv6)	<pre>interface port-channel 110.2511 encapsulation dot1q 2511 vrf member <demo_name> ip address 10.10.10.10/24 no shut interface port-channel 110.2575 encapsulation dot1q 2575 vrf member <demo_name> ipv6 address 10:10:10:10:10:10:10:10/64</pre>	N9K and N7K	
	<pre>interface port-channel 110.2577 ip policy route-map FROM_VPER interface port-channel 110.2577 ipv6 policy route-map FROM_VPERv6 ! EEM to track both VPERs, when one is up restore traffic event manager applet VPER_TRACK_UP event track 111 state up action 1.0 syslog msg "BOTH VPERS ARE UP. REMOVING BYPASS!" action 2.0 cli command "config t" action 3.0 cli command "route-map TO_VPER_OR_FW permit 20" action 4.0 cli command "no continue 30" action 5.0 cli command "exit" action 6.0 cli command "route-map TO_VPER_OR_FWv6 permit 20" action 7.0 cli command "no continue 30" action 8.0 cli command "exit" action 9.0 cli command "route-map FROM_FW_TO_VPER_OR_MOBILE permit 10" action 10.0 cli command "no continue 20" action 11.0 cli command "end" action 12.0 cli command "route-map FROM_FW_TO_VPER_OR_MOBILEv6 permit 10" action 13.0 cli command "no continue 20" action 14.0 cli command "end" action 15.0 syslog msg "TRAFFIC HAS BEEN RESTORED TO VPER"</pre>	N7K	

	<pre>! EEM to track both VPERs, when one is up restore traffic event manager applet VPER_TRACK_UP event track 111 state up action 1.0 syslog msg BOTH VPERs ARE UP. REMOVING BYPASS action 2.0 cli command "config t" action 3.0 cli command "route-map TO_VPER_OR_FW permit 20" action 4.0 cli command "no continue 30" action 5.0 cli command "exit" action 6.0 cli command "route-map TO_VPER_OR_FWv6 permit 20" action 7.0 cli command "no continue 30" action 8.0 cli command "exit" action 9.0 cli command "route-map FROM_FW_TO_VPER_OR_MOBILE permit 10" action 10.0 cli command "no continue 20" action 11.0 cli command "end" action 12.0 cli command "route-map FROM_FW_TO_VPER_OR_MOBILEv6 permit 10" action 13.0 cli command "no continue 20" action 14.0 cli command "end" action 15.0 syslog msg TRAFFIC HAS BEEN RESTORED TO VPER</pre>	N9K	
IP access list (IPv4 and IPv6)	<pre>ip access-list ALL_POOLS 10 permit ip 1.0.0.0/8 any 20 permit ip any 1.0.0.0/8 30 permit ip 2.0.0.0/8 any 40 permit ip any 2.0.0.0/8 ! Need to configure a ACL for all All POOLS ipv6 access-list ALL_POOLSv6 10 permit ipv6 2001:1::/32 any 20 permit ipv6 any 2001:1::/32 30 permit ipv6 2001:2::/32 any 40 permit ipv6 any 2001:2::/32</pre>	N9K and N7K	
Route-map (IPv4 and IPv6)	<pre>set ip next-hop verify-availability 69.83.32.35 track 2 route-map TO_VPER_OR_FW permit 30 match ip address ALL_POOLS ! Set the ip next-hop to the FW VIP set ip next-hop 69.83.136.129 route-map TO_VPER_OR_FW_v6 permit 10 ! Leave room here for the pilot packets route-map TO_VPER_OR_FW_v6 permit 20 match ipv6 address VPER_POOLSv6 set ipv6 next-hop verify-availability 2001:4888:16:2078:1e1:210:: track 1 set ipv6 next-hop verify-availability 2001:4888:16:207a:1e1:210:: track 2 route-map TO_VPER_OR_FW_v6 permit 30 match ipv6 address ALL_POOLSv6 ! Set the ipv6 next-hop to the FW VIP set ipv6 next-hop 2001:4888:39:3080:308:25::</pre>	N9K and N7K	
	<pre>route-map FROM_FW_TO_VPER_OR_MOBILE permit 10 match ip address VPER_POOLS set vrf <demo_name>_VPER</pre>	N7K	

Monitor erspan	<pre>monitor session 1 type erspan-source erspan-id 5 vrf <demo_name> ip ttl 25 ip dscp 42 monitor erspan origin ip-address 10.0.0.1 global</pre>	N9K and N7K	
QOS- class-map	<pre>class-map type qos match-any TEST1 match packet length 5</pre>	N9K and N7K	
QOS class-map policy-map	<pre>class-map type control-plane match-any cust1-copp-system-p-class-exception match exception ip option match exception ip icmp unreachable match exception ipv6 option match exception ipv6 icmp unreachable class-map type control-plane match-any cust1-copp-system-p-class-fcoe match access-group name cust1-copp-system-p-acl-mac-fcoe policy-map type control-plane cust1-copp-system-p-policy-strict class cust1-copp-system-p-class-exception set cos 1 police cir 360 kbps bc 250 ms conform transmit violate drop class cust1-copp-system-p-class-fcoe set cos 6 police cir 1060 kbps bc 1000 ms conform transmit violate drop</pre>		
Tunnel Interface	<pre>interface Tunnel1 vrf member <demo_name> ip address 10.10.10.10/24 tunnel source 1.1.1.201 tunnel destination 1.1.1.200 no shutdown</pre>	N9K and N7K	



APPENDIX **C**

OpenStack Configuration for SR-IOV Support

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

On Controller Node

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

```
[ml2]

tenant_network_types = vxlan, vlan

# Add support for vlan, vxlan and flat type drivers

type_drivers = vlan, vxlan, flat

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

mechanism_drivers = sriovnicswitch, cisco_vts
```

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

```
[ml2_type_vlan]

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

# reserved for the provider networks for those physnets

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

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

ml2_type_flat]

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

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

flat_networks = *
```

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

```
[ml2_sriov]

# The default supported_pci_vendor_devs value for the installation may have the value
```

```

only for the PFs

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

# as well in this.

supported_pci_vendor_devs = 8086:10fb,8086:10ed

agent_required = True

```

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

```

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

```

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

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

```

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

```

For multiple SR-IOV nics, there should be a mapping entry per `physnet/NIC` card

Enable the OpenStack Networking SR-IOV agent

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

Then proceed to create the SR-IOV port instances.



Note

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

- [Sample for SR-IOV Trunk \(No-Bonding\), on page 205](#)
- [Sample for SR-IOV Trunk \(Bonding\), on page 205](#)

Sample for SR-IOV Trunk (No-Bonding)

This section provides an example for enabling SR-IOV trunk (no-bonding).

Step 1 Create a Flat network.

```
openstack network create --provider-network-type flat --provider-physical-network SRIOV-B vma-flat-net
flatneta=$(openstack network list -f value -c ID -c Name | grep vma-flat-net | awk '{print $1}')
```

Step 2 Create a VLAN network.

```
openstack network create --provider-network-type vlan --provider-physical-network SRIOV-A
--provider-segment xxx vma-vlanxxx-net
vlanneta=$(openstack network list -f value -c ID -c Name | grep vma-vlanxxx-net | awk '{print $1}')
```

Step 3 Create subnet on the Flat and VLAN network.

```
openstack subnet create --network $flatneta --subnet-range 1.1.1.0/24 --gateway 1.1.1.1 --no-dhcp
vma-flat-subnet
openstack subnet create --network $vlanneta --subnet-range 2.1.1.0/24 --gateway 2.1.1.1 --no-dhcp
vma-vlanxxx-subnet
```

Step 4 Create port on Flat and VLAN network.

```
openstack port create --vnic-type direct --network $ flatneta flatneta-port
openstack port create --vnic-type direct --network $vlanneta vma-child-port-vlanxxx-a
```

Step 5 Create trunk with parent (Flat) and sub-port (VLAN).

```
openstack network trunk create --parent-port flatneta-port \
--subport port=vma-child-port-vlanxxx-a, segmentation-type=vlan, segmentation-id=xxx vma-trunk-a
```

Step 6 Attach the parent port to Nova instance.

```
parentaid=$( openstack port list -f value -c ID -c Name | grep flatneta-port | awk '{print $1}')
nova boot --flavor m1.medium --image cents7-1800-custom --nic port-id=$parentaid vma
```

Sample for SR-IOV Trunk (Bonding)

This section provides an example for enabling SR-IOV trunk (bonding).

Step 1 Create Flat networks for each SR-IOV physnet, but use the same network name for both the networks. Keep the subnets identical.

Note It is important to keep the network and subnet names exactly the same.

a. Create a Flat network on the first physnet.

```
openstack network create --provider-network-type flat --provider-physical-network SRIOV-A
vma-flat-net
```

b. Find and save the ID of the network created.

```
flatneta=$(openstack network list -f value -c ID -c Name | grep vma-flat-net | awk '{print $1}')
```

- c. Create a Flat network on the second physnet.

```
openstack network create --provider-network-type flat --provider-physical-network SRIOV-B
vma-flat-net
```

- d. Find and save the ID of the network created, there should be two matching, so exclude the \$flatneta network.

```
flatnetb=$(openstack network list -f value -c ID -c Name | grep vma-flat-net | grep -v $flatneta
| awk '{print $1}')
```

- e. Create subnet on the Flat networks.

```
openstack subnet create --network $flatneta --subnet-range 1.1.1.0/24 --gateway 1.1.1.1 --no-dhcp
vma-flat-subnet
openstack subnet create --network $flatnetb --subnet-range 1.1.1.0/24 --gateway 1.1.1.1 --no-dhcp
vma-flat-subnet
```

Step 2 Create VLAN Networks for each VLAN tag that should be allowed in this trunk.

Create a pair of networks, one on each physnet for each VLAN. Replace *xxx* with the VLAN ID to be used in the following examples.

Note It is important to keep the network name and subnet name exactly the same for the pair created for a specific VLAN.

- a. Create a VLAN network on the first physnet.

```
openstack network create --provider-network-type vlan --provider-physical-network SRIOV-A
--provider-segment xxx vma-vlanxxx-net
vlanneta=$(openstack network list -f value -c ID -c Name | grep vma-vlanxxx-net | awk '{print
$1}')
```

- b. Create a VLAN network on the second physnet.

```
openstack network create --provider-network-type vlan --provider-physical-network SRIOV-B
--provider-segment xxx vma-vlanxxx-net
vlannetb=$(openstack network list -f value -c ID -c Name | grep vma-vlanxxx-net | grep -v $vlanneta
| awk '{print $1}')
```

- c. Create subnet on VLAN networks.

```
openstack subnet create --network $vlanneta --subnet-range 2.1.1.0/24 --gateway 2.1.1.1 --no-dhcp
vma-vlanxxx-subnet
openstack subnet create --network $vlannetb --subnet-range 2.1.1.0/24 --gateway 2.1.1.1 --no-dhcp
vma-vlanxxx-subnet
```

Step 3 Create ports on the Flat as well as VLAN networks created above.

```
openstack port create --vnic-type direct --network $flatneta vma-parent-port-a
openstack port create --vnic-type direct --network $flatnetb vma-parent-port-b
```

- a. Create as many child ports as VLANs needed.

```
openstack port create --vnic-type direct --network $vlanneta vma-child-port-vlanxxx-a
openstack port create --vnic-type direct --network $vlannetb vma-child-port-vlanxxx-b
```

Step 4 Create two trunks, one per physnet. The parent port is the Flat network port, and all the VLAN network ports are the child ports.

```
openstack network trunk create --parent-port vma-parent-port-a \
--subport port=vma-child-port-vlanxxx-a, segmentation-type=vlan, segmentation-id=xxx vma-trunk-a
openstack network trunk create --parent-port vma-parent-port-b \
--subport port=vma-child-port-vlanxxx-b, segmentation-type=vlan, segmentation-id=xxx vma-trunk-b
```


Step 5 Create an instance with NICs mapping to the trunk parent ports.

```
parentaid=$( openstack port list -f value -c ID -c Name | grep vma-parent-port-a | awk '{print $1}')  
parentbid=$( openstack port list -f value -c ID -c Name | grep vma-parent-port-b | awk '{print $1}')
```

```
nova boot --flavor m1.medium --image cents7-1801-custom --nic port-id=$parentaid --nic  
port-id=$parentbid vma
```

Step 6 Attach networks/subnets to routers.

Create a router and attach the interfaces of the VLAN networks—Only one of the VLAN networks in the pair for each VLAN should be attached. The Flat network subnet (one in the pair) may need to be attached if there are untagged packets from the VM/instance to be routed as well.

Step 7 Configure instance interfaces with bonded NICs in active/standby.



APPENDIX **D**

collectd Plugin Configuration for VTC and VTF

This appendix provides collectd plugin configuration for VTC and VTF.

- [collectd Plugin Configuration, on page 209](#)

collectd Plugin Configuration

The following sections provide the default collectd plugin configuration.

Policy Plane (VTC) Plugin Configuration

1. CPU

```
LoadPlugin cpu
<Plugin cpu>
ReportByCpu true
ReportByState true
ValuesPercentage false
ReportNumCpu false
ReportGuestState false
SubtractGuestState true
</Plugin>
```

2. python

```
LoadPlugin python
<Plugin python>
ModulePath "/opt/vts/lib/python/"
LogTraces true
Import "vtsCollectD.vts_collectd_plugin"
</Plugin>
```

3. write_log

```
LoadPlugin write_log
<Plugin write_log>
Format JSON
</Plugin>
```

4. logfile

```
LoadPlugin logfile
<Plugin logfile>
LogLevel info
File "/opt/vts/log/collectd/metrics.log"
```

```

Timestamp true
PrintSeverity true
</Plugin>

```

5. interface

```

LoadPlugin interface

<Plugin interface>
Interface "eth0"
Interface "eth1"
Interface "lo"
IgnoreSelected false
ReportInactive true
UniqueName false
</Plugin>

```

6. disk

```

LoadPlugin disk

```

7. memory

```

LoadPlugin memory
<Plugin memory>
ValuesAbsolute true
ValuesPercentage false
</Plugin>

```

8. load

```

LoadPlugin load
<Plugin load>
ReportRelative true
</Plugin>

```

Data Plane (VTF)—Plugin Configs

1. python

```

LoadPlugin python
<Plugin python>
ModulePath "/opt/cisco/vpe/collectd/"
LogTraces true
Interactive false
Import "cisco-vpfa-collectd-plugin"
</Plugin>

```

2. interface

```

LoadPlugin interface
<Plugin interface>
Interface "br-ctlplane"
Interface "br-ex"
Interface "lo"
Interface "br-tenant"
IgnoreSelected false
ReportInactive true
UniqueName false
</Plugin>

```

3. disk

```

LoadPlugin disk

```

4. load

```
LoadPlugin load
<Plugin load>
  ReportRelative true
</Plugin>
```

5. memory

```
LoadPlugin memory
<Plugin memory>
ValuesAbsolute true
ValuesPercentage false
</Plugin>
```

6. cpu

```
LoadPlugin cpu
<Plugin cpu>
ReportByCpu true
ReportByState true
ValuesPercentage false
ReportNumCpu false
ReportGuestState false
SubtractGuestState true
</Plugin>
```

7. write_log

```
LoadPlugin write_log
<Plugin write_log>
Format JSON
</Plugin>
```

Write_Http Plugin Format

```
<LoadPlugin write_http>
  FlushInterval 60
</LoadPlugin>
<Plugin write_http>
  <Node "example">
    URL "http://10.2.1.1:8100/configure/collectd"
    Format "JSON"
    BufferSize 10240
  </Node>
</Plugin>
```




APPENDIX **E**

collectd Output JSON Examples

This appendix provides examples collectd JSON output.

VTC has the following information to emit to the Centralized Collect-D :

1. Master or Slave
2. IP or Hostname of the VTC
3. Stats Category—say number of tenants
4. Stats Sub-Category—say the vtep name if we have tenants per vtep
5. Count—count of tenants

Collect-D offers minimum attributes to set. They are:

1. Host
2. Plugin
3. Plugin-Instance
4. Type
5. Type-Instance

Of this the default values of Type can be only gauge and counter. To have custom types, the collect-d types.db needs to be updated with the custom types.

JSON Format

```
{
  "values": [1],
  "dstypes": ["counter"],
  "dsnames": ["value"],
  "time": 1515406938.687,
  "interval": 10.000,
  "host": "vtc-master-192.168.133.126",
  "plugin": "tenants",
  "plugin_instance": "",
  "type": "vtep",
  "type_instance": "vtep1"
}
```

- [Default Plugins—JSON Examples, on page 214](#)
- [Custom Plugin—JSON Examples, on page 214](#)

Default Plugins—JSON Examples

CPU

```
{
  "values": [1270783],
  "dstypes": ["derive"],
  "dsnames": ["value"],
  "time": 1515406948.326,
  "interval": 10.000,
  "host": "vtc126",
  "plugin": "cpu",
  "plugin_instance": "0",
  "type": "cpu",
  "type_instance": "user"
}
```

Memory

```
{
  "values": [335605760],
  "dstypes": ["gauge"],
  "dsnames": ["value"],
  "time": 1515406938.400,
  "interval": 10.000,
  "host": "vtc126",
  "plugin": "memory",
  "plugin_instance": "",
  "type": "memory",
  "type_instance": "slab_recl"
}
```

Custom Plugin—JSON Examples

Total Number of Tenants

```
{
  "values": [ 3],
  "dstypes": ["counter"],
  "dsnames": ["value"],
  "time": 1515406938.687,
  "interval": 10.000,
  "host": "vtc126",
  "plugin": "vtc",
  "plugin_instance": "",
  "type": "tenants",
  "type_instance": "vtc"
}
```

Total Number of Tenant per VTEP

```
{
  "values": [1],
  "dstypes": ["counter"],
  "dsnames": ["value"],
  "time": 1515406938.687,
  "interval": 10.000,
  "host": "vtc126",
  "plugin": "vtc",
  "plugin_instance": "",
  "type": "tenants",
}
```



```
    "type_instance": "vtep1"  
  }
```

