

Cisco SD-WAN Cloud OnRamp for Colocation Solution Guide, Release 20.5.1

First Published: 2018-12-20

Last Modified: 2021-04-22

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

© 2019–2021 Cisco Systems, Inc. All rights reserved.

Full Cisco Trademarks with Software License

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 NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE.

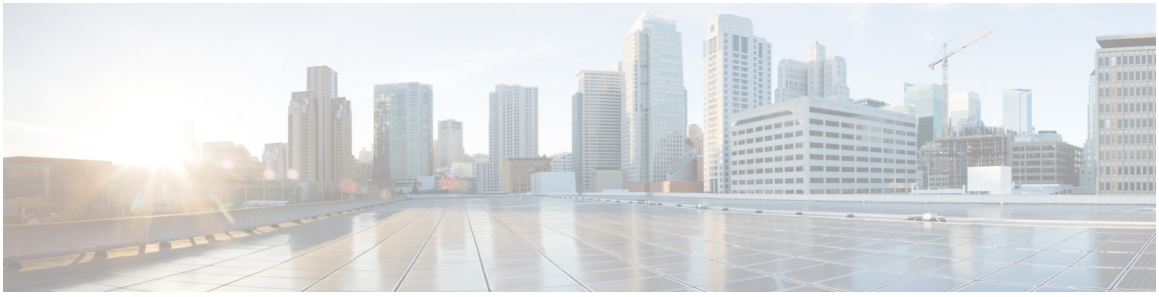
IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THIS MANUAL, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY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



CONTENTS

PREFACE

Preface ix

Audience ix

Related Documentation ix

List of Acronyms and Abbreviations x

Communications, Services, and Additional Information x

CHAPTER 1

Information About Cisco SD-WAN Cloud onRamp for Colocation Solution 1

Cisco SD-WAN Cloud onRamp for Colocation Solution 1

Cisco SD-WAN Cloud onRamp for Colocation Solution Components 2

CHAPTER 2

Prerequisites and Requirements of Cisco SD-WAN Cloud onRamp for Colocation Solution 7

Cisco SD-WAN Cloud onRamp for Colocation Solution Requirements 7

Hardware Requirements 7

Software Requirements 9

Wiring Requirements 10

Prescriptive Connections 10

Flexible Connections 11

Prerequisites for Deploying Solution 13

Ordering and Sizing of Cisco SD-WAN Cloud onRamp for Colocation Solution Devices 14

CHAPTER 3

Get Started with Cisco SD-WAN Cloud onRamp for Colocation Solution 15

Cisco SD-WAN Cloud onRamp for Colocation Solution–Deployment Workflow 15

Install Cisco NFVIS Cloud OnRamp for Colocation on Cisco CSP 16

Log Into CIMC User Interface 16

Activate Virtual Device 18

Map NFVIS Cloud OnRamp for Colocation Image 19
Bring up Cisco Cloud Services Platform Devices 19
Onboard CSP Devices Using Plug-and-Play Process 19
Onboard CSP Devices Using USB Bootstrapping Process 20
Verify Onboarded Devices and Activate Devices 22
Bring up Switch Devices 23
Bring up Cisco Colo Manager 25
Provision and Configure Cisco SD-WAN Cloud onRamp for Colocation Solution 26
Provision DHCP Server Per Colocation 26
Device Port Connectivity Details and Service Chaining for Prescriptive Connections 27
Validated Service Chains 31
Validated VM Packages 32
Customized Service Chains 33

CHAPTER 4 **Configure Cisco SD-WAN Cloud onRamp for Colocation Solution Devices from vManage** 35
Add Cloud OnRamp Colocation Devices into vManage 35
Delete Cloud OnRamp for Colocation Devices from vManage 35
Manage Clusters in vManage 36
Provision and Configure Cluster 37
Create and Activate Clusters 38
Cluster Configuration 40
Login Credentials 40
Resource Pool 41
Port Connectivity 42
NTP 43
Syslog Server 44
TACACS Authentication 44
Backup Server Settings 46
Progress of Cluster Activation 48
View Cluster from vManage 51
Edit Cluster in vManage 51
Add CSP Device to Cluster 52
Delete CSP Devices from Cluster 54
Delete CSP with CCM 55

Replace Cisco CSP Devices After RMA 56
Return of Materials of Cisco CSP Devices 57
RMA Process for Cisco CSP Devices 57
Prerequisites and Restrictions for Backup and Restore of CSP Devices 58
Remove PNF Devices from Cluster 60
Remove Cluster from vManage 60
Remove and Replace Switch 61
Reactivate Cluster from vManage 63
Manage Service Groups 64
VNF Placement for Service Chains in vManage 64
Create Service Chain in a Service Group 64
QoS on Service Chains 70
Clone Service Groups 71
Create Custom Service Chain 73
Physical Network Function Workflow 74
Custom Service Chain with Shared PNF Devices 75
Configure PNF and Cisco Catalyst 9500 Switches 78
Custom Service Chain with Shared VNF Devices 78
Shared VNF Use Cases 80
View Service Groups in vManage 86
Edit Service Group in vManage 86
Attach and Detach Service Group with Cluster 87
Day-N Configuration Workflow of Cisco SD-WAN Cloud onRamp for Colocation Solution 88

CHAPTER 5 **Software Image Management for Cluster Components and SWIM 91**
Manage VM Catalog and Repository 91
VNF Image Format 92
Upload VNF Images in Cisco vManage Repository 92
Create Customized VNF Image 93
View VNF Images in vManage Repository 98
Delete VNF Images from vManage Repository 98
Upgrade NFVIS Software Through vManage 99
Upload NFVIS Upgrade Image 99
Upgrade CSP Device with NFVIS Upgrade Image 100

Upgrade Cisco Catalyst 9500 Switches 101
Supported Upgrade Scenarios and Recommended Connections 104

CHAPTER 6 **Monitor Cisco SD-WAN Cloud onRamp for Colocation Solution Devices** 107
Monitor Operational Status of Cloud OnRamp for Colocation Devices from vManage 107
View Information About VNFs from vManage 109
View Cisco Colo Manager Health from vManage 111
Monitor Cloud OnRamp Colocation Clusters from Cisco vManage 113
Cisco Colo Manager States for Switch Configuration 117
Cisco Colo Manager States and Transitions from Host 118
Cisco Colo Manager Notifications 118
VM Alarms 121
VM States 123
Cloud Services Platform Real-Time Commands 123

CHAPTER 7 **High Availability** 125
Redundancy 125
Redundancy of Network Fabric 126
Redundancy of x86 Compute Hardware 126
Redundancy of Physical NIC or Interface 126
Redundancy of NFVIS, Virtualization Infrastructure 126
Redundancy of Service Chain or VNF 126
Recovery of Cisco Colo Manager 127
Handle Various Failure Scenarios 128

CHAPTER 8 **Cisco SD-WAN Cloud onRamp for Colocation Multitenancy** 129
Overview of Colocation Multitenancy 129
Roles and Functionalities in Multitenant Environment 130
Recommended Specifications in Multitenant Environment 131
Assumptions and Restrictions in Colocation Multitenancy 132
Service Provider Functionalities 133
Provision New Tenants 133
Create Colocation Group 133
View Permissions of a User Group 134

Create an RBAC User and Associate to Colocation Group 134
Delete RBAC User From Colocation User Group 135
Delete Tenants 135
Manage Tenant Colocation Clusters 136
Tenant Functionalities 137
Manage Colocation Clusters as Tenants 137
Monitor in Comanaged Multitenant Environment 138

CHAPTER 9 **Troubleshoot Cisco SD-WAN Cloud onRamp for Colocation Solution** 139
Troubleshoot Colocation Multitenancy Issues 139
Troubleshoot Catalyst 9500 Issues 140
Troubleshoot Cisco Cloud Services Platform Issues 145
DHCP IP Address Assignment 152
Troubleshoot Cisco Colo Manager Issues 153
Troubleshoot Service Chain Issues 155
Troubleshoot Physical Network Function Management Issues 157
Log Collection from CSP 157
Troubleshoot vManage Issues 157

CHAPTER 10 **Custom Packaging Details for Shared VNF** 159
Cisco vEdge Router Variable List 159
Cisco CSR1000V Variable List 163
ASAv Variable List 167



Preface

This guide provides information about how to configure and deploy Cisco SD-WAN Cloud onRamp for Colocation solution on a supported Cisco hardware device. The guide also provides details on virtual machine deployments, configuration of software features.

This guide assumes that readers have a broad understanding of networking terminologies and principles. It also assumes prior exposure to current trends in multi-cloud offerings.

- Audience, on page ix
- Related Documentation, on page ix
- List of Acronyms and Abbreviations, on page x
- Communications, Services, and Additional Information, on page x

Audience

This guide is intended for network administrators and operators who are familiar with basic Linux installation and configuration requirements.

Related Documentation

- Configuration Guide for Cisco Enterprise Network Function Virtualization Infrastructure Software
- Release Notes for Cisco Enterprise Network Function Virtualization Infrastructure Software
- Configuration Guide for Cisco Catalyst 9500 Switches
- Cisco Cloud Services Platform 5000 Hardware Installation Guide
- Configuration Guide for Cisco Network Plug and Play on Cisco APIC-EM

List of Acronyms and Abbreviations

Table 1: List of Acronyms and Abbreviations

Acronym or Abbreviation	Expansion
CSP	Cisco Cloud Services Platform
CIMC	Cisco Integrated Management Controller
Cisco SD-WAN Cloud onRamp for Colocation	Name of the solution
DHCP	Dynamic Host Configuration Protocol
DNA	Digital Network Architecture
DMZ	Demilitarized Zone
NAT	Network Address Translation
NSO	Network Services Orchestrator
FP	Function Pack (NSO)
NIC	Network Interface Controller
NFVIS	Network Function Virtualization Infrastructure Software
OVS	Open Virtual Switch
SD-WAN	Software-defined Wide Area Networking
SR-IOV	Single Root IO Virtualization
STUN	Session Traversal Utilities for NAT
SWIM	Software Image management
VEPA	Virtual Ethernet Port Aggregator
VM	Virtual Machine
VNF	Virtual Network Function
PNF	Physical Network Function
vNIC	Virtual Network Interface Controller

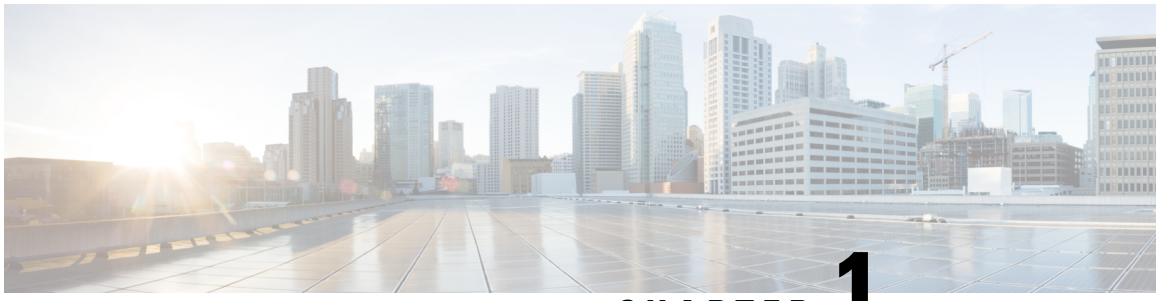
Communications, Services, and Additional Information

- To receive timely, relevant information from Cisco, sign up at Cisco Profile Manager.

- To get the business impact you're looking for with the technologies that matter, visit Cisco Services.
- To submit a service request, visit Cisco Support.
- To discover and browse secure, validated enterprise-class apps, products, solutions and services, visit Cisco Marketplace.
- To obtain general networking, training, and certification titles, visit Cisco Press.
- To find warranty information for a specific product or product family, access Cisco Warranty Finder.

Cisco Bug Search Tool

Cisco Bug Search Tool (BST) is a web-based tool that acts as a gateway to the Cisco bug tracking system that maintains a comprehensive list of defects and vulnerabilities in Cisco products and software. BST provides you with detailed defect information about your products and software.



CHAPTER 1

Information About Cisco SD-WAN Cloud onRamp for Colocation Solution

- Cisco SD-WAN Cloud onRamp for Colocation Solution, on page 1
- Cisco SD-WAN Cloud onRamp for Colocation Solution Components, on page 2

Cisco SD-WAN Cloud onRamp for Colocation Solution

As more applications move to the cloud, the traditional approach of backhauling traffic over expensive WAN circuits to a data center is no longer relevant. The conventional WAN infrastructure was not designed for accessing applications in the cloud. The infrastructure is expensive and introduces unnecessary latency that degrades the experience.

Network architects are reevaluating the design of the WANs to achieve the following:

- Support a cloud transition.
- Reduce network costs.
- Increase the visibility and manageability of the cloud traffic.

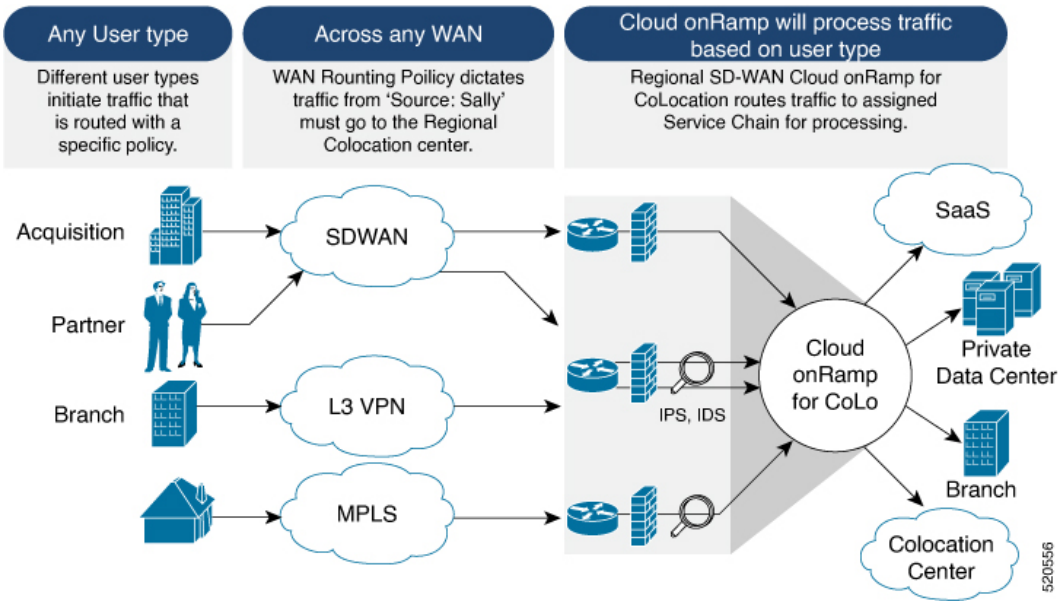
The architects are turning to Software-Defined WAN (SD-WAN) fabric to take advantage of inexpensive broadband Internet services and to route intelligently a trusted SaaS cloud-bound traffic directly from remote branches.

With Cisco SD-WAN Cloud onRamp for Colocation solution built specifically for colocation facilities, the traffic is routed to the best-permissible path from branches and remote workers to where those applications are hosted. The solution also allows distributed enterprises to have an alternative to enabling direct internet access at the branch and enhance their connectivity to infrastructure-as-a-service (IaaS) and software-as-a-service (SaaS) providers.

The solution provides enterprises with multiple distributed branch offices that are clustered around major cities or spread over several countries the ability to regionalize the routing services in colocation facilities. Reason being, these facilities are physically closer to the branches and can host the cloud resources that the enterprise needs to access. So, essentially by distributing a virtual Cisco SD-WAN over a regional architecture of colocation centers, the processing power is brought to the cloud edge.

The following image shows how you can aggregate the access to the multicloud applications from multiple branches to regional colocation facilities.

Figure 1: Cisco SD-WAN Cloud onRamp for CoLocations



The solution can serve four specific types of enterprises:

- Multinational companies that cannot use direct internet connections to the cloud and SaaS platforms due to security restrictions and privacy regulations.
- Partners and vendors without Cisco SD-WAN but still need connectivity to their customers. They do not want to install SD-WAN routing appliances in their site.
- Global organizations with geographically distributed branch offices that require high bandwidth, optimum application performance, and granular security.
- Remote access that need secure VPN connections to an enterprise over inexpensive direct internet links.

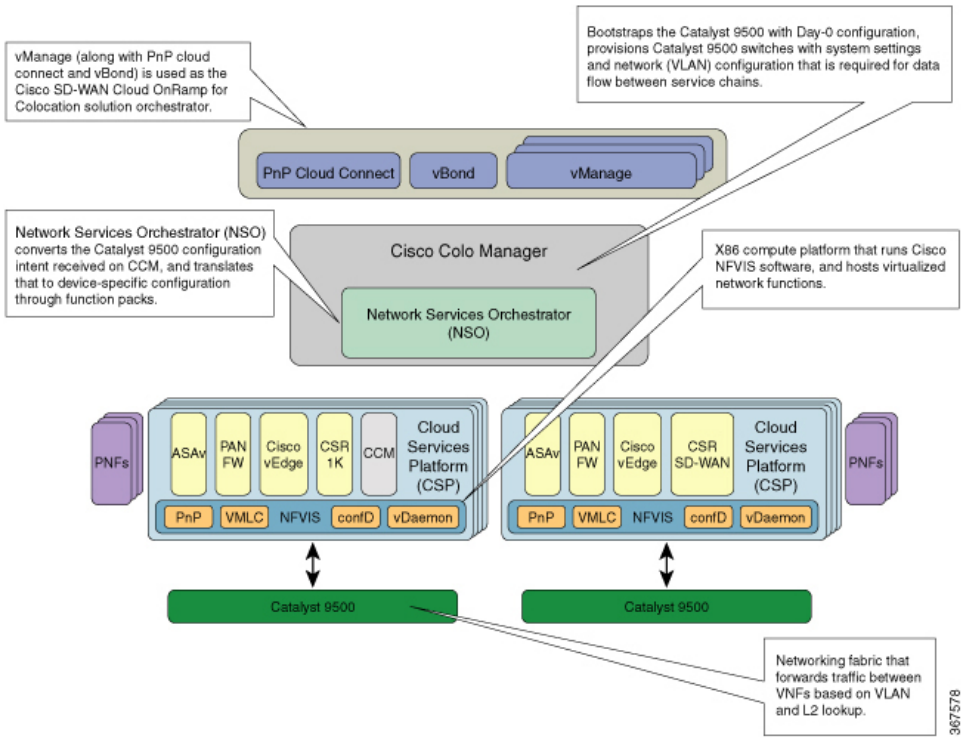
The Cisco SD-WAN Cloud onRamp for Colocation solution can be hosted within certain colocation facilities by a colocation IaaS provider. You can select the colocation provider that meets your needs in a region on a regional basis as long as it supports the necessary components.

Cisco SD-WAN Cloud onRamp for Colocation Solution Components

The Cisco SD-WAN Cloud onRamp for Colocation solution can be deployed in multiple colocations. A colocation is a stack of compute and networking fabric that brings up multiple virtual networking functions and multiple service chains on them. This stack connects branch users, endpoints to a hybrid cloud or data center. Cisco vManage is used as the orchestrator to provision the devices in a colocation. Each colocation does not have visibility of other colocations in the same site or across sites.

The following image shows the components of Cisco SD-WAN Cloud onRamp for Colocation solution.

Figure 2: Cisco SD-WAN Cloud onRamp for Colocation Solution Architectural Overview



- **Cisco Cloud Services Platform, CSP-5444 and CSP-5456**—Cloud Services Platform (CSP) is an x86 Linux hardware platform that runs NFVIS software. It is used as the compute platform for hosting the virtual network functions in the Cisco SD-WAN Cloud onRamp for Colocation solution. Multiple CSP systems can be used in a Cisco SD-WAN Cloud onRamp for Colocation deployment.

  Cisco Network Function Virtualization Infrastructure Software—The Cisco Network Function Virtualization Infrastructure Software (NFVIS) software is used as the base virtualization infrastructure software running on the x86 compute platform. The Cisco NFVIS software provides VM lifecycle management, VM service chaining, VM image management, platform management, PNP for bootstrapping a device, AAA features, and syslog server. See the NFVIS Functionality Changes for SD-WAN Cloud OnRamp for Colocation in NFVIS documentation.

- **Virtual Network Functions** —The Cisco SD-WAN Cloud onRamp for Colocation solution supports both Cisco-developed and third-party Virtual Network Functions (VNFs). The following table includes the validated VNFs and their versions:

Table 2: Validated Virtual Network Functions

Virtual Network Functions	Version
Cisco CSR1000V	17.1.1, 17.2, 17.3
Cisco Catalyst 8000V	17.4.1a
Cisco IOS XE SD-WAN Device	16.12.1, 16.12.2r, 17.2.1r, 17.3.1a
Cisco ASAv	9.12.2, 9.13.1, 9.15.1
CheckPoint	R80.30, R80.40

Virtual Network Functions	Version
Cisco FTDv/NGFW	6.4.0.1, 6.5.0-115
Cisco vEdge Cloud Router	19.2.1, 20.1.1, 20.3.1, 20.4.1
Palo Alto Firewall (PAFW)	9.0.0
Fortinet Firewall	6.0.2

To validate third-party VNFs on the Cisco SD-WAN Cloud onRamp for Colocation solution, you can use the Cisco certification program. For more information about validating third-party VNFs, see https://developer.cisco.com/site/nfv/#the-ecosystem-program .

- **Physical Network Functions**—A Physical Network Function (PNF) is a physical device that is dedicated to provide a specific network function as part of a colocation service chain such as a router or a firewall. The following are the validated PNFs and their versions:

Table 3: Validated Physical Network Functions

Physical Network Functions	Version
Cisco FTD Model: FPR-9300	6.4.0.1, 6.5
Cisco ASR 1000 Series	16.12.1, 17.1, 17.2, 17.3

- **Network Fabric** —Forwards traffic between the VNFs in a service chain by using a L2 and VLAN-based lookup. The last VNF can forward traffic to the network fabric either through L2 or L3 forwarding. The network fabric can include either of the following:
  - Cisco Catalyst 9500-40X switch: Supports 40 10G ports and two 40G ports, which is used as the network fabric
  - Cisco Catalyst 9500-48Y4C switch: Supports 48 1G/10G/25G ports and four 40G/100G ports, which is used as the network fabric.

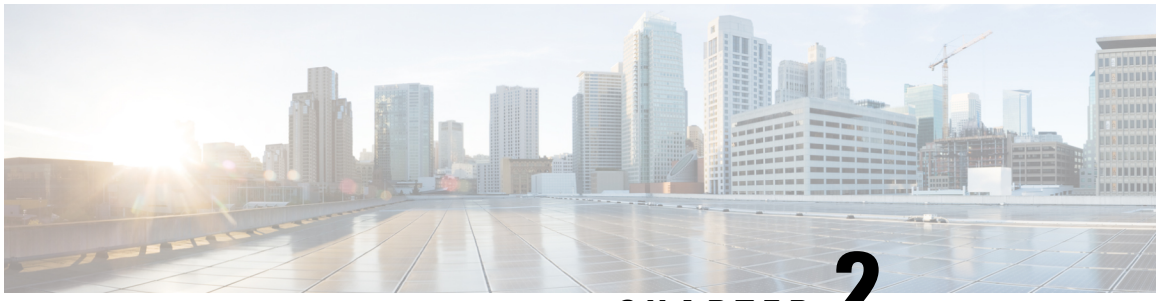
- **Management Network**—A separate management network connects the NFVIS software running on the CSP systems, the virtual network functions, and the switches in fabric. This management network is also used for transferring files and images into and out of the systems. The Out of Band management switch configures the management network. The IP addresses assigned to the CSP devices, Cisco Catalyst 9500-40X or Cisco Catalyst 9500-48Y4C switches are acquired by the management network pool through DHCP configuration. The orchestrator manages VNF management IP addresses and assigns through the VNF Day-0 configuration file.
- **Virtual Network Function Network Connectivity** — A VNF can be connected to the physical network by using either Single Root IO Virtualization (SR-IOV) or through a software virtual switch. A VNF can have one or more virtual network interfaces (VNICs), which can be directly or indirectly connected to the physical network interfaces. A physical network interface can be connected to a software virtual switch and one or more VNFs can share the virtual switch. The Cisco SD-WAN Cloud onRamp for Colocation solution manages the creation of virtual switch instances and the virtual NIC membership to create connectivity. By default, all the physical interfaces and the management interface in the CSP system are available for use by VNFs.

In Cisco SD-WAN Cloud onRamp for Colocation deployments, SR-IOV interfaces are configured in Virtual Ethernet Port Aggregator (VEPA) mode. In this mode, the NIC sends all the traffic that is received from the VNFs to the external Cisco Catalyst 9500-40X or Cisco Catalyst 9500-48Y4C switches. The Cisco Catalyst 9500-40X or Cisco Catalyst 9500-48Y4C transfers the traffic that is based on the L2 MAC address and VLAN. It can send the traffic back to the CSP or to an external connected network. The Catalyst 9500 switch ports that are connected to the CSP interfaces are configured in VEPA mode. When a VLAN is configured on a VNF VNIC, the VLAN must be configured on the connected port on Cisco Catalyst 9500-40X or Cisco Catalyst 9500-48Y4C switches.

A VNF using a SR-IOV interface and a VNF using the software switch can be service chained through the external switch fabric.

- **Physical Network Function Network Connectivity**— A PNF can be connected to the Cisco Catalyst 9500-40X or Cisco Catalyst 9500-48Y4C switch ports, which are the free data ports available from the right side.
- **Service Chains** —In Cisco SD-WAN Cloud onRamp for Colocation solution deployment, the traffic between the VNFs is service chained externally through Cisco Catalyst 9500-40X or Cisco Catalyst 9500-48Y4C. The service chaining requirement provides service chaining functionality to the traffic across VNFs running either on a single CSP or across multiple CSP systems in a cluster. The service chaining is based on the source and destination endpoints in the service chain and is not based on the provider application. In Cisco SD-WAN Cloud onRamp for Colocation solution, L2 (VLAN, destination MAC address) based service chaining has been used.
- **Cisco Colocation Manager** —The Cisco Colocation Manager (CCM) component is a software stack that manages the Cisco Catalyst 9500-40X or Cisco Catalyst 9500-48Y4C switches. In this solution, Cisco Colocation Manager is hosted on NFVIS software in a docker container. The CSP devices host Cisco Colocation Manager along with PNFs and VNFs as shown in the solution architectural overview

  A single CCM instance per cluster is brought up in one of the CSP devices after activating a cluster. The CCM software accepts the Cisco Catalyst 9500-40X or Cisco Catalyst 9500-48Y4C configuration and monitors them. See Configure Cisco SD-WAN Cloud onRamp for Colocation Solution Devices from vManage for more information.
- **Orchestration through Cisco vManage** —Cisco vManage server is used for orchestrating the Cisco SD-WAN Cloud onRamp for Colocation solution. See Cisco SD-WAN Configuration Guides.



CHAPTER 2

Prerequisites and Requirements of Cisco SD-WAN Cloud onRamp for Colocation Solution

- Cisco SD-WAN Cloud onRamp for Colocation Solution Requirements, on page 7
- Prerequisites for Deploying Solution, on page 13
- Ordering and Sizing of Cisco SD-WAN Cloud onRamp for Colocation Solution Devices, on page 14

Cisco SD-WAN Cloud onRamp for Colocation Solution Requirements

The following are the hardware, software, Cloud OnRamp for Colocation cluster, and cabling requirements for deploying Cisco SD-WAN Cloud onRamp for Colocation solution.

Hardware Requirements

The following table lists the hardware requirements:

Table 4: Feature History

Feature Name	Release Information	Description
Support for Cisco Cloud Services Platform, CSP-5456	Cisco SD-WAN Release 20.4.1	Starting from this release, Cisco CSP-5456 is supported on the Cloud onRamp for Colocation solution. The CSP-5456 offers a higher capacity of 56 cores, which maximizes the placement of VNFs in service chains.

Table 5: Hardware Requirements

Components	Hardware Requirements
Compute platform	CSP- 5444 and CSP-5456
Physical form factor	Cisco UCS C240 M5SX (2RU)

Components	Hardware Requirements
Processor cores	CSP-5444: 44 physical cores CSP-5456: 56 physical cores
PCIe NIC slots	6
Disk	8 * 1.2 TB = 9.6 TB
Disk slots	26 (24 useable)
Memory	192 GB of RAM
RAID	12-Gbps SAS HW controller, 4 GB flash-backed write cache (FBWC), RAID 10.
Base Networking	4x1PCIE card in M5 6x1GE Intel i350 ports, 2x1GE LoM **Note** 2-GigE interfaces in a port channel configuration are required for the NFVIS and VM management traffic.
Network Interface Cards (NIC)	2xIntel X520 2-port 10G (Niantic) and Intel XL710 4-port 10G SFP+ (Fortville) **Note** Two Fortville 10G interfaces in port-channel configuration and connected to a virtual switch. This connectivity is required for production traffic to or from the VMs, which support only virtio interface. **Note** Two Fortville 10G interfaces in port-channel configuration and connected to a virtual switch. This configuration is required for VNF HA state synchronization between VNFs hosted on two different CSP systems. **Note** Four Niantic 10G interfaces in SR-IOV mode. The VMs that need high performance and low latency network connectivity to bypass the hypervisor or virtual switch require these interfaces. The VMs that can support SR-IOV must be connected to the SR-IOV virtual function (VFs). Link redundancy is not available in this mode. **Note** For prescriptive connections, ensure that the Fortville NIC (X710) is placed in riser 1, slot-2 and Niantic cards (X520) in riser1, slot 1; and riser 2, slot 4.
Processors (2)	2xIntel Xeon Gold 6152 Series
Power Supplies	Dual power

Components	Hardware Requirements
Network fabric	**Catalyst 9500-40X** Supports forty10G ports and two 40G ports
	Catalyst 9500-48Y4C Supports forty-eight 1G/10G/25G ports and four 40G/100G ports
Management network	Any switch with sufficient number of 1G ports and port channel feature can be used as the management switch. Two switches are recommended to support hardware and link redundancy.

Software Requirements

The following table lists the software requirements:

Table 6: Software Requirements

Components	Software Requirements
Virtualization infrastructure software	Cisco NFVIS Cloud OnRamp for Colocation See Release Notes for Cisco SD-WAN Cloud OnRamp for Colocation Solution.
Orchestration	Cisco vManage See • Cisco SD-WAN Product Documentation for more information. • Cisco SD-WAN Release Notes for more information about the latest vManage features.

All CSP devices and switches must run same version of the software in the Cloud OnRamp for Colocation solution. Any new software version for all devices in a colocation is hosted on Cisco vManage, upon availability.

Supported Platforms and Firmware

The following table lists the supported platform and firmware versions of Cisco NFVIS:

Platform	Firmware	Version
CSP-5444, CSP-5456	BIOS	C240M5.4.1.2b.0.0917201934
	CIMC	4.1(2b)

To upgrade a CIMC version, see the Cisco Host Upgrade Utility User Guide.

Note We recommend that you reach out the Technical Assistance Center (TAC) when upgrading the CIMC version.

Wiring Requirements

Table 7: Feature History

Feature Name	Release Information	Description
Flexible Topologies	Cisco vManage Release 20.3.1 Cisco NFVIS Release 4.2.1	This feature provides the ability to flexibly insert the NIC cards and interconnect the devices (CSP devices and Catalyst 9500 switches) within the Cloud onRamp for Colocation cluster. Any CSP ports can be connected to any port on the switches. The Stackwise Virtual Switch Link (SVL) ports can be connected to any port and similarly the uplink ports can be connected to any port on the switches.

The solution supports both flexible and prescriptive connections between Cisco CSP devices and Cisco Catalyst 9500 switches.

Prescriptive Connections

Prescriptive connections are supported on both Cisco Catalyst 9500-48Y4C and Cisco Catalyst 9500-40X switches.

Ensure that you connect the SVL ports and uplink ports of the Catalyst 9500 switches based on the following information:

Cisco Catalyst 9500-40X

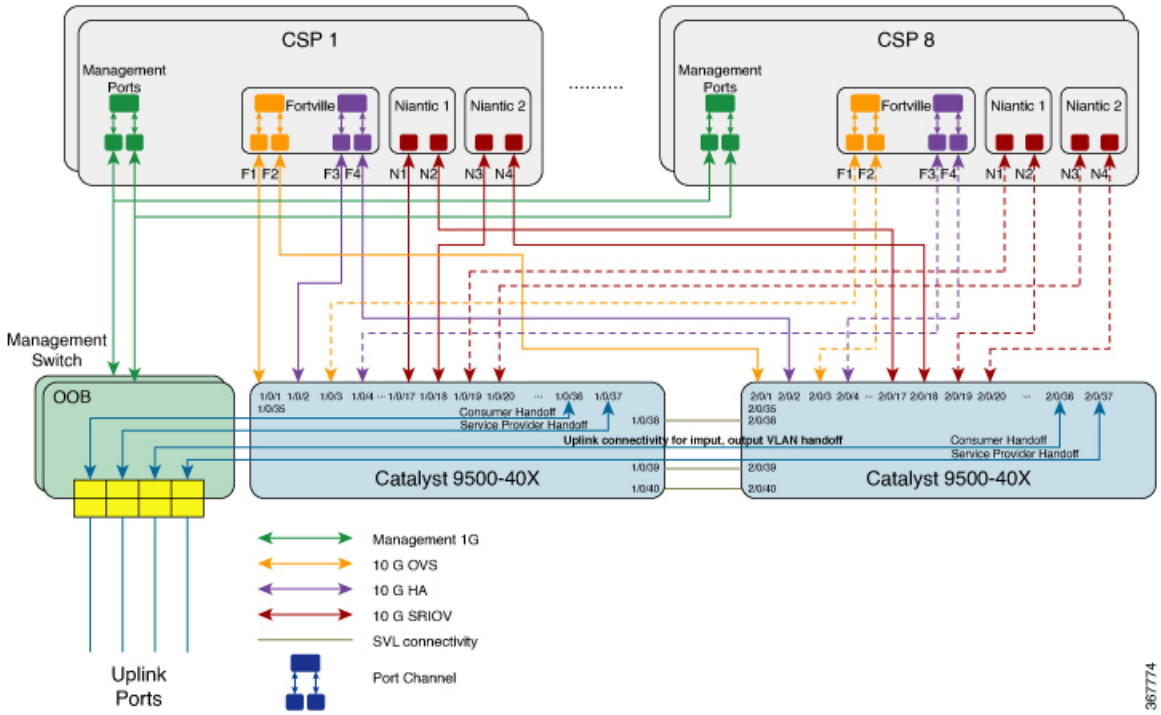
- Stackwise Virtual Switch Link (SVL) ports: 1/0/38-1/0/40, and 2/0/38-2/0/40
- Uplink ports: 1/0/36, 2/0/36 (input VLAN handoff) and 1/0/37, 2/0/37 (output VLAN handoff)

Cisco Catalyst 9500-48Y4C

- SVL ports: 1/0/46-1/0/48, and 2/0/46-2/0/48
- Uplink ports: 1/0/44, 2/0/44 (input VLAN handoff) and 1/0/45, 2/0/45 (output VLAN handoff) for 10G/25G throughput.

The following image shows the high-level design of the physical connectivity for Cisco Catalyst 9500-40X switch.

Figure 3: Prescriptive Connections for Cisco Catalyst 9500-40X



In the preceding topology, each CSP has two 1-GB management ports configured as port channels to the OOB management switch. Each of the Cisco Catalyst 9500-40X switch is connected to the 1-GB port. This connectivity requires two ports on the Management switch per cloud onramp for colocation. The service provider handoff is connected to 10-GB ports on this switch. All service providers ports are trunked into the Cisco Catalyst 9500-40X switch. All the VLANs are configured on all ports of Cisco Catalyst 9500-40X switch.

You can similarly connect the CSP devices with the Cisco Catalyst 9500-48Y4C switches in a prescribed manner.

Note The management switches are not orchestrated and must be manually provisioned. Although the management switches are not orchestrated, ensure that the management switches and devices are connected as per the defined connections.

Flexible Connections

Flexible connections are supported on Cisco Catalyst 9500-40X and Cisco Catalyst 9500-48Y4C switches. For flexible connections:

- Exactly two Niantic cards and one Fortville card should be inserted into a Cisco CSP device in any riser card slot.

Note If you insert the Niantic cards into slots other than riser slots 1 and 4, and Fortville card into any slot other than slot 2, then clean install Cisco NFVIS on the Cisco CSP device after connecting all the cards.

- All data ports on a Cisco CSP device connected to any available ports on Cisco Catalyst 9500-40X or Cisco Catalyst 9500-48Y4C switches.

Note Ensure that you connect all ports on Cisco CSP devices and they are connected in a redundant manner to the primary and secondary switch ports. If all Cisco CSP ports are not connected, the cluster activation process fails.

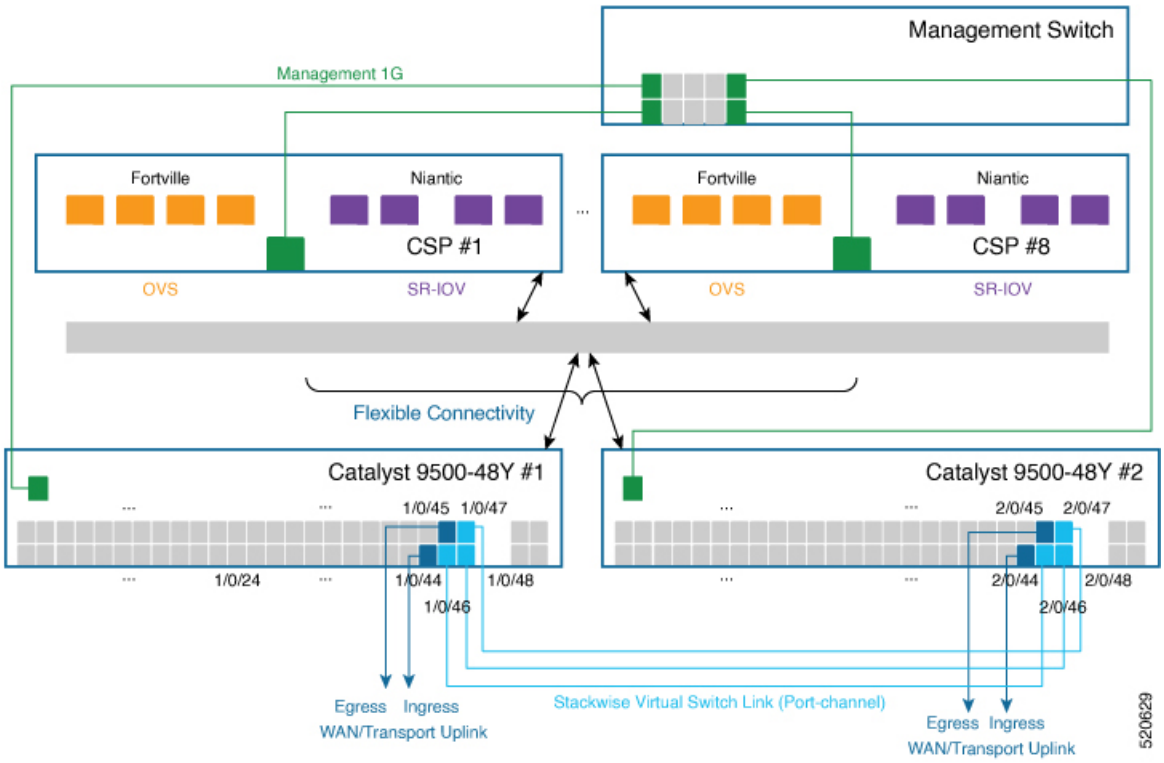
- Connect SVL ports anywhere between 1/0/1-1/0/48 and 2/0/1-2/0/48. It cannot be connected between 1/0/48-1/0/52 or 2/0/48-2/0/52.
- Connect Uplink ports anywhere between 1/0/1-1/0/48 and 2/0/1-2/0/48 for 10G/25G throughput, or between 1/0/49-1/0/52 and 2/0/49-2/0/52 for 40G/100G throughput
- Connect all Niantic and Fortville ports of a Cisco CSP device for redundancy. For example, if Niantic ports are plugged into riser slots 1 and 2 and Fortville ports are plugged into riser slot 4, then you can connect the Cisco CSP interfaces to the switches in either of the following ways:
  - Primary switch: eth1-1, eth2-1, eth4-1, eth4-3

    Secondary Switch: eth1-2, eth2-2, eth4-2, eth4-4
  - Primary switch: eth1-2, eth2-1, eth4-1, eth4-2

    Secondary Switch: eth1-1, eth2-2, eth4-3, eth4-4
- Connect Physical Network functions (PNFs) to any available Cisco Catalyst 9500-40X or Cisco Catalyst 9500-48Y4C switches
- Connect each of the Cisco Catalyst 9500-40X or Cisco Catalyst 9500-48Y4C switch to the 1-GB management port. Each Cisco CSP device has two 1-GB management ports configured as port channels to the OOB management switch. The management switches are not orchestrated through Cisco vManage. Therefore, ensure that you connect the management switches and management ports as shown in the following image.

The following image shows the flexible connectivity between the Cisco CSP devices and Cisco Catalyst 9500-48Y4C switches where the SVL and uplink ports are connected to the default ports.

Figure 4: Cisco SD-WAN Cloud onRamp for Colocation Solution Flexible Connections



Prerequisites for Deploying Solution

The following are prerequisites for deploying the Cisco SD-WAN Cloud onRamp for Colocation solution:

- A minimum of two CSP PID (two Niantics and one Fortville) required. You can order more CSP devices as per the number of service chains that are required per cluster (including HA instances). Also, consider the throughput requirement or number of sessions terminating the cloud onramp for colocation when ordering the number of CSP devices.
- A smart account that is required to propagate the ordered devices to the PNP cloud and vOrchestrator.
- Two Cisco Catalyst 9500-40X or Cisco Catalyst 9500-48Y4C and OOB switches, and a DHCP server per cluster are required.
- Port channel, RJ45 and data SFP along with cables for connectivity are required.
- Cisco ASR 1000 series or a router for WAN termination are required.
- Terminal server for configuring switches and CIMC is required.
- Split management IP pool per cluster into two parts. Configure one part on a DHCP server by considering number of physical devices in a cluster and IP addresses required for broadcast and gateway. Configure the other part of management IP pool on the vManage for VNFs and Cisco Colo Manager. The first IP address in the vManage management pool is used for Cisco Colo Manager. Ensure that you configure this IP address and PNP server for the switch.

- OOB management switch is required to terminate the traffic.

Ordering and Sizing of Cisco SD-WAN Cloud onRamp for Colocation Solution Devices

The cloud onramp for colocation cluster requirements can be categorized into small, medium, large, and extra large clusters that are based on throughput and compute demands.

Consider the following criteria to determine the various cloud onramp for colocation size categories:



Note The cloud onramp for colocation size must be determined before orchestration when ordering the devices such as, CSP devices, and Cisco Catalyst 9500-40X and Cisco Catalyst 9500-48Y4C switches.

- Depending on the number of connections that are required for public clouds and the number of customers trying to reach these clouds, decide the number of required service chains.
- Depending on the policies that must be enforced, decide the number of VMs required in each service chain.
- From the preceding two criteria, you can determine on an average the throughput that is required per service chain.

In a single Cisco SD-WAN Cloud onRamp for Colocation solution deployment, you can deploy between two to eight CSP systems in a cluster.



CHAPTER 3

Get Started with Cisco SD-WAN Cloud onRamp for Colocation Solution

- Cisco SD-WAN Cloud onRamp for Colocation Solution–Deployment Workflow, on page 15
- Install Cisco NFVIS Cloud OnRamp for Colocation on Cisco CSP , on page 16
- Bring up Cisco Cloud Services Platform Devices, on page 19
- Bring up Switch Devices, on page 23
- Bring up Cisco Colo Manager , on page 25
- Provision and Configure Cisco SD-WAN Cloud onRamp for Colocation Solution, on page 26

Cisco SD-WAN Cloud onRamp for Colocation Solution–Deployment Workflow

This topic outlines the sequence of how to get started with the colo devices and build clusters on vManage. Once a cluster is created and configured, you can follow the steps that are required to activate the cluster. Understand how to design service groups or service chains and attach them to an activated cluster. The supported Day-N operations are also listed in this topic.

1. Complete the solution prerequisites and requirements. See Prerequisites and Requirements of Cisco SD-WAN Cloud onRamp for Colocation Solution, on page 7.
    - Complete wiring the CSP devices (set up CIMC for initial CSP access) and Cisco Catalyst 9500-40X or Cisco Catalyst 9500-48Y4C switches (set up console server) along with OOB or management switches. Power on all devices.
    - Set up and configure DHCP server. See Provision DHCP Server Per Colocation, on page 26 .
2. Verify the installed version of Cisco NFVIS and install NFVIS, if necessary. See Install Cisco NFVIS Cloud OnRamp for Colocation on Cisco CSP , on page 16 .
3. Set up or provision a cluster. A cluster constitutes of all the physical devices including CSP devices, and Cisco Catalyst 9500-40X or Cisco Catalyst 9500-48Y4C switches. See Get Started with Cisco SD-WAN Cloud onRamp for Colocation Solution, on page 15.
    - Bring up CSP devices. See Onboard CSP Devices Using Plug-and-Play Process , on page 19.
    - Bring up Cisco Catalyst 9500-40X or Cisco Catalyst 9500-48Y4C switches. See Bring up Switch Devices, on page 23.

- Provision and configure a cluster. See Provision and Configure Cluster, on page 37.

  Configure a cluster through cluster settings. See Cluster Configuration, on page 40.

4. Activate a cluster. See Create and Activate Clusters, on page 38.
5. Design service group or service chain. See Manage Service Groups, on page 64.

Note You can design a service chain and create a service group anytime before creating clusters or activating clusters after all VMs are uploaded to the repository.

6. Attach or Detach service group and service chains to a cluster. See Attach and Detach Service Group with Cluster, on page 87.

Note Service chains can be attached to a cluster after the cluster is active.

7. (Optional) Perform all Day-N operations.
   - Detach a service group to detach service chains. See Attach and Detach Service Group with Cluster, on page 87.
   - Add and delete CSP devices from a cluster. See Add Cloud OnRamp Colocation Devices into vManage, on page 35 and Delete Cloud OnRamp for Colocation Devices from vManage, on page 35.
   - Deactivate a cluster. See Remove Cluster from vManage, on page 60.
   - Reactivate a cluster. See Reactivate Cluster from vManage, on page 63.
   - Design more service group or service chain. See Create Service Chain in a Service Group, on page 64.

Install Cisco NFVIS Cloud OnRamp for Colocation on Cisco CSP

This section provides information about a series of tasks you need to perform to install NFVIS Cloud OnRamp for Colocation on a Cisco CSP device.

Log Into CIMC User Interface

Before you begin

- Ensure that you have configured the IP address to access CIMC.
- If not installed, install Adobe Flash Player 10 or later on your local system.

For details on how to configure an IP address for CIMC, see the *Set up CIMC for UCS C-Series Server* guide on cisco.com.

For information about upgrading CIMC, see the CIMC Firmware Update Utility guide on cisco.com.

Procedure

Step 1 In your web browser, enter the IP address that you configured to access CIMC during initial setup.

Step 2 If a security dialog box displays, do the following:

a) **Optional**: Select the check box to accept all content from Cisco.

b) Click **Yes** to accept the certificate and continue.

Step 3 In the log in window, enter your username and password.

When logging in for the first time to an unconfigured system, use **admin** as the username and **password** as the password.

Step 4 Click **Log In**.

The **Change Password** dialog box only appears the first time you log into CIMC.

Step 5 Change the password as appropriate and save.

The CIMC home page is displayed.

Step 6 From the **CIMC Server** tab, select **Summary**, and click **Launch KVM Console**.

The KVM Console opens in a separate window.

Step 7 From the **Virtual Media** menu on the KVM Console, select **Activate Virtual Devices**.

If prompted with an unencrypted virtual media session message, select **Accept this session**, and click **Apply**. The virtual devices are activated now.

Step 8 From the **Virtual Media** menu on the KVM Console, select **Map CD/DVD**.

Step 9 Browse for the installation file (ISO) on your local system, and select it.

Step 10 Click **Map Device**.

The ISO image file is now mapped to the CD/DVD.

Step 11 From the **CIMC Server** tab, select **BIOS**.

For more information about upgrading BIOS, see the BIOS Upgrade guide on cisco.com.

Step 12 From the **BIOS Actions** area, select **Configure Boot Order**.

The Configure Boot Order dialog box appears.

Step 13 From the **Device Types** area, select **CD/DVD Linux Virtual CD/DVD**, and then click **Add**.

Step 14 Select **HDD**, and then click **Add**.

Step 15 Set the boot order sequence using the **Up** and **Down** options. The **CD/DVD Linux Virtual CD/DVD** boot order option must be the first choice.

Step 16 To complete the boot order setup, Click **Apply**.

Step 17 Reboot the server by selecting the **Power Off Server** option from the Server Summary page in CIMC.

Step 18 After the server is down, select the **Power On Server** option in CIMC.

When the server reboots, the KVM console will automatically install Cisco Enterprise NFVIS from the virtual CD/DVD drive. The entire installation might take 30 minutes to one hour to complete.

Step 19 After the installation is complete, the system is automatically rebooted from the hard drive. Log into the system when the command prompt changes from "localhost" to "nfvis" after the reboot.

Wait for some time for the system to automatically change the command prompt. If it does not change automatically, press **Enter** to manually change the command prompt from "localhost" to "nfvis". Use **admin** as the login name and **Admin123#** as the default password.

Note
The system prompts you to change the default password at the first login. You must set a strong password as per the on-screen instructions to proceed with the application. You cannot run API commands or proceed with any tasks unless you change the default password at the first login. API will return 401 unauthorized error if the default password is not reset.

Step 20 You can verify the installation using the System API or by viewing the system information from the Cisco Enterprise NFVIS portal.



Note Ensure that the RAID configuration is 4.8 TB RAID-10. To configure RAID through CIMC, see the Cisco UCS Servers RAID Guide on cisco.com.

Activate Virtual Device

You will have to launch the KVM Console to activate virtual devices.

Before you begin

Ensure that you have the Java 1.6.0_14 or a higher version installed on your local system.

Procedure

Step 1 Download the Cisco Enterprise NFVIS image from a prescribed location to your local system.

Step 2 From CIMC, select the **Server** tab, and click **Launch KVM Console**.

Note
A JNLP file will be downloaded to your system. You must open the file immediately after it is downloaded to avoid the session timeout.

Step 3 Open the renamed *.jnlp* file. When it prompts you to download Cisco Virtual KVM Console, click **Yes**. Ignore all security warnings and continue with the launch.

The KVM Console is displayed.

Step 4 From the **Virtual Media** menu on the KVM Console, select **Activate Virtual Devices**.

If prompted with an unencrypted virtual media session message, select **Accept this session**, and click **Apply**. The virtual devices are activated now.

Map NFVIS Cloud OnRamp for Colocation Image

Procedure

Step 1 From the **Virtual Media** menu on the KVM Console, select **Map CD/DVD...**.

Step 2 Browse for the installation file (ISO) on your local system, and select it .

Step 3 Click **Map Device**.
The ISO image file is now mapped to the CD/DVD.

Step 4 From the KVM console, power cycle (warm reboot) and system installation process starts and NFVIS is installed.

Bring up Cisco Cloud Services Platform Devices

Table 8: Feature History

Feature Name	Release Information	Description
Onboarding CSP Device with Day-0 Configuration Using USB Drive	Cisco SD-WAN Release 20.4.1	This feature enables you to onboard CSP devices by loading the Day-0 configuration file to a USB drive. Use this onboarding option when you can't access the Internet to reach the Plug-and-Play Connect server.

To bring up the Cisco Cloud Services Platform (CSP) devices, you can use the following options:

- Automated deployment: Securely onboards and deploys CSP devices with factory settings into the Cisco SD-WAN network during the Day-0 configuration. The deployment dynamically discovers the IP address of Cisco vBond Orchestrator using the Plug-and-Play (PnP) process for Cisco CSP devices.
- Bootstrap deployment: Requires you to share the configuration files with the CSP devices. You can either create a configuration file and copy it to a bootable USB, or add the configuration file to the USB. The bootable USB is connected and available on the devices at the time of bootup.

Onboard CSP Devices Using Plug-and-Play Process

This topic describes how the bringing up of Cisco CSP devices are automated using the PnP process.

Before you begin

- Ensure that you connect the CSP devices as per the prescribed topology, and power them on.
- Connect the Plug-and-Play (PnP) supported interface to the WAN transport (typically Internet).

Power on a Cisco CSP device. The following process occurs:

Procedure

Step 1 When the device boots up, it obtains the IP address, default gateway, and DNS information through the DHCP process on the supported PnP interface of the device.

Step 2 The device connects with the Cisco cloud hosted PnP Connect server and shares its chassis or serial number with the PnP server to be authenticated by it.

Step 3 After authentication, the PnP Connect portal provides the device with information about the Cisco vBond Orchestrator, organization name, and root certificates.

For deployments that use enterprise root-ca certificate, information about Cisco vBond Orchestrator IP address or DNS, organization-name, and enterprise root-ca certificate are downloaded on the device from the PnP Connect portal using the HTTPS protocol. The device uses this information to initiate control connections with the Cisco vBond Orchestrator.

You can view the availability of the device and association with the Cisco vBond Orchestrator on the PnP interface through the PnP Connect portal.

Step 4 The PnP Connect portal then displays a **Redirect Successful** status when the device is redirected through PnP to the Cisco vBond Orchestrator.

Step 5 After authentication with the Cisco vBond Orchestrator, the device is provided with Cisco vManage and Cisco vSmart Controller information to register and establish a secure connection.

Step 6 The device attempts to establish a secure control connection with the Cisco vManage server.

Step 7 After authentication with the Cisco vBond Orchestrator, the Cisco vManage server responds to the device with the system IP of the device and reauthenticates the device using the shared system-ip information.

Step 8 To join the Cisco SD-WAN overlay network, the device reinitiates control connections to all the SD-WAN controllers using the configured `system-ip` IP address.

Onboard CSP Devices Using USB Bootstrapping Process

If you're unable to use the automated discovery option, use this deployment option to configure the factory-shipped device, which comes without any configuration.

We recommend this deployment option when:

- The device is connected to a private WAN transport (MPLS) that can’t provide a dynamic IP address.
- Internet access isn't available to reach the Plug-and-Play Connect server.

Points to Consider

- The USB drive can have multiple Day-0 configuration files, which are identified by the serial number of the device in the file name. This naming convention enables you to use the same USB drive for bootstrapping multiple devices.
- The supported Day-0 configurations included in the configuration file are:
  - Static IP configuration of the device
  - Cisco vBond Orchestrator IP address and the port configuration
  - DNS server and domain name configuration

- The bootstrap configuration can be uploaded to a USB key and inserted into a device at the install site.

Before you begin

- The device must be in factory default state with no added configuration.
- The device must be installed with a fresh image of Cisco NFVIS.
- The USB drive must be Virtual File Allocation Table (VFAT) formatted to recognize and automount the drive. Insert the USB drive into a laptop or desktop to format it.
- The device should be able to reach the Cisco vBond Orchestrator.

Procedure

Step 1 Create a configuration file on the root folder of the USB drive.

Ensure that the configuration file name is, *nfvis_config_SERIAL.xml*, where

SERIAL represents the serial number of the CSP device.

For example,

nfvis_config_ WZP232903K6.xml

Step 2 Copy the following to the configuration file.

```
<config xmlns="http://tail-f.com/ns/config/1.0">
       <vm_lifecycle xmlns="http://www.cisco.com/nfvis/vm_lifecycle">
  <networks>
  <network>
    <name>int-mgmt-net</name>
    <subnet>
      <name>int-mgmt-net-subnet</name>
      <address>192.168.30.6</address>
      <netmask>255.255.255.0</netmask>
      <gateway>192.168.30.1</gateway>
    </subnet>
  </network>
  </networks>
       </vm_lifecycle>

<system xmlns="http://viptela.com/system">
        <organization-name>vIPtela Inc Regression</organization-name>
        <sp-organization-name>vIPtela Inc Regression</sp-organization-name>
       <vbond>
          <remote>172.23.191.87</remote>
          <port>12346</port>
        </vbond>
      </system>

<vpn xmlns="http://viptela.com/vpn">.
        <vpn-instance>
          <vpn-id>0</vpn-id>
          <interface>
            <if-name>colo-mgmt</if-name>
            <tunnel-interface>
              <encapsulation>
                <encap>ipsec</encap>
```

```
    </encapsulation>
  </tunnel-interface>
  <shutdown>false</shutdown>
</interface>
</vpn-instance>
</vpn>
</config>
```

Note
It's mandatory to copy the above-mentioned static IP configuration of the device to the configuration file. The static IP configuration of the device is represented by the following Day-0 configurations:

`<address></address>, <netmask></netmask>, and <gateway></gateway>`

Step 3 Insert the USB drive into the Cisco CSP device and power on the device.

When the device boots up, the device searches for the configuration file in the bootable USB drive. After the file is located, the device suspends the PnP process and loads the bootstrap configuration file.

Step 4 Remove the USB drive.

Note
If you don't unmount the USB drive and reboot the device after the configuration has been applied, the USB drive configuration isn't reapplied. The CSP device isn't in Factory Data Reset (FDR) state or restored to its original system state.

Step 5 To access a CSP device, SSH to a static IP address provided in Step 2 such as, 192.168.30.6.

Step 6 Change the default password at the first login when the system prompts you to change.

Ensure that you set a strong password based on the on-screen instructions. You can't run API commands or proceed with any tasks unless you change the default password at the first login.

What to do next

To verify the device onboarding process, proceed to Verify Onboarded Devices and Activate Devices, on page 22.

Verify Onboarded Devices and Activate Devices

Procedure

Step 1 Log in to Cisco vManage with admin credentials using the URL HTTPS: //vManage-ip-address/.

Step 2 Click **Configuration** > **Devices**.

From the list of devices, the CSP devices that have the serial number with the word token aren't yet onboarded. To authenticate these devices with the SD-WAN controllers, Cisco vManage provides a One-Time Password (OTP). The OTP is autogenerated by Cisco vManage after adding the CSP device in the SD-WAN controller authorized device list.

Step 3 Under the **Valid** column, verify the validity of the installed certificate of all the listed CSP devices. See Failures with Certificate installation, on page 147. Also, verify if root CA has been installed. See CSP hasn't established connectivity with vManage, on page 149.

Note
For device onboarding using enterprise root-ca certificates, the CSP device receives the root certificates, along with the Cisco vBond Orchestrator and organization name information from the PnP Connect portal.

Step 4 To activate the CSP device and associate the chassis number and the Serial No (one-time password) with the CSP device, on the CLI of the CSP device, use the following command:

request activate chassis-number *chassis-number* **token** *token-number*

For more information about the **request device** command, see request device.

Example:

request activate chassis-number *CSP-5444-serial-number* **token** *70d43cfbd0b3b426da63dba2dd4f4c49*

Step 5 To bring up the remaining CSP devices, repeat Steps 1–4 for each of the CSP devices.

Bring up Switch Devices

This section describes about how Cisco Catalyst 9500-40X or Cisco Catalyst 9500-48Y4C switch devices are brought up through the Day-0 configuration.

Before you begin

Ensure that you note the following before bringing up the switch devices:

- Cisco Catalyst 9500-40X or Cisco Catalyst 9500-48Y4C switch devices have both Network-Advantage and DNA-Advantage licenses. To verify the available licenses on the switch devices, use the following command:

  ```
  Device# show license status
  ```

  To know more about the license usage information, see the **show license usage** command.

- Either PNP redirect setup or manual PNP profile being set on the switch devices is required. For a PNP redirect setup, add switches SN and Cisco Colo Manager IP address to PNP, and add entries of devicehelper.cisco.com to OOB router of the network if the DHCP server is on OOB router. For example,

  ```
  #conf t
  #ip host devicehelper.cisco.com <OOB router of the network>
  ```

- Ensure that both switches are connected as per the SVL mode configuration.

Procedure

Step 1 Clean the switch configuration if they have been previously used.

a) Renumber switch, which is required for SVL stack mode.

Note
Ensure that you do not touch the switches during SVL mode. Also, do not perform any action such as, pressing enter or space, which can cause switches to complete SVL.

Use the **show switch** command to determine the switch number and whether the provisioned switch exists in the switch stack. If the switch number is two, then use the **switch 2 renumber 1** command, and then erase the configuration.

b) To erase the switch startup configuration and return it to its initial state, use the **write erase** command.

c) To reload the switch with a new configuration, use the following commands in privileged EXEC mode and enter **no** for not saving the modified configuration:

```
switch(config)#reload
```

Note
You do not need to save the configuration.

d) Perform steps b and c on the secondary switch device after the switch stack reloading has been completed. This action ensures that the secondary switch device is reloaded twice.

Step 2 After Cisco Catalyst 9500-40X or Cisco Catalyst 9500-48Y4C switch boots up, it gets an IP address from the local DHCP server and initiates PNP discovery.

Step 3 The DHCP server with option 43 enables Cisco Catalyst 9500-40X or Cisco Catalyst 9500-48Y4C switch to reach the PNP server in Cisco Colo Manager.

The Cisco Colo Manager IP address is the PNP server IP address of a cluster on vManage. Ensure that DHCP server with option 43 always point to the port, 9191.

Example:

The following is an example of local PNP server for switches:

```
ip dhcp pool Cat9k
network 10.114.11.39 255.255.255.0
dns-server 172.31.232.182
default-router 172.31.232.182
option 43 ascii "5A;B2;K4;I10.114.11.40;J9191”
```

Where, 10.114.11.40 is the local PNP server or Cisco Colo Manager IP address.

The output after setting DHCP server with option 43 to port, 9191 is:

```
ip dhcp excluded-address 172.31.232.182 172.31.232.185
ip dhcp excluded-address 172.31.233.182
ip dhcp excluded-address 172.31.232.254
ip dhcp excluded-address 172.31.23.10 172.31.23.49
ip dhcp excluded-address 172.31.23.52 172.31.23.100
ip dhcp excluded-address 172.31.23.252
ip dhcp excluded-address 172.31.23.253
ip dhcp excluded-address 172.31.23.230 172.31.23.250
!
```

Step 4 After the switches reach the PNP server on Cisco Colo Manager, it pushes the Day-0 configuration. The Day-0 configuration push happens if a cluster is activated on vManage. If a cluster is not activated, the Cisco Catalyst 9500-40X or Cisco Catalyst 9500-48Y4C switches reach the PNP server on Cisco Colo Manager every minute and stays in backoff mode.

After the switch devices are brought up, the SSH connection and NETCONF sessions on the switch devices are enabled for Cisco Colo Manager to push Day-N configuration and ongoing switch management is continued.

Example

About Uplink Ports 36 and 37 in Prescriptive Connections

For prescriptive connections, ports 36 (input VLAN handoff) and 37 (output VLAN handoff) are reserved for uplink ports.

Note The 1/0/36, 1/0/37 and 2/0/36, 2/0/37 switch ports are configured in "active" mode. If a user is not using port channel and not connected to ports 36 and 37, the OOB switch ports that are connected to Cisco Catalyst 9500-40X on ports 36 or 37 must be configured as "passive" mode.

For example,

- **interface Port-channel1 switchport trunk allowed VLAN 100-106**

```
example VLANs
 switchport mode trunk
!
```

- **interface TenGigabitEthernet1/0/1**

```
port connected to cat9k 1/0/36 or 1/0/37
 switchport mode trunk
 channel-group 1 mode passive
 spanning-tree portfast
!
```

- **interface TenGigabitEthernet1/0/2**

```
interface TenGigabitEthernet1/0/2
 switchport mode trunk
 channel-group 1 mode passive
 spanning-tree portfast
!
```

What to do next

To bring up another switch, repeat all the mentioned steps in sequence for the next switch.

Bring up Cisco Colo Manager

This section describes about how Cisco Colo Manager is brought up. The Cisco Colo Manager acts as a PNP agent for the Catalyst 9K switches in a cluster. It takes care of the Day-0 configuration push to the Catalyst 9K switches and also relays the configuration from vManage to Catalyst 9K.

Note During cluster activation process, Cisco Colo Manager is automatically brought up.

Procedure

Step 1 All CSP devices in the cloud onramp for colocation establish a DTLS tunnel with vManage.

Step 2 vManage selects one CSP device by sending a NETCONF action API to bring up Cisco Colo Manager on that CSP device.

Step 3 Cisco Colo Manager is in "Starting" state when it is brought up. Cisco Colo Manager can then move to "Healthy" or "Unhealthy" state depending on the health check status.

What to do next

After switch configuration and once colo manager is up, both switches reach the colo manager. Ensure that you check the PNP list on Cisco Colo Manager to verify that both the switch devices have called home. See Switch devices are not calling home to PNP or Cisco Colo Manager, on page 140.



Note For activation to continue, both switches must call home.

Provision and Configure Cisco SD-WAN Cloud onRamp for Colocation Solution

To order Cisco SD-WAN Cloud onRamp for Colocation PID, choose Cisco SD-WAN Cloud onRamp for Colocation on Cisco Commerce Workspace (CCW).

Customer-specific order details such as, Smart Account name, Virtual Account name must be provided while ordering.

To provision and configure the Cisco SD-WAN Cloud onRamp for Colocation solution, perform the following:

1. Ensure that Cloud Service Platform (CSP) devices and Cisco Catalyst 9500-40X or Cisco Catalyst 9500-48Y4C switches are cabled as per the prescribed or flexible connections, and powered on.
2. The Smart Account synchronizes customer-specific device order details with PNP Connect and vOrchestrator.

Provision DHCP Server Per Colocation

To manage IP addresses of the physical devices such as switches, VNFs, and CSP devices, you must configure a DHCP server per colocation. The Cisco Colo Manager IP address can be configured in DHCP option 43 for Cisco Catalyst 9500-40X or Cisco Catalyst 9500-48Y4C to reach Cisco Colo Manager.

vManage fixes and assigns Cisco Colo Manager IP addresses for a colocation. It manages and assigns IP addresses of all VNFs through Day-0 configuration.



Note The subnet for both physical (CSP devices, switches) and virtual appliances (Cisco Colo Manager, VNF) must be same.

You can pick an appropriate subnet for a colocation and limit the pool for IP addresses depending on the number of CSP devices and switches in a colocation. vManage picks the first IP address entered in the VNF management IP pool in the vManage interface and configures it as the (Switch PNP Server IP) Cisco Colo Manager IP address. The second and third IP addresses form the management pool are used for switch management IP addresses. The **Switch PNP Server IP** field can be edited to provide an alternative IP address

if a different IP address is configured in the DHCP server for PNP of switches. The remaining IP addresses from the vManage pool are assigned to the remaining VNFs in the colocation.

Note Ensure that you set up a DNS server in each colocation.

Device Port Connectivity Details and Service Chaining for Prescriptive Connections

In Cisco SD-WAN Cloud onRamp for Colocation solution deployments, the Cisco Catalyst 9500-40X switches connected to CSP systems perform service chaining. If VMs support SR-IOV, Cisco Catalyst 9500-40X switches perform service chaining, whereas VMs without SR-IOV support, service chaining is done by Open Virtual Switch (OVS).

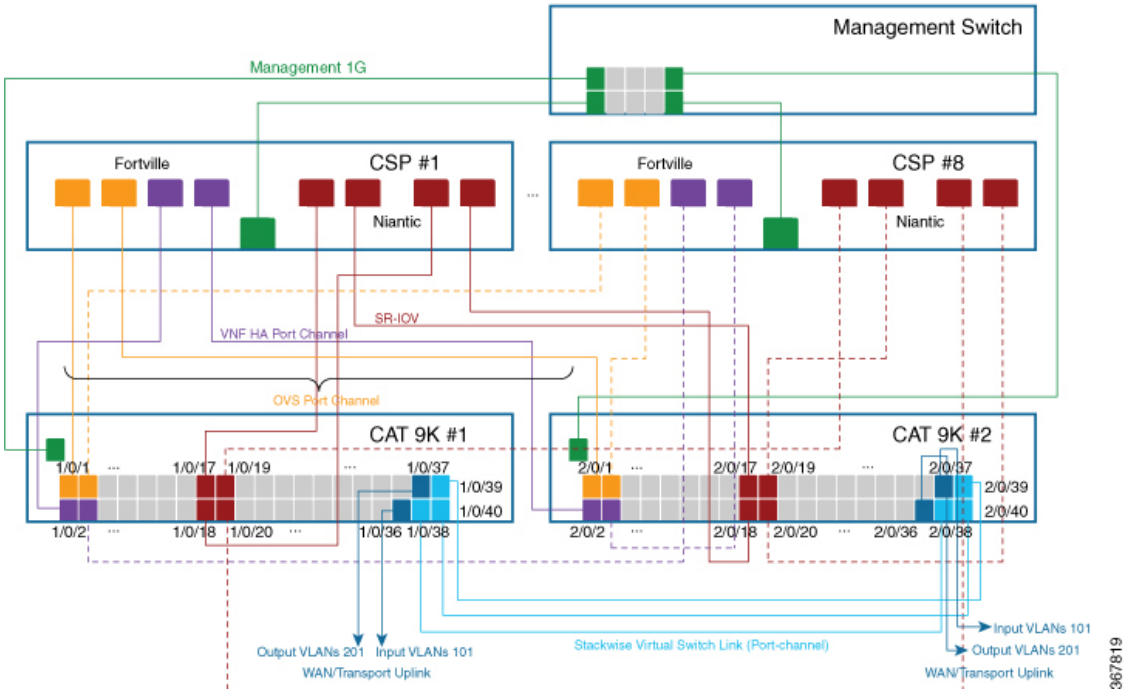
Virtual switch-based service chains are used for High Availability traffic and control traffic.

VLAN-based L2 service chaining from Cisco Catalyst 9500-40X switch is used for Cisco SD-WAN Cloud onRamp for Colocation solution. In this service chaining, each virtual NIC interface of a VM in a service chain is configured on the same access VLAN on a CSP virtual switch. The switch pushes the VLAN tag of the packets entering and leaving the vNIC interface. The VNF can remain unaware of the next service in the service chain. To forward traffic between the VNFs hosted either on the same CSP or across different CSP devices in a cluster, the physical switch with the matching VLAN gets configured.

In Cisco SD-WAN Cloud onRamp for Colocation solution deployments, the deja-vu check is disabled on the switch ports that are connected to the CSP system for unicast traffic.

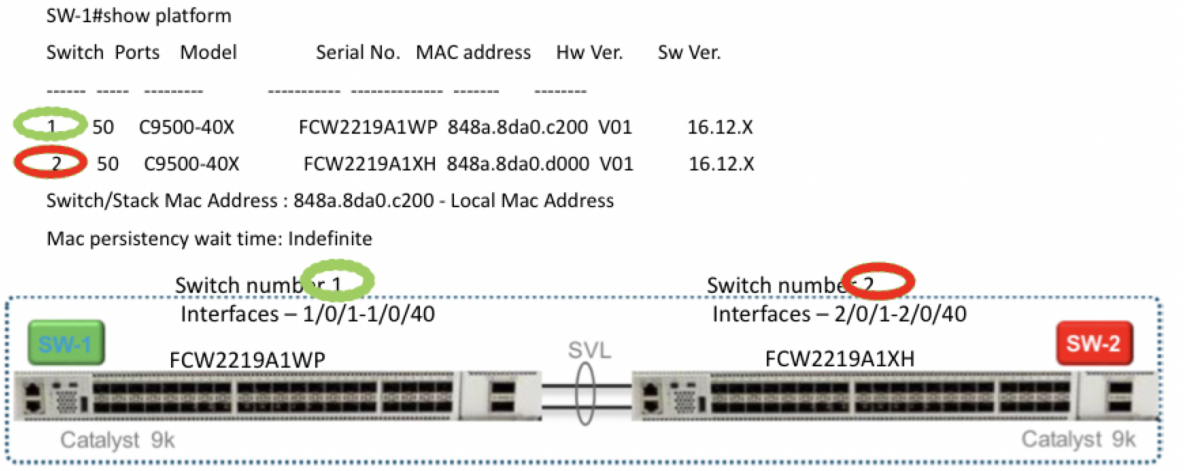
The following topology displays connectivity of the CSP ports to Cisco Catalyst 9500-40X switches and OOB switch.

Figure 5: Service Chain Connectivity with OVS, VEPA Enabled Switch Ports



The following is the location of an interface in switches:

Note The location of an interface is applicable once switches are in SVL mode after successful cluster activation.



The following ports are VEPA disabled and configured with port channels:

- 1/0/1-1/0/16
- 2/0/1-2/0/16

The following ports are VEPA enabled and port channels configuration is disabled:

- 1/0/17-1/0/32
- 2/0/17-2/0/32

Note VEPA ports are only applicable to SRIOV interfaces.

The following ports are the WAN connectivity ports:

- 1/0/36, 2/0/36—Connect port 1/0/36 to receive outside traffic from branch/VPN connections (via an OOB switch).
- 1/0/37, 2/0/37—Connect port 1/0/37 to forward service chain traffic to specific VLANs that is mapped to provider networks on an OOB switch.

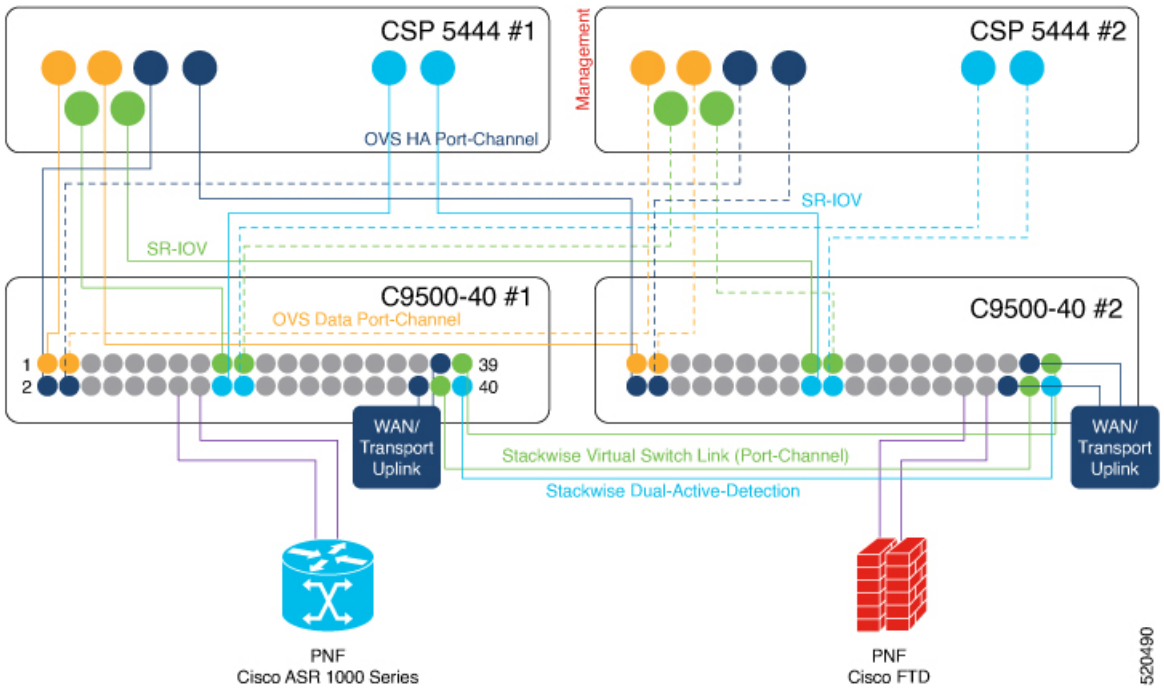
You can connect the ports as follows:

- Data ports—Connect ports 1/0/1-1/0/35 to CSP devices. To achieve redundancy and HA across switches, you can connect two ports to one CSP and the other two can be connected to next CSP. For example, ports 1/0/1 and 2/0/1 is used for data and HA respectively can be connected to the first CSP, CSP #1. Next, 1/0/2 and 2/0/2 is another port channel that is connected to the next CSP, CSP #2, and so on. Hence, the OVS ports consume all eight CSP devices.
- WAN connectivity ports—Connect port 1/0/36 on configured VLAN/s to receive outside traffic (Input VLAN handoff). Connect port 1/0/37 to forward service chain traffic to specific VLANs that is mapped to provider networks (Output VLAN handoff). External input or output VLAN traffic can come from branch or VPN connections and provider networks terminate at the Cloud OnRamp for Colocation through the OOB switch. For each service chain configured in the cluster and input or output VLAN configured for each service chain, the configuration on the ports, 36 and 37 occurs during service chain deployment.

  If ports 36 or 37 are connected to the OOB switch and not using port channels, ensure that all VLAN handoffs are configured either on input or output VLAN handoffs correspondingly. For example, if port 36 is connected, configure all VLAN handoff on input VLAN handoff for a service chain. If port 37 is s connected, configure all VLAN handoff on output VLAN handoff for a service chain.
- Connect ports 1/0/38-1/0/40 in Stackwise Virtual Switch Link (SVL) configuration.

  The following cabling image shows how the physical network functions are connected to the Cisco Catalyst 9500-40X switches.

Figure 6: PNF Cabling Image



The following table provides the ports available for PNF:

Table 9: Ports on Cisco Catalyst 9500-40X Switches for PNF

Number of CSP Devices	Number of PNFs	Switch Ports available for PNFs on First Switch	Switch Ports available for PNFs on Second Switch
7	1	1/0/15-1/0/16, 1/0/31-1/0/32	2/0/15-2/0/16, 2/0/31-2/0/32
6	2	1/0/13-1/0/16, 1/0/29-1/0/32	2/0/13-2/0/16, 2/0/29-2/0/32
4	4	1/0/11-1/0/16, 1/0/27-1/0/32	2/0/11-2/0/16, 2/0/27-2/0/32

To remove CSP devices and shuffle ports, perform the following steps:

1. If all eight CSP devices are connected to switches and if you want to connect a PNF device to the switches:
   a. Deactivate or remove the eighth CSP (CSP connected to the right most data ports on switch) from the cluster by using the RMA workflow on vManage.
   b. Disconnect the CSP physical connections on Cisco Catalyst 9500-40X switches.
   c. Connect the PNF device in place of the disconnected CSP.

2. If one of the first seven CSP devices must be removed to make additional ports available for PNF, perform the following steps:
   a. Perform the steps mentioned in 1.
   b. Move the right most connected CSP that is the eighth CSP to the ports that are made available by the removed CSP.

      For example, if the first CSP is removed, move the eighth CSP to the position of the first CSP and connect the PNF in place of the eighth CSP.

For this phase of Cisco SD-WAN Cloud onRamp for Colocation solution deployment, full chain VNF configuration is supported. In a full chain configuration, all the VNFs for the producer and consumer chains are part of a single service chain. The VNFs are not shared across different types of producers and consumers. A separate instance of a service chain supports each combination of consumer and producer type. For a full chain configuration, all the VNFs in a chain are L2 service chained.

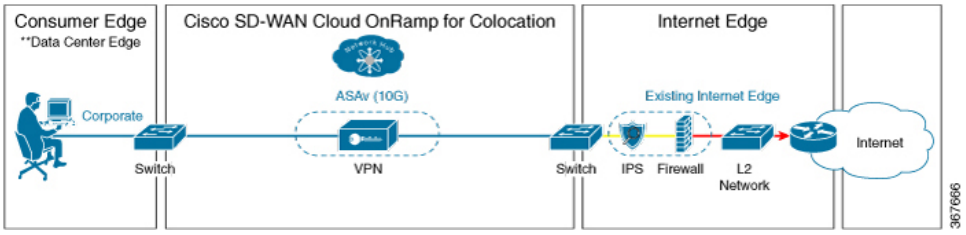
The vManage orchestrator manages the Cisco SD-WAN Cloud onRamp for Colocation solution service chain configuration. vManage assigns the VLANs from the VLAN pool that is provided for the colocation to the individual VM vNICs and configures the switch with appropriate VLANs. The VNFs can remain unaware about the service chain. Apart from the Day-0 VNF configuration, vManage does not configure the individual VNFs that take part in the service chain. .

Validated Service Chains

In Cisco SD-WAN Cloud onRamp for Colocation solution deployments, the following are the four validated service chains that you can deploy within a cluster from vManage. For all the validated service chains, each VM can be instantiated in HA or standalone modes.

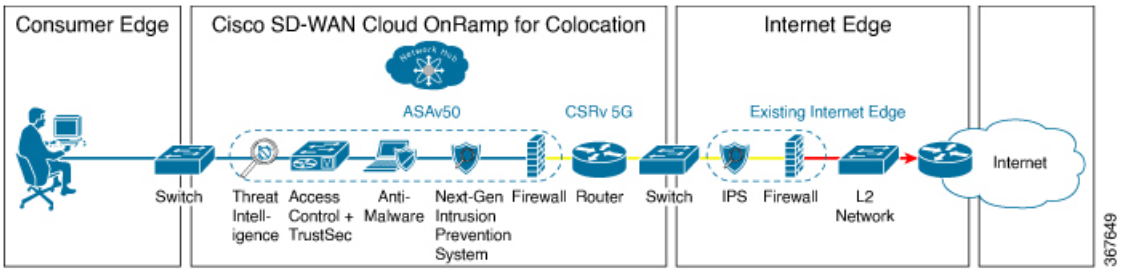
- Employee Remote VPN Access—In this service chain, there is a firewall, which can be in L3 VPN HA or L3 VPN non-HA modes. The firewall VNFs can be ASAv, Palo Alto Networks Firewall, Firepower_Threat_Defense_Virtual (FTDv). Here, ASAv is in routed mode, no Day-0 configuration support for the VPN connect, no BGP on consumer chain, and no VLANs.

  Figure 7: Employee Remote VPN Access Service Chain



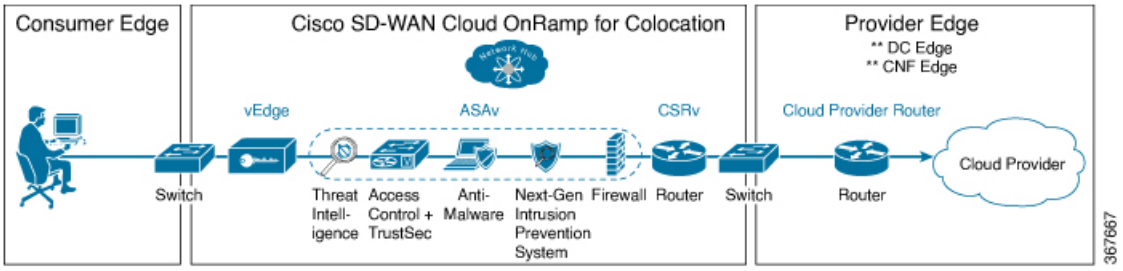
- Internet Edge (Outbound Internet, eCommerce, SaaS)—In this service chain, a firewall is followed with a router. The firewall modes can be L3-VLAN HA and L3-VLAN non-HA. The routers can be in L3 HA and L3 non-HA modes. Here, ASAv is always in routed mode. One VLAN handoff is required and inbound subinterfaces can be up to four. The termination can be in routed mode or in a trunk mode with subinterfaces up to four. You can choose the hypervisor tagged VLANs versus VNF to do the VLAN tagging. In VNF VLAN tagging, you can terminate to a minimum of 1 VLAN and maximum of 4 VLANs. In hypervisor tagged VLANs, all VLANs are tagged in the same inbound VNF interface.

Figure 8: Internet Edge Service Chain



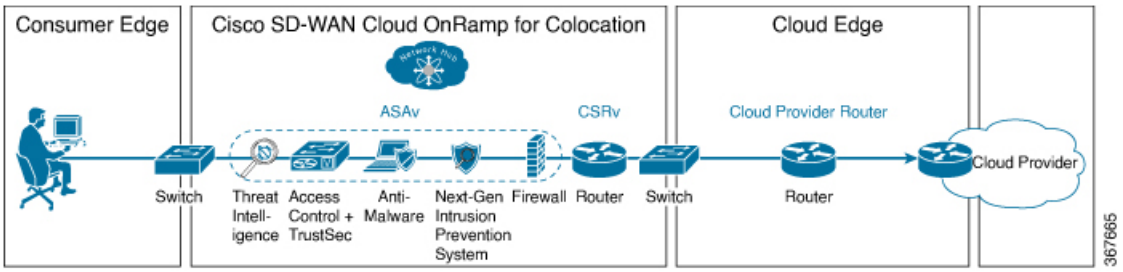
- SD-WAN Access—In this service chain, vEdge is followed by a firewall, which is followed by a router. The firewall modes can be L2 HA, L2 non-HA, L3 HA and L3 non-HA. The routers can be in L3 HA and L3 non-HA modes.

Figure 9: SD-WAN Access Service Chain



- Cloud Edge (Public Cloud Access)—In this service chain, firewall is followed by a router, where the firewall is in routed mode. The firewall modes can be, L3 HA and L3 non-HA. The routers can be in L3 HA and L3 non-HA modes. This service chain is Internet Edge (Outbound Internet, eCommerce, SaaS) with firewall mode being L3.

Figure 10: Cloud Edge (Public Cloud Access) Service Chain



See Create Service Chain in a Service Group, on page 64 topic about how you can choose the validated service chains through vManage.

Validated VM Packages

VM packages are created as per use cases. These packages have recommended Day-0 configuration for each supported use case. Any user can bring the required custom Day-0 configuration and package the VM as per their requirement. In the validated packages, various Day-0 configurations are bundled into a single VM package. For example, if a VM is a firewall VM, it can be used in transparent or routed mode if it is in the

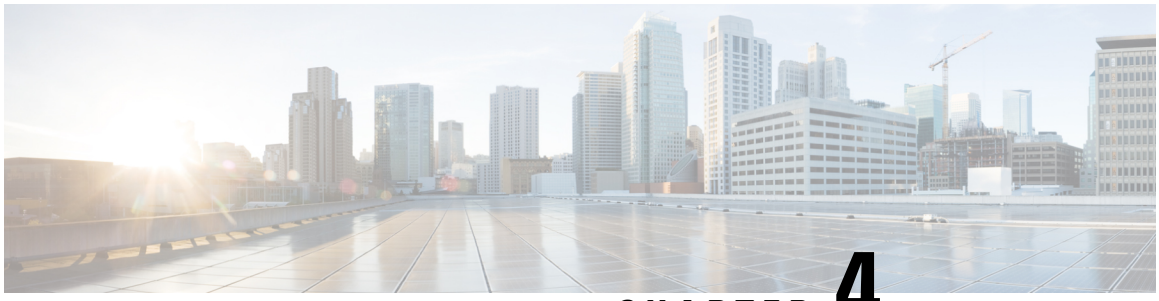
middle of a service chain. If a VM is the first or last VM in a service chain, it can be a terminating tunnel to a branch or provider, or routed traffic, or can terminate multiple branches, or a provider. Each use case is set up as a special tag in image metadata for a user to make a selection at deployment or while provisioning a service chain. If a VM is in the center of a service chain, vManage can automate the IP addresses and VLANs for those segments. If VM is terminating to a branch or provider, user must configure the IP addresses, peer addresses, autonomous system numbers, and others.

Customized Service Chains

Service chains are a named list of service-functions and associated endpoint-group through which packets flow. You can customize service chains and create service chain templates. A service chain template is a chain of VMs serving the intent of connecting the ingress traffic to the cloud. Service chain templates can have predefined service chains containing validated VMs .

The first VNF and the last VNF in a customized service chain can be a router (or firewall). In SD-WAN case, the first VM is a vEdge, which is orchestrated. In non-SD-WAN case, the first VM can be modeled as a gateway router, which is not orchestrated.

You can choose a service chain template and modify the template by inserting one or more VMs and delete one or more VMs. For each VM in the service chain, you can select the VM image that has been brought up from the VM catalog. For example, if the first VM in the service chain is a ROUTER, you can select either Cisco 1000v, or choose from VM repository, or any third-party router.



CHAPTER 4

Configure Cisco SD-WAN Cloud onRamp for Colocation Solution Devices from vManage

- Add Cloud OnRamp Colocation Devices into vManage, on page 35
- Delete Cloud OnRamp for Colocation Devices from vManage, on page 35
- Manage Clusters in vManage, on page 36
- Manage Service Groups, on page 64
- Attach and Detach Service Group with Cluster, on page 87
- Day-N Configuration Workflow of Cisco SD-WAN Cloud onRamp for Colocation Solution, on page 88

Add Cloud OnRamp Colocation Devices into vManage

You can add CSP devices, switch devices, and VNFs through vManage.

Before you begin

Ensure that you have the following setup details such as:

- Cisco SD-WAN setup details such as, vManage IP address and credentials, vBond IP address and credentials.
- NFVIS setup details such as, CSP device CIMC IP address and credentials or UCSC CIMC IP address and credentials.
- Able to access both the switch consoles.

When you order the Cisco SD-WAN Cloud onRamp for Colocation solution PID, all cloud onramp for colocation device information is received though smart account that can be accessed by vManage.

Delete Cloud OnRamp for Colocation Devices from vManage

To delete the CSP devices from vManage, perform the following steps:

Before you begin

Ensure that you consider the following:

- If any service chains are attached to a device that is being deleted, detach service groups. See Attach and Detach Service Group with Cluster, on page 87.
- If a CSP device that is being deleted is hosting Cisco Colo Manager, see Recovery of Cisco Colo Manager, on page 127.

Procedure

Step 1 Log into Cisco vManage, at the URL HTTPS: //vManage-ip-address/, as a user with "admin" credentials.
Step 2 In Cisco vManage, in the **Configuration** > **Certificates** screen, locate the device, click **More Actions**, and click **Invalid**.
Step 3 In the **Configuration** > **Certificates** screen, click **Send to Controller**.
Step 4 In the **Configuration** > **Devices** screen, in the **WAN Edge List** tab, select the device.
Step 5 Click the **More Actions** icon to the right of the row and click **Delete WAN Edge**.
Step 6 Click **OK** to confirm deletion of the device.

Deleting a device removes its serial and chassis numbers from the **WAN edge router serial number** list, and also permanently removes its configuration from the Cisco vManage.

Manage Clusters in vManage

Use the Cloud onRamp for Colocation screen to configure a colocation cluster and service groups that can be used with the cluster.

The three steps to configure are:

- Create a cluster. See Create and Activate Clusters, on page 38.
- Create a service group. See Create Service Chain in a Service Group, on page 64.
- Attach a cluster with a service group. See Attach and Detach Service Group with Cluster, on page 87.

A colocation cluster is a collection of two to eight CSP devices and two switches. The supported cluster templates are:

- Small cluster—2 Cisco Catalyst 9500-40X or Cisco Catalyst 9500-48Y4C+2 CSP
- Medium Cluster—2 Cisco Catalyst 9500-40X or Cisco Catalyst 9500-48Y4C+4 CSP
- Large Cluster—2 Cisco Catalyst 9500-40X or Cisco Catalyst 9500-48Y4C+6 CSP
- X-Large Cluster—2 Cisco Catalyst 9500-40X or Cisco Catalyst 9500-48Y4C+8 CSP



Note Ensure that you add a minimum of two CSP devices one-by-one to a cluster. You can keep adding three, four, and so on, up to a maximum of eight CSP devices. You can edit a Day-N configuration of any cluster, and add pairs of CSP devices to each site up to a maximum of eight CSP devices.

Ensure that all devices that you bring into a cluster have the same software version.

Note You can't use the CSP-5444 and CSP-5456 devices in the same cluster.

Following are the cluster states:

- Incomplete—When a cluster is created from the vManage interface without providing the minimum requirement of two CSP devices and two switches. Also, cluster activation is not yet triggered.
- Inactive—When a cluster is created from the vManage interface after providing the minimum requirement of two CSP devices and two Switches, and cluster activation is not yet triggered.
- Init—When the cluster activation is triggered from the vManage interface and Day-0 configuration push to the end devices is pending.
- Inprogress—When one of the CSP devices within a cluster comes up with control connections, the cluster moves to this state.
- Pending—When the Day-0 configuration push is pending or VNF install is pending.
- Active—When a cluster is activated successfully and NCS has pushed the configuration to the end device.
- Failure—If Cisco Colo Manager has not been brought up or if any of the CSP devices that failed to receive an UP event.

A cluster transitioning to an active state or failure state is as follows:

- **Inactive** > **Init** > **Inprogress** > **Pending** > **Active**—Success
- **Inactive** > **Init** > **Inprogress** > **Pending** > **Failure**—Failure

During a cluster creation, cluster clearing, and cluster deletion, ensure that you clean the configurations of both switches. See Troubleshoot Catalyst 9500 Issues, on page 140 for more information about cleaning switch configuration that has been used previously.

Provision and Configure Cluster

This topic describes about activating a cluster that enable deployment of service chains.

To provision and configure a cluster, perform the following:

1. Create a cluster by adding two to eight CSP devices and two switches.

   CSP devices can be added to a cluster and configured through vManage before bringing them up. You can configure CSP devices and Catalyst 9K switches with the global features such as, AAA, default user (admin) password, NTP, syslog, and more.

2. Configure cluster parameters including IP address pool input such as, service chain VLAN pool, VNF management IP address pool, management gateway, VNF data plane IP pool, and system IP address pool.

3. Configure a service group.

   A service group consists of one or more service chains.

Note You can add a service chain by selecting one of the predefined or validated service chain template, or create a custom one. For each service chain, configure input and output VLAN handoff and service chain throughput or bandwidth, as mentioned. The service chain throughput is in MBPS, and you can assign as high as 10 GB, and as low as 10 MB. The default service chain bandwidth is 10 MBPS. See the Ordering and Sizing of Cisco SD-WAN Cloud onRamp for Colocation Solution Devices topic.

4. Configure each service chain by selecting each VNF from the service template. Choose a VNF image that is already uploaded to the VNF repository to bring up the VM along with required resources (CPU, memory, and disk). Provide the following information for each VNF in a service chain:
   - The specific VM instance behavior such as, HA, shared VM can be shared across service chains.
   - Day-0 configuration values for tokenized keys and not part of the VLAN pool, management IP address, or data HA IP address. The first and last VMs handoff-related information such as peering IP and autonomous system values must be provided. The internal parameters of a service chain are automatically filled by the orchestrator from the VLAN or Management or Data Plane IP address pool provided.
5. Add the required number of service chains for each service group and create the required number of service groups for a cluster.
6. To attach a cluster to a site or location, activate the cluster after all configuration has been completed.

   You can watch the cluster status change from in progress to active or error.

To edit a cluster, perform the following:

1. Modify the activated cluster by adding or deleting service groups or service chains.
2. Modify the global features configuration such as, AAA, system setting, and more.

You can predesign a service group and service chain before creating a cluster. They can be attached with a cluster after the cluster is active.

Create and Activate Clusters

This topic provides the steps on how you can form a cluster with CSP devices, Cisco Catalyst switches as a single unit, and provision the cluster with cluster-specific configuration.

Before you begin

- Ensure that you synchronize the clocks for Cisco vManage and CSP devices. To synchronize a clock for CSP devices, configure the NTP server for CSP devices when you enter information about cluster settings.
- Ensure that you configure the NTP server for Cisco vManage and Cisco vBond Orchestrator. To configure the NTP server, see the Cisco SD-WAN System and Interface Configuration Guide.
- Ensure that you configure the OTP for the CSP devices to bring up the CSP devices.
- Ensure that you power on both the Catalyst 9500 switches and ensure that they are operational.

Procedure

Step 1 Log in to Cisco vManage with admin credentials using the URL HTTPS: //vManage-ip-address/.

Step 2 In Cisco vManage, click **Configuration** > **Cloud OnRamp for Colocation**.

a) In the **Cluster** tab, click the **Configure & Provision Cluster** button.

b) Provide the following information:

Table 10: Cluster Information

Field	Description
Cluster Name	The cluster name can contain 128 alphanumeric characters.
Description	The description can contain 2048 alphanumeric characters.
Site ID	Specifies the overlay network site identifier. Ensure that the value you enter for Site ID is similar to the organizations Site ID structure for the other Cisco SD-WAN overlay elements.
Location	The location can contain 128 alphanumeric characters.
Cluster Type	To configure a cluster in a multitenant mode so that it can be shared across multiple tenants, select **Shared**. **Note** In the single-tenant mode, the cluster type **Non Shared** is selected by default.

c) To configure switches, click a switch icon in the **Switches** box. The **Edit Switch** dialog box is displayed. Provide a name and choose the switch serial number from the drop-down list. Click **Save**.

The switch name can contain128 alphanumeric characters.

The switch serial numbers that you view in the drop-down list are obtained and integrated with Cisco vManage using the PnP process. These serial numbers are assigned to switches when you order Cisco SD-WAN Cloud onRamp for Colocation solution PID on CCW and procure the switches.

Note
You can keep the serial number field blank for switches and CSP devices, design your cluster, and then edit the cluster later to include the serial number after you procure the devices. However, you can't activate a cluster with the CSP devices or switches without the serial numbers.

d) To configure another switch, repeat step c.

e) To configure CSP devices, click a CSP icon in the **Appliances** box. The **Edit CSP** dialog box is displayed. Provide a CSP device name and choose the CSP serial number from the drop-down list. Click **Save**.

The CSP device name can contain 128 alphanumeric characters.

f) Configure OTP for the CSP devices to bring up the devices.

g) To add remaining CSP devices, repeat Step e.

h) Click **Save**.
After you create a cluster, on the cluster configuration page, an ellipsis enclosed in a yellow circle appears next to a device where the serial number hasn't been assigned for the device. You can edit a device to enter the serial numbers.
i) To edit a CSP device configuration, click a CSP icon, and perform the process in substep e.
j) To set the mandatory and optional global parameters for a cluster, on the cluster configuration page, enter the parameters under **Cluster Configuration**. See Cluster Configuration, on page 40.
k) Click **Save**.
You can view the cluster that you created in a table.

Step 3 To activate a cluster, on the **Cluster** tab,
a) Click a cluster.
b) Click the **More Actions** icon at the right of the row in the table.
c) Click **Activate** against the cluster.

When you activate the cluster, Cisco vManage establishes a DTLS tunnel with the CSP devices in the cluster, where it connects with the switches through Cisco Colo Manager. When the DTLS tunnel connection is running, a CSP device in the cluster is chosen to host the Cisco Colo Manager. Cisco Colo Manager is brought up and vManage sends global parameter configurations to the CSP devices and Cisco Catalyst 9500 switches. For information about cluster activation progress, see Progress of Cluster Activation, on page 48.

Cluster Configuration

The cluster configuration parameters are:

Login Credentials

1. In the **Cluster Topology** screen, click **Add** against **Credentials**. In the **Credentials** configuration window, enter the following:
   - (Mandatory) Template Name—The template name can contain 128 alphanumeric characters.
   - (Optional) Description—The description can contain 2048 alphanumeric characters.
2. Click **New User**.
   - In the **Name** field, enter the username.
   - In the **Password** field, enter the password and confirm the password in the **Confirm Password** field.
   - In the **Role** drop-down list, select administrators.
3. Click **Add**.
   The new user with username, password, and role with action appears.
4. Click **Save**.
   The login credentials for the new user are added.
5. To cancel the configuration, click **Cancel**.
6. To edit the existing credential for the user, click **Edit** and save the configuration.

Resource Pool

1. In the **Cluster Topology** screen, click **Add** against **Resource Pool**. In the **Resource Pool** configuration window, enter the following:
   - Name—The name of the IP address pool should contain 128 alphanumeric characters.
   - Description—The description can contain 2048 alphanumeric characters.
2. In the **DTLS Tunnel IP** field, enter the IP addresses to be used for the DTLS tunnel. To enter multiple IP addresses, separate them by commas. To enter a range, separate the IP addresses with a hyphen (for example, 172.16.0.180-172.16.255.190).
3. In the **Service Chain VLAN Pool** field, enter the VLAN numbers to be used for service chains. To enter multiple numbers, separate them by commas. To enter a numeric range, separate the numbers with a hyphen (for example, 1021-2021).

   Consider the following when entering the VLAN information:

   1002-1005 are reserved VLAN values, and they shouldn't be used in the cluster creation VLAN pool.

Note Valid VNF VLAN pool: 1010-2000 and 1003-2000

Invalid: 1002-1005 (shouldn't be used)

Caution 1002-1005 isn't allowed for configuration. Allowed VLANs should be contiguous.

Example: Enter data VLAN pool as 1006-2006. Ensure that this VLAN range isn't used in the Input/Output VLAN during service chain creations.

4. In the **VNF Data Plane IP Pool** field, enter the IP addresses to be used for auto configuring data plane on a VNF interface. To enter multiple IP addresses, separate them by commas. To enter a range, separate the IP addresses with a hyphen (for example, 10.0.0.1-10.0.0.100).
5. In the **VNF Management IP Pool** field, enter the IP addresses to be used for theVNF. To enter multiple IP addresses, separate them by commas. To enter a range, separate the IP addresses with a hyphen (for example, 192.168.30.99-192.168.30.150).

Note These addresses are IP addresses for secure interfaces.

6. In the **Management Subnet Gateway** field, enter the IP address of the gateway to the management network. It enables DNS to exit the cluster.
7. In the **Management Mask** field, enter the mask value for the failover cluster. For example, /24 and not 255.255.255.0
8. In the **Switch PNP Server IP** field, enter the IP address of the switch device.

Note The IP address of the switch is automatically picked from the management pool, which is the first IP address. You can change it if a different IP is configured in the DHCP server for the switch.

9. Click **Save**.

Port Connectivity

Prerequisites

- When configuring the SVL and uplink ports, ensure that the port numbers you configure on Cisco vManage match the physically cabled ports.
- Ensure that you assign serial numbers to both the switches. See Create and Activate Clusters, on page 38 .

To configure the SVL and uplink ports:

In the **Cluster Topology** screen, click **Add** next to **Port Connectivity**. In the **Port Connectivity** configuration window, both the configured switches appear. Hover over a switch port to view the port number and the port type.

Before you change the default port number and port type, note the following information about Cisco Catalyst 9500-40X and Cisco Catalyst 9500-48Y4C switches:

- The following are the default SVL and uplink ports of Cisco Catalyst 9500 switches.

  Cisco Catalyst 9500-40X

  - Stackwise Virtual Switch Link (SVL) ports: 1/0/38-1/0/40, and 2/0/38-2/0/40
  - Uplink ports: 1/0/36, 2/0/36 (input VLAN handoff) and 1/0/37, 2/0/37 (output VLAN handoff)

- *Cisco Catalyst 9500-48Y4C*

  - SVL ports: 1/0/46-1/0/48, and 2/0/46-2/0/48
  - Uplink ports: 1/0/44, 2/0/44 (input VLAN handoff) and 1/0/45, 2/0/45 (output VLAN handoff) for 10G/25G throughput.

- I, E, and S represent the ingress, egress, and SVL ports.
- If the physical cabling is same as the default configuration, click **Save**.

To change the default ports when the connectivity is different for SVL or uplink ports, perform the following:

1. If both the switches are using the same ports:
   a. Click a port on a switch that corresponds to a physically connected port.
   b. To add the port configuration to the other switch, check the **Apply change** check box.

   If both the switches are not using the same ports:
   a. Click a port on **Switch1**.

b. Choose a port type from **Port Type** drop-down list.

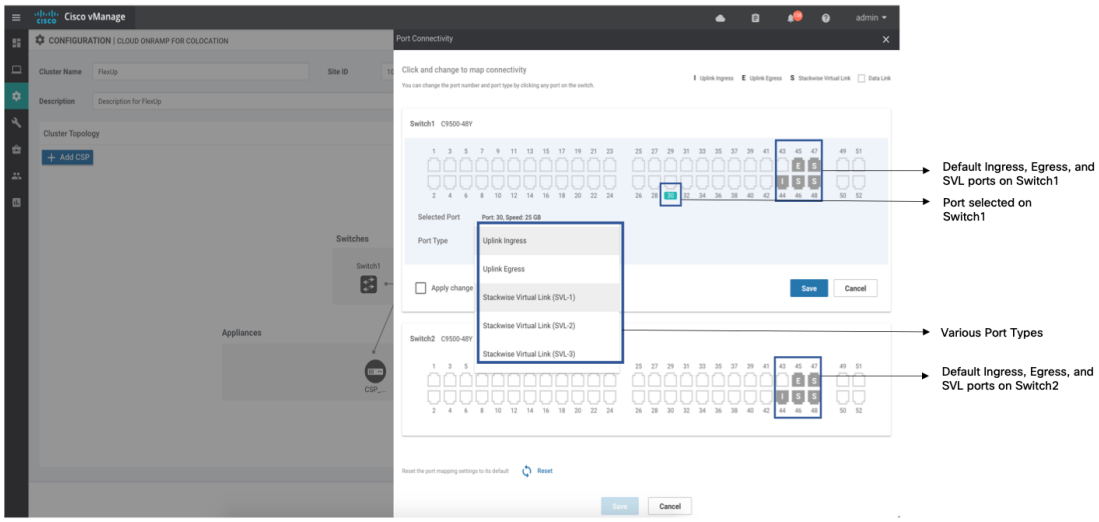
c. Click a port on **Switch2** and then choose the port type.

2. To add another port, repeat step 1.
3. Click **Save**
4. To edit port connectivity information, in the **Cluster Topology** screen, click **Edit** next to **Port Connectivity**.

Note You can modify the SVL and uplink ports of a cluster when the cluster has not been activated.

5. To reset the ports to default settings, click **Reset**.

Figure 11: Port Connectivity for Flexible Connections



The remaining ports (SR-IOV and OVS) on the Cisco CSP devices and the connections with switches are automatically discovered using Link Layer Discovery Protocol (LLDP) when you activate a cluster. You don't need to configure those ports.

Cisco Colo Manager (CCM) discovers switch neighbor ports and identifies whether all Niantic and Fortville ports are connected. If any port is not connected, CCM sends notifications to Cisco vManage that you can view in the task view window.

NTP

Optionally, configure the NTP server for the cluster:

1. In the **Cluster Topology** screen, under **Other Settings**, click **Add** against **NTP**. In the **NTP** configuration window, enter the following:

- Template Name—Name of the NTP template should be in alphanumeric characters and the name should be within 128 characters.
- Description—The description should be in alphanumeric characters and can be upto 2048 characters.

2. In the **Preferred server** field, enter the IP address of the primary NTP server.
3. In the **Backup server** field, enter the IP address of the secondary NTP server.
4. Click **Save**.

   The NTP servers are added.
5. To cancel the NTP server configuration, click **Cancel**.
6. To edit the NTP server configuration details, click **Edit**.

Syslog Server

Optionally, configure the syslog parameters for the cluster:

1. In the **Cluster Topology** screen, under **Other Settings**, click **Add** against **Syslog**. In the **Syslog** configuration window, enter the following:
   - Template Name—Name of the system template should be in alphanumeric characters and the name should be within 128 characters.
   - Description—The description can be up to 2048 characters and can contain only alphanumeric characters.
2. In the **Severity** drop-down list, select the severity of syslog messages to be logged.
3. To add a new syslog server, click **New Server**.

   Type the IP address of a syslog server.
4. Click **Save**.
5. To cancel the configuration, click **Cancel**.
6. To edit the existing syslog server configuration, click **Edit** and save the configuration.

TACACS Authentication

Table 11: Feature History

Feature Name	Release Information	Description
TACACS Authentication	Cisco SD-WAN Release 20.3.1 Cisco vManage Release 20.3.1	This feature allows you to configure the TACACS authentication for users accessing the Cisco CSP and Cisco Catalyst 9500 devices. Authenticating the users using TACACS validates and secures their access to the Cisco CSP and Cisco Catalyst 9500 devices.

The TACACS authentication determines the valid users who can access the Cisco CSP and Cisco Catalyst 9500 devices after a cluster is active.

Points to consider

- By default, the admin users with Role-based access control (RBAC) are authorized to access the Cisco CSP and Cisco Catalyst 9500 devices.
- Do not configure the same user with different passwords when configuring using TACACS and RBAC. If same user with a different password is configured on TACACS and RBAC, the RBAC user and password authentication is used. For information about how to configure RBAC on the devices, see Login Credentials, on page 40.

To authenticate users:

1. To add TACACS server configuration, in the **Cluster Topology** screen, click **Other Settings** > **Add** next to **TACACS**.

   To edit TACACS server configuration, in the **Cluster Topology** screen, click **Other Settings** > **Edit** next to **TACACS**.

   In the **TACACS** configuration window, enter information about the following:

   - Template Name—The TACACS template name can contain 128 alphanumeric characters.
   - (Optional) Description—The description can contain 2048 alphanumeric characters.

2. To add a new TACACS server, click **+ New TACACS SERVER**.
   - In **Server IP Address**, enter the IPv4 address.

     Use IPv4 addresses for hostnames of TACACS server.
   - In **Secret**enter the password and confirm the password in **Confirm Secret**.

3. Click **Add**

   The new TACACS server details are listed in the **TACACS** configuration window.



Note You can add a maximum of four TACACS servers.

4. To add another TACACS server, repeat step 2 to step 3.

   When authenticating users, if the first TACACS server is not reachable, the next server is verified until all the four servers are verified.

5. Click **Save**
6. To delete a TACACS server configuration, choose a row from the TACACS server details list and click **Delete** under **Action**.



Note To modify an existing TACACS server information, ensure to delete a TACACS server and then add a new server.

7. To view the TACACS server configuration, in Cisco vManage, click **Configuration** > **Devices**.

From the device table, choose a Cisco CSP device or Cisco Catalyst 9500 switch, click **More Actions** > **Running Configuration**.

Backup Server Settings

Points to Consider

- If you don't use an NFS server, Cisco vManage can’t successfully create backup copies of a CSP device for future RMA requirements.
- The NFS server mount location and configurations are same for all the CSP devices in a cluster.
- Don't consider an existing device in a cluster as the replacement CSP device.



Note If a replacement CSP device isn’t available, wait until the device appears in Cisco vManage.

- Don't attach further service chains to a cluster after you identify that a CSP device in the cluster is faulty.
- The backup operation on a CSP device creates backup files containing NFVIS configuration and VMs (if VMs are provisioned on the CSP device). You can use the following information for reference.
  - An automated backup file is generated and is in the format:

    serial_number + "_" + time_stamp + ".bkup"

    For example,

    `WZP22180EW2_2020_06_24T18_07_00.bkup`

  - An internal state model is maintained that specifies the status of the overall backup operation and internal states of each backup component:
    - NFVIS: A configuration backup of the CSP device as an xml file, config.xml.
    - VM_Images: All VNF tar.gz packages in `data/intdatastore/uploads` which are listed individually.
    - VM_Images_Flavors: The VM images such as, img_flvr.img.bkup.
    - Individual tar backups of the VNFs: The files such as, vmbkp.
  - The backup.manifest file contains information of files in the backup package and their checksum for verification during restore operation.

To create backup copies of all CSP devices in a cluster, perform the following steps:

1. On the **Cluster Topology** screen, click **Other Settings** > **Add** next to **Backup**.

   To edit backup server settings, on the **Cluster Topology** screen, click **Other Settings** > **Edit** next to **Backup**

   In the **Backup** configuration window, enter information about the following fields:
   - Mount Name—Enter the name of the NFS mount after mounting an NFS location.
   - Storage Space—Enter the disk space in GB.

- Server IP: Enter the IP address of the NFS server.
- Server Path: Enter the folder path of the NFS server such as, `/data/colobackup`
- Backup: Click **Backup** to enable it.
- Time: Set a time for scheduling the backup operation.
- Interval: Choose from the options to schedule a periodic backup process.
  - Daily: The first backup is created a day after the backup configuration is saved on the device, and everyday thereafter.
  - Weekly: The first backup is created seven days after the backup configuration is saved on the device, and every week thereafter.
  - Once: The backup copy is created on a chosen day and it's valid for the entire lifetime of a cluster. You can choose a future calendar date.

2. Click **Save**.
3. To view the status of the previous five backup operations, use the **show hostaction backup status** command. To know about the backup status configuration command, see Backup and Restore NFVIS and VM Configurations. To use this command:
   a. In Cisco vManage, click the **Tools** > **SSH Terminal** screen to start an SSH session with Cisco vManage.
   b. Choose the CSP device.
   c. Enter the username and password for the CSP device and click **Enter** to log in to the CSP device and run the **show hostaction backup status** command.

Restore CSP Device

You can perform the restore operation only by using the CLI on the CSP device that you're restoring.

1. Use the **mount nfs-mount storage** command to mount NFS:

   For more information, see Network File System Support.



Note To access the backup file, the configuration for mounting an NFS file system should match the faulty device. You can view this information from other healthy CSP devices as the NFS mount location and configurations are same for all the CSP devices. To view and capture the information, you can do one of the following:

- In Cisco vManage, go to the **Cluster Topology** screen, and click **Add** next to **Backup** under **Cluster Configuration**.
- Use the **show running-config** command to view the active configuration that is running on a CSP device. See Prerequisites and Restrictions for Backup and Restore of CSP Devices.

mount nfs-mount storage { *mount-name* | **server_ip** *server_ip* | **server_path** *server_path* | **storage_space_total_gb** *storage_space_total_gb* | **storage_type** *storage_type* }

For example, `mount nfs-mount storage nfsfs/ server_ip 172.19.199.199 server_path /data/colobackup/ storage_space_total_gb 100.0 storagetype nfs`

2. Restore the backup information on a replacement CSP device using the **hostaction restore** command:

   For example,

   ```
   hostaction restore except-connectivity file-path
   nfs:nfsfs/WZP22180EW2_2020_06_24T18_07_00.bkup
   ```

Note Specify the except-connectivity parameter to retain the connectivity with the NFS server mounted in Step 2.

3. Use the **show hostaction backup status** command to view the status of the previous five backup images and their operational status.

   Also, you can view the backup images from the notifications available on the Cisco vManage **Monitor** > **Network** > **Events** page

4. Use the **show hostaction restore-status** command on the CSP device to view the status of the overall restore process and each component such as system, image and flavors, VM and so on.

5. To fix any failure after viewing the status, perform a factory default reset of the device.

Note The factory default reset sets the device to default configuration. Therefore, before performing the restore operation from Steps 1-4 on the replacement device, verify that all the restore operation prerequisites are met. See Prerequisites and Restrictions for Backup and Restore of CSP Devices, on page 58.

To know more about how to configure the restore operation on CSP devices, see Backup and Restore NFVIS and VM Configurations.

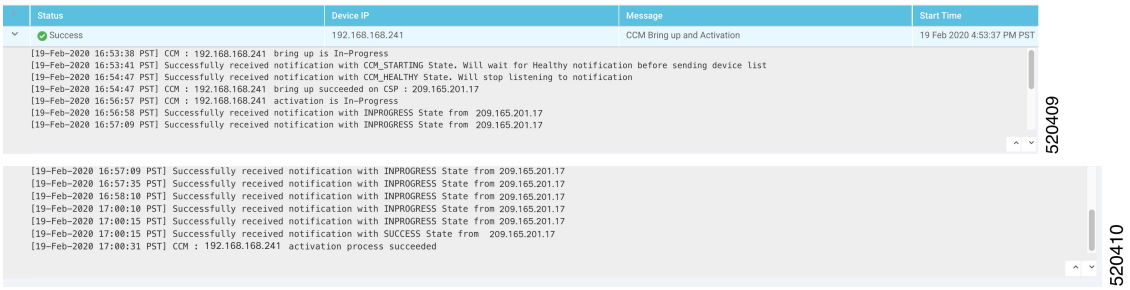
Progress of Cluster Activation

Table 12: Feature History

Feature Name	Release Information	Description
Monitor Cluster Activation Progress	Cisco SD-WAN Release 20.1.1	This feature displays the cluster activation progress at each step and shows any failures that may occur during the process. The process of activating a cluster takes approximately 30 minutes or longer, and you can monitor the progress using the vManage task view window and events from the Monitoring page.

To check cluster activation status after activating a cluster, view the progress on the task view window:

Figure 12: Cluster Activation



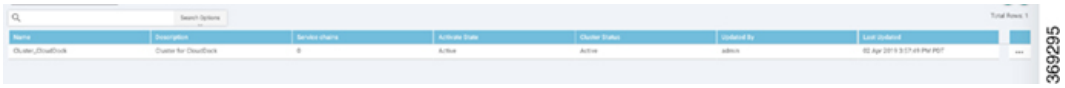
Perform the following verification steps:

1. To view cluster state and change the state:

   a. On the **Cluster** tab of Cloud OnRamp for Colocation screen, if a cluster goes into a "PENDING" state, click the **More Actions** icon, and click **Sync**. This action moves a cluster back to an "ACTIVE" state.

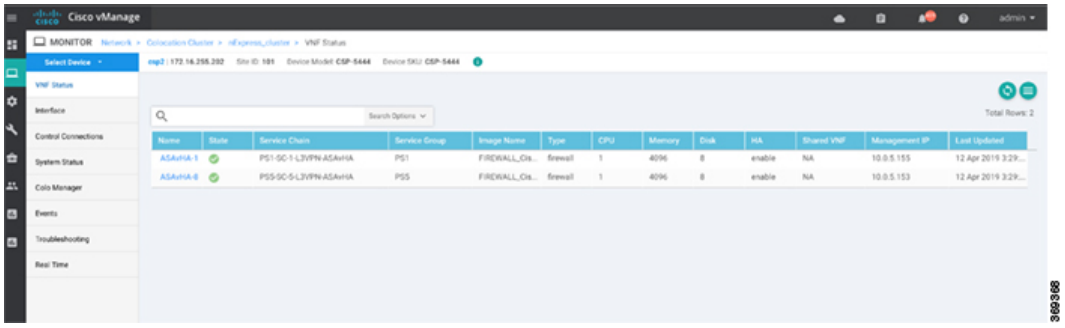
   b. To view if a cluster moves back to an "ACTIVE" state, you can view the successful activation on the **Cluster** tab.

   Figure 13: Cluster Tab



2. To view the service groups present on CSP devices, click **Monitor** > **Network** > **Colocation Cluster**.

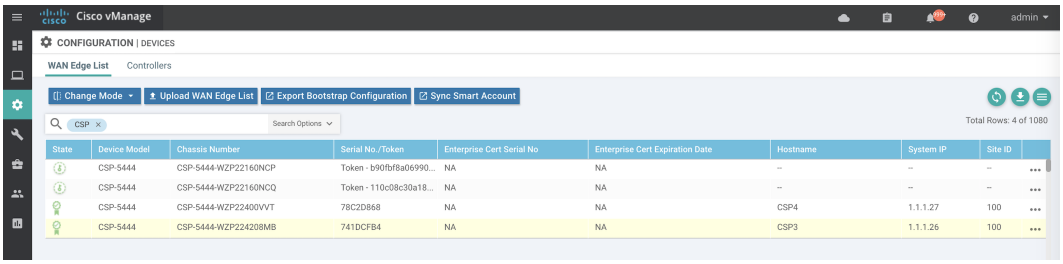
   Choose a cluster and then choose a CSP device as shown in the following image. You can choose and view other CSP devices.



3. To check if cluster has been activated from a CSP device:

   a. In Cisco vManage, click **Configuration** > **Devices**.

   b. View device status of all the CSP devices and ensure that they are in sync with Cisco vManage.



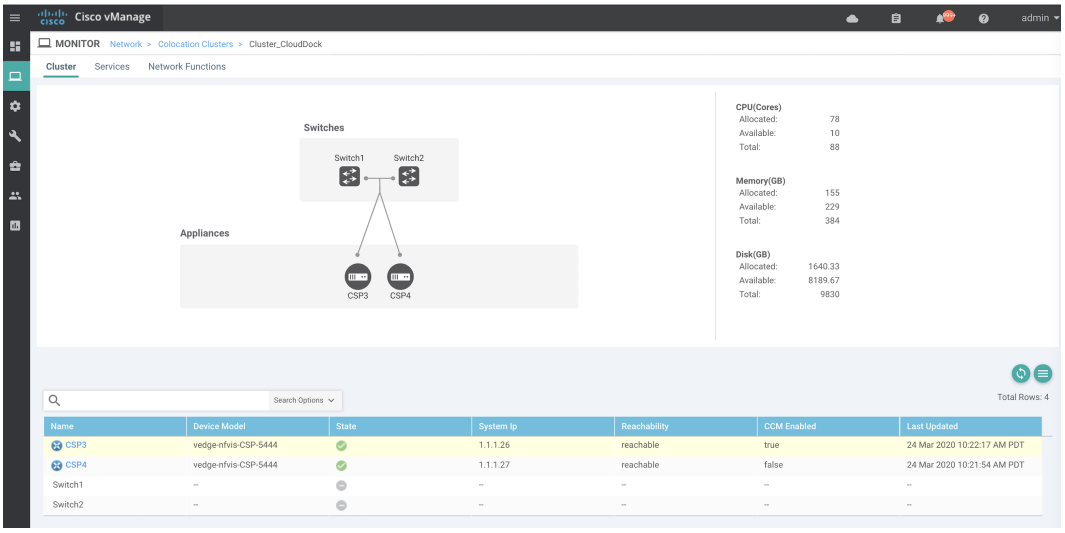
c. View the state of CSP devices and verify that the certificates are installed for CSP devices.

Note If the state of CSP devices does not show "cert installed" for more than five minutes after CSP activation through OTP, see Troubleshoot Cisco Cloud Services Platform Issues, on page 145 .

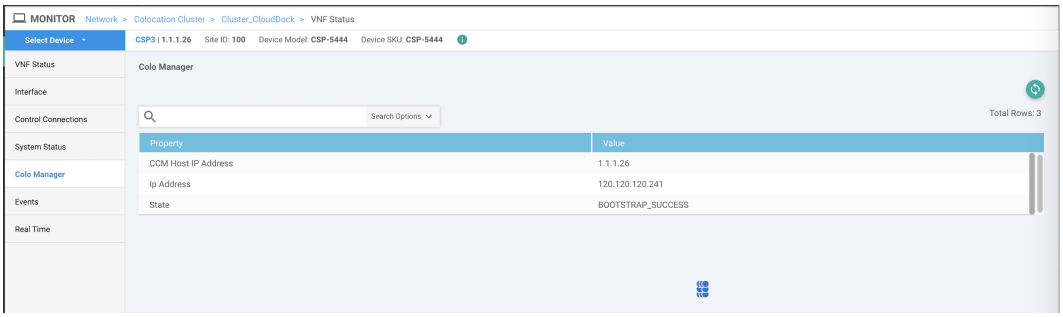
After cluster has been activated from CSP device, the Cisco Colo Manager (CCM) performs the cluster activation tasks on the Cisco NFVIS host.

4. To view if CCM is enabled for a CSP device,
   a. In Cisco vManage, click **Monitor** > **Network**.
   b. Click **Colocation Clusters**.

      In the detail part, view if CCM is enabled for specific CSP devices.



5. To monitor CCM health,
   a. In Cisco vManage, click **Monitor** > **Network**.
   b. To view CCM health, from the left pane, click **Colo Manager**.



If the Cisco Colo Manager status does not go to "HEALTHY" after "STARTING", see Troubleshoot Cisco Colo Manager Issues, on page 153 .

If the status of Cisco Colo Manager goes to "HEALTHY" after "STARTING" but the status of Cisco Colo Manager shows IN-PROGRESS for more than 20 minutes after the switch configurations are already complete, see Switch devices are not calling home to PNP or Cisco Colo Manager, on page 140 .

View Cluster from vManage

To view a cluster configuration, perform the following steps:

Procedure

Step 1 Log into Cisco vManage, at the URL HTTPS: //vManage-ip-address/, as a user with "admin" credentials.

Step 2 In vManage, choose **Configuration** > **Cloud OnRamp for Colocation**. The CLOUD ONRAMP FOR COLOCATION Cluster screen appears, and the **Configure & Provision Cluster** button is highlighted.

Step 3 In the **Cluster** tab, click a cluster, click the **More Actions** icon to the right of its row, and click **View** against the cluster.

The Cluster window opens, displaying the switches and CSP devices in the cluster and showing which cluster settings have been configured.

Step 4 You can only view the global parameters being set, configuration of switches and CSP devices.

Step 5 Click the **Cancel** button to return to the CLOUD ONRAMP FOR COLOCATION Cluster screen.

Edit Cluster in vManage

To modify any existing cluster configuration such as global parameters, perform the following steps:

Procedure

Step 1 Log into Cisco vManage, at the URL HTTPS: //vManage-ip-address/, as a user with "admin" credentials.

Step 2 In vManage, select **Configuration** > **Cloud OnRamp for Colocation**. The CLOUD ONRAMP FOR COLOCATION Cluster screen appears, and the **Configure & Provision Cluster** button is highlighted.

Step 3 In the **Cluster** tab, click a cluster, click the **More Actions** icon to the right of its row, and click **Edit** against the cluster.

The Cluster window opens, displaying the switches and CSP devices in the cluster and showing which cluster settings have been configured.

Step 4 In the cluster design window, you can modify some of the global parameters. Based on whether a cluster is in active or inactive state, you can perform the following operations on a cluster:

a. Inactive state:

- Edit all global parameters, and the Resource pool parameter.
- Add more CSP devices (up to eight).
- Cannot edit the name or serial number of a switch or CSP device. Instead, delete the CSP or switch and add another switch or CSP with a different name and serial number.
- Delete an entire cluster configuration.

b. Active state:

- Edit all global parameters, except the Resource pool parameter.

  Note
  The Resource pool parameter cannot be changed when the cluster is activated. However, the only way to change the Resource pool parameter is to delete the cluster and recreate it again with the correct Resource pool parameter.

- Cannot edit the name or serial number of a switch or CSP device. Instead, delete the CSP or switch and add another switch or CSP with a different name and serial number.
- Cannot delete a cluster in active state.
- Add more CSP devices (up to eight).

Step 5 Click the **Save Cluster** button.

Add CSP Device to Cluster

You can add and configure the CSP devices using Cisco vManage.

Before you begin

Ensure that the Cisco NFVIS version that you use is same for all the CSP devices in the cluster.

Procedure

Step 1 In Cisco vManage, choose **Configuration** > **Cloud OnRamp for Colocation**.

Step 2 In the **Cluster** tab, choose a cluster from the table, click the **More Actions** icon at the right of the row, and choose **Add/Delete CSP**.

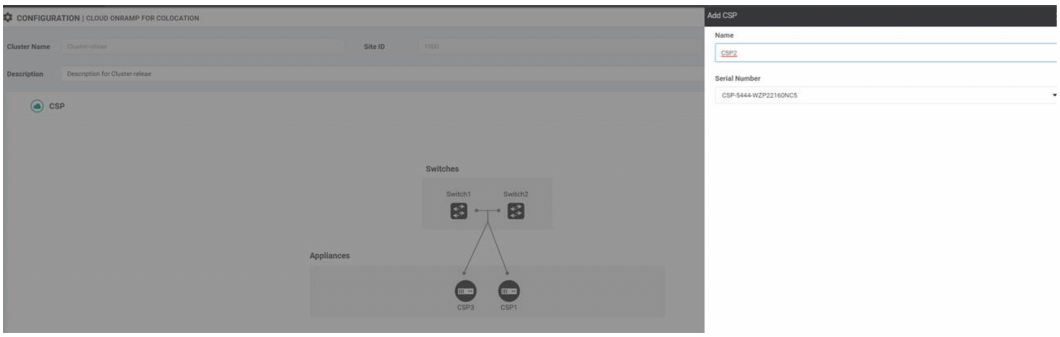
Step 3 To add a CSP device, click **+ Add CSP**. The **Add CSP** dialog box appears. Provide a name and choose the CSP device serial number. Click **Save**.

Step 4 To configure a CSP device, click the CSP icon in the CSP box. The **Edit CSP** dialog box appears. Provide a name and choose the CSP device serial number. Click **Save**.

The name can contain 128 alphanumeric characters.

Note
To bring up the CSP devices, ensure that you configure the OTP for the devices.

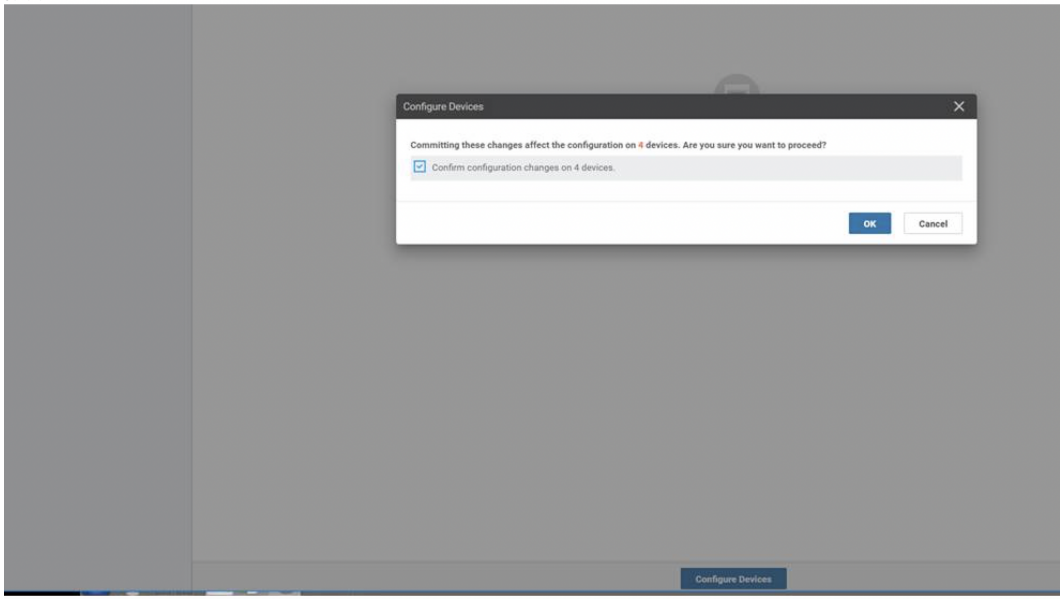
Figure 14: Add CSP



Step 5 Click **Save**.

Step 6 After saving, perform the onscreen configuration instructions as shown in the following images:





Step 7 To check if the CSP device is added, use the task view page that displays a list of all running tasks.

You can go to the Cluster or Network page, and verify if the CSP device is added.

Delete CSP Devices from Cluster

You can delete CSP devices using Cisco vManage.

Procedure

Step 1 In Cisco vManage, choose **Configuration** > **Cloud OnRamp for Colocation**.

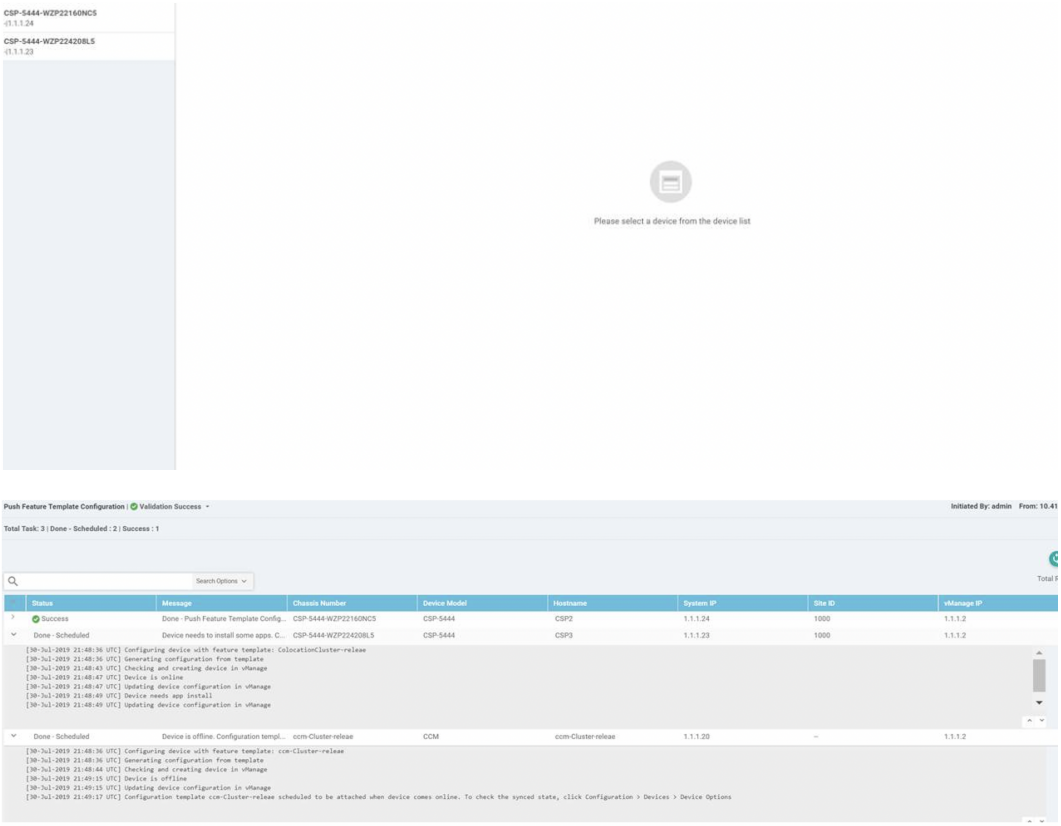
Step 2 In the **Cluster** tab, choose a cluster from the table, click the **More Actions** icon at the right of the row, and click **Add/Delete CSP**.

Step 3 To delete a CSP device, click a CSP icon from the **Appliances** box.

Step 4 Click **Delete**.

Step 5 Click **Save**.

Step 6 Perform the onscreen instructions to proceed with the deletion as shown in the following images.



Step 7 Reset the devices to factory-default settings. .

Step 8 To decommission invalid CSP devices in Cisco vManage, click **Configuration** > **Devices**.

Step 9 Locate the CSP devices that are in the deactivated cluster, click the **Add/Delete CSP**, and click **Decommission WAN Edge**.

This action provides new tokens to the devices.

If an HA service chain is deployed on a CSP device that is removed, the corresponding HA service chains are removed from the CSP device that hosts the HA instances.

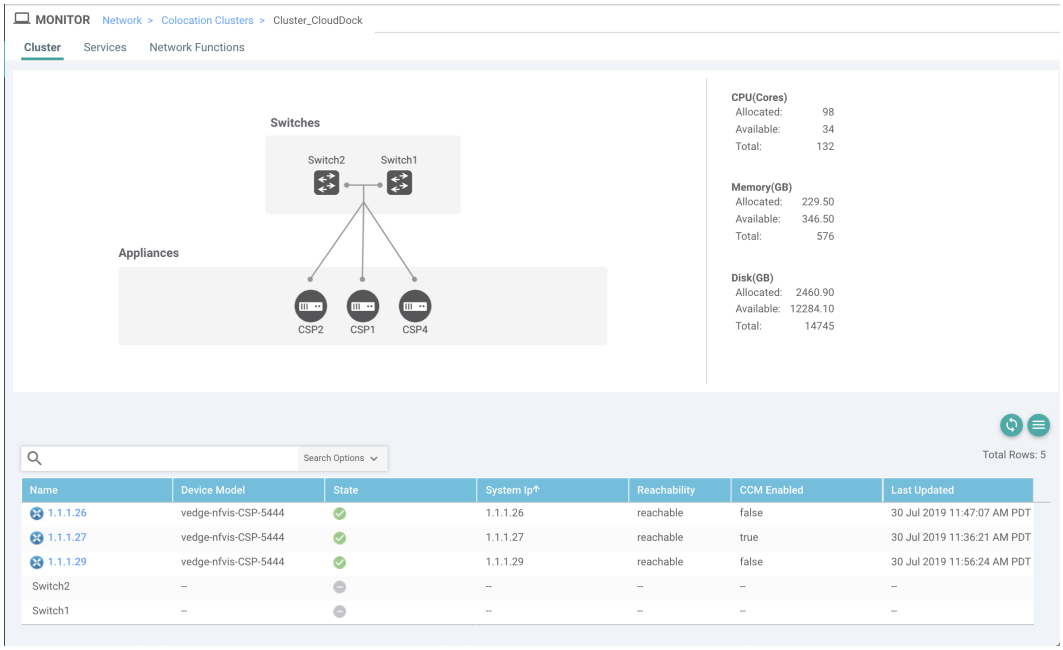
Delete CSP with CCM

Procedure

Step 1 Determine the CSP that hosts the CCM.

Step 2 If **CCM Enabled** is true on a CSP and you choose to delete the corresponding CSP, click the **More Actions** icon at the right of the row, and choose **Add/Delete CSP**. In the following image from the Montior screen, you can view the CCM enabled state from the detail part.

Figure 15: CSP Device with CCM



When the CSP that runs the service chain mornitoring service and CCM is removed from a cluster, ensure that you click the **Sync** button of the cluster. Clicking the sync button starts the service chain health monitoring service on a different CSP to continue monitoring the existing service chain health.

Ensure that Cisco vManage has control connections to all CSP devices so that it can bring up CCM instance on another CSP device.

After you delete a CSP device with CCM, when a service chain is configured and attached to this cluster, the CCM instance is brought up on one of the CSP devices on this cluster. The service chain monitoring is disabled until the CCM instance isn't brought up in any of the remaining CSP devices.

Replace Cisco CSP Devices After RMA

SUMMARY STEPS

1. In Cisco vManage , click **Configuration** > **Cloud OnRamp for Colocation**.
2. Click a cluster from the table, click the **More Actions** option, and choose **RMA**.
3. In the RMA dialog box:

DETAILED STEPS

Procedure

Step 1 In Cisco vManage , click **Configuration** > **Cloud OnRamp for Colocation**.

Step 2 Click a cluster from the table, click the **More Actions** option, and choose **RMA**.

Step 3 In the RMA dialog box:

a) Select Appliance: Choose a CSP device that you want to replace.

All CSP devices in a specific colocation cluster are displayed in the format, CSP Name-<Serial Number>.

b) Choose a serial number for a new CSP device from the drop-down list.

c) Click **Save**.

After saving, you can view the configuration.

Return of Materials of Cisco CSP Devices

Table 13: Feature History

Feature Name	Release Information	Description
RMA Support for Cisco CSP Devices	Cisco SD-WAN Release 20.5.1 Cisco vManage Release 20.5.1	This feature allows you to replace a faulty CSP device by creating backup copies of the device, and then restoring the replacement device to a state it was in before the replacement. The VMs running in HA mode operate uninterrupted with continuous flow of traffic during device replacement.

You can now create backup copies and restore NFVIS configurations and VMs.

Points to Consider

- You can use Network File Storage (NFS) servers to create regular backup copies of the CSP devices.
- If you're using an external NFS server for the backup operation, ensure that you maintain and clean the NFS directory regularly. This maintenance ensures that the NFS server has sufficient space for the incoming backup packages.
- If you don't use NFS servers, don't configure the backup server settings using Cisco vManage. However, if you're not configuring the backup server settings, you can't restore the replacement device. You can use delete CSP to remove the faulty device, add a new CSP device, and then start provisioning the service chains onto the added CSP device.

RMA Process for Cisco CSP Devices

Ensure that you perform the Return of Materials (RMA) process in the following order:

1. Create a backup copy of all the CSP devices in a cluster using Cisco vManage. See Backup Server Settings, on page 46.

Note During CSP device replacement, create a backup copy of the device in the NFS server when creating a cluster using Cisco vManage. Perform one of the following if you're bringing up a cluster or editing an existing cluster.

- Bring up a colocation cluster: At the time of cluster creation and activation, provide information about the NFS storage server and backup intervals. If the backup task fails on a CSP device, the device returns an error, but the cluster activation continues. Ensure that you update the cluster after addressing the failure and wait for a successful cluster activation.
- Edit a colocation cluster: For an existing active cluster, edit the cluster and provide information about the NFS storage server and backup intervals.

2. Contact Cisco Technical Support to get a replacement CSP device. See Cisco Cloud Services Platform 5000 Hardware Installation Guide for more information about replacing a CSP device.
3. Rewire the replacement Cisco CSP device with the Cisco Catalyst 9500 switches to move the wiring of the faulty device to the replacement device. See Wiring Requirements, on page 10.
4. Verify that the Cisco CSP ISO image running on the replacement device is the same that was running on the faulty device.
5. Restore the replacement device using CLI.

Prerequisites and Restrictions for Backup and Restore of CSP Devices

Prerequisites

Backup Operation

- The connectivity to the NFS server from CSP devices should be established before configuring the backup server settings using Cisco vManage.
- The backup directory on the NFS server should have write permission.
- The external NFS server should be available, reachable, and maintained. The maintenance of the external NFS server requires you to check the available storage space and network reachability regularly.
- The schedule for the backup operation should be synced with the local date and time on the CSP device.

Restore Operation

- The replacement device should have the same resources as the faulty device. These resources are, Cisco NFVIS image version, CPU, memory and storage as the faulty CSP device.
- The connectivity between the replacement device and switch ports should be same as the faulty device and switches.
- The PNIC wiring of the replacement device should match the faulty device on the Catalyst 9500 switches.

  For example,

  If slot-1/port-1 (eth1-1) on the faulty device is connected to switch-1 and port, 1/0/1, then connect slot-1/port-1 (eth1-1) of the replacement device to the same switch port, such as switch-1 and port, 1/0/1.

- The onboarding of the replacement device should be completed using the PnP process for CSP devices.
- To prevent the loss of backup access during the restore operation, the configuration for mounting an NFS server to access the backup package should match the configuration on the faulty device.

  You can view configuration information from other CSP devices as the NFS mount location and configurations are same for all the CSP devices. To view the active configuration that is running on a healthy CSP device, use the **show running-config** command. Use this active configuration information when creating a mount point during the restore operation.

  For example,

```
nfvis# show running-config mount
mount nfs-mount storage nfsfs/
storagetype            nfs
storage_space_total_gb 123.0
server_ip              172.19.199.199
server_path            /data/colobackup/
!
```

- The authentication of the replacement device with the Cisco SD-WAN controllers using the OTP process should be completed after restoring the replacement device.



Note Use the **request activate chassis-number** *chassis-serial-number* **token** *token-number* command to authenticate a device by logging in to Cisco NFVIS.

- The replacement device shouldn't have any configuration other than the configuration of the faulty device.

Restrictions

Backup Operation

- The periodic backup operation doesn't start during the upgrade of a CSP device.
- If the NFS folder path isn't available on the NFS server, the backup operation doesn't start.
- Only one backup operation can occur at a specific time.
- The backup operation fails if the available disk space on the NFS server is less than the combined size of the VM export size and tar.gz VM packages.
- The backup device information can only be restored on a replacement CSP device and not on any existing device that is already part of the cluster.
- The NFS mount configurations can’t be updated after they are configured for a CSP device. To update, delete the NFS configuration and reapply an updated configuration to the NFS server and reconfigure the backup schedule. Perform this update when the backup operation isn't in progress.

Restore Operation

- Only one restore operation can occur at a specific time.
- If a backup file doesn’t exist in the NFS server, the restore operation doesn't start.
- The restore operation isn't supported when you convert a cluster from a single tenant mode to multitenant mode, and conversely.

Remove PNF Devices from Cluster

Procedure

Step 1 Detach all service groups and service chains that has the PNF.

Step 2 (Optional) Delete the service groups.

If the deleted PNF is an ASR router, which is orchestrated through vManage, invalidate and decommission the device from the **Device** page.

Step 3 Remove the cables that connect the PNF with the Cisco Catalyst 9500-40X or Cisco Catalyst 9500-48Y4C switches and manually remove the VLAN configuration from the Cisco Catalyst 9500-40X or Cisco Catalyst 9500-48Y4C corresponding interfaces.

Remove Cluster from vManage

To decommission an entire cluster from vManage, perform the following steps:

Procedure

Step 1 Log into Cisco vManage, at the URL HTTPS: //vManage-ip-address/, as a user with "admin" credentials.

Step 2 In Cisco vManage, in the **Configuration** > **Certificates** screen, locate and verify status of devices to be deleted, and click **Invalid** against the devices.

Step 3 In the **Configuration|Ceritificates** screen, click **Send to Controllers**.

Step 4 In vManage, click **Configuration** > **Cloud OnRamp for Colocation**. The CLOUD ONRAMP FOR COLOCATION Cluster screen appears, and the **Configure & Provision Cluster** button is highlighted.

Step 5 In the **Cluster** tab, locate the cluster that has invalid devices, click the **More Actions** icon to the right of its row, and click **Deactivate** against the cluster.

If the cluster is attached to one or more service groups, you are prompted with a message that service chains hosting the VMs are running on this device and whether you can continue with the cluster deletion. However, although you confirm deletion of a cluster, you are not allowed to remove the cluster without detaching the service groups that are hosted on this device. If the cluster is not attached to any service group, you are prompted with a message to confirm the cluster deletion.

Note
You can delete the cluster, if necessary, or can keep it in deactivated state.

Step 6 To delete the cluster, select **Delete**.

Step 7 Click the **Cancel** button to return to the CLOUD ONRAMP FOR COLOCATION Cluster screen without deleting the cluster.

Step 8 To decommission invalid devices, in vManage, click **Configuration** > **Devices**.

Step 9 Locate the devices that are in the deactivated cluster, click the **More Actions** icon to the right of the device row, and click **Decommission WAN Edge**.

This action provides new tokens to your devices.

Step 10 Reset the devices to the factory default by using the command:

factory-default-reset all

Step 11 Log into NFVIS by using **admin** as the login name and **Admin123#** as the default password.

Step 12 Reset switch configuration and reboot switches. See Clean switches configuration and reset switches to factory defaults, on page 144 .

Remove and Replace Switch

The Cisco Catalyst 9500-40X or Cisco Catalyst 9500-48Y4C series of switches are used in the data path for switching traffic between the different VNF devices in a service chain. There are two switches that are stacked by using Stackwise Virtual (SVL) technology.

To achieve a redundant stack, the switches use a set of two stackwise virtual links (SV links) and one dual active detection (DAD link). For prescriptive connections on Cisco Catalyst 9500-40X, ports 38, 39 are SVL links and port 40 is the DAD link. For prescriptive connections on Cisco Catalyst 9500-48Y4C, ports 46, 47 are SVL links and port 48 is the DAD link.

In a stack, there are two switches in which one of the switches is active and the other is the standby. The control plane databases are synchronized between the switches. Each switch is assigned a switch number as part of the stack. The switches are numbered 1 and 2 in the current scenario. For more information on SVL redundancy, see High Availability Switch Configuration Guide.

Note In the case of a switch failure, ensure that you know the switch number that failed. This switch can be used to set up as the replacement.

To replace a switch in the stack:

Procedure

Step 1 On the switch 1 console, use the **show switch** command.

```
Switch#sh switch
Switch/Stack Mac Address : c4b3.6a70.f480 - Foreign Mac Address
Mac persistency wait time: Indefinite
                                             H/W   Current
Switch#   Role    Mac Address     Priority Version  State
-------------------------------------------------------------------------------
*1       Active   c4b3.6a71.0b00     1      V01     Ready
 2       Member   0000.0000.0000     0      V01     Removed
```

Note
Here, the switch number that is removed is two. This switch number is required when configuring the new switch.

Step 2 On the switch that replaces the failed unit, ensure that the switch number is one. This is achieved by using the **show switch** command again on the new unit.

```
Switch#sh switch
Switch/Stack Mac Address : 5486.bc78.c900 - Local Mac Address
Mac persistency wait time: Indefinite
                                             H/W   Current
Switch#   Role    Mac Address     Priority Version  State
------------------------------------------------------------------------------------
*1       Active   5486.bc78.c900     1      V01     Ready
```

Step 3 If the new switch is numbered two, ensure that you renumber it to 1 and then reload the switches by using the following commands:

```
Switch#sh switch
Switch/Stack Mac Address : 5486.bc78.c900 - Local Mac Address
Mac persistency wait time: Indefinite
                                              H/W   Current
Switch#   Role    Mac Address     Priority Version  State
------------------------------------------------------------------------------------
*2       Active   5486.bc78.c900     1      V01     Ready


Switch#switch 2 renumber 1
WARNING: Changing the switch number may result in a configuration change for that switch.  The
interface configuration associated with the old switch number will remain as a provisioned
configuration. New Switch Number will be effective after next reboot. Do you want to continue?[y/n]?
 [yes]:
Switch#reload

System configuration has been modified. Save? [yes/no]: no
Reload command is being issued on Active unit, this will reload the whole stack
Proceed with reload? [confirm]

Jun 17 19:41:01.793: %SYS-5-RELOAD: Reload requested by console. Reload Reason: Reload Command
```

Step 4 Connect the required cables for SVL; which are ports 38, 39, and 40 from the first Cisco Catalyst 9500-40X switch to the second switch.

Step 5 On the second switch, configure and save the configuration.

```
Switch(config)#
stackwise-virtual
domain 10
!
interface TenGigabitEthernet1/0/38
stackwise-virtual link 1
!
interface TenGigabitEthernet1/0/39
stackwise-virtual link 1
!
interface TenGigabitEthernet1/0/40
stackwise-virtual dual-active-detection
```

Step 6 Renumber the new unit to be the same as the one it is replacing, and then reload the box.

```
Switch#switch 1 renumber 2
WARNING: Changing the switch number may result in a configuration change for that switch.  The
interface configuration associated with the old switch number will remain as a provisioned
configuration. New Switch Number will be effective after next reboot. Do you want to continue?[y/n]?
 [yes]: yes
Switch# reload
```

After the new switch comes up, it joins the stack and synchronizes with the configuration.

```
Switch#sh switch
Switch/Stack Mac Address : c4b3.6a70.f480 - Foreign Mac Address
Mac persistency wait time: Indefinite
                                           H/W   Current
Switch#   Role    Mac Address     Priority Version  State
-------------------------------------------------------------------------------------
*1       Active   c4b3.6a71.0b00     1      V01     Ready
 2       Member   5486.bc78.c900     1      V01     Ready



Switch#
*Jun 17 21:00:57.696: %IOSXE_REDUNDANCY-6-PEER: Active detected switch 2 as standby.
*Jun 17 21:00:57.694: %STACKMGR-6-STANDBY_ELECTED: Switch 1 R0/0: stack_mgr: Switch 2 has
been elected STANDBY.
*Jun 17 21:01:02.651: %REDUNDANCY-5-PEER_MONITOR_EVENT: Active detected a standby insertion
 (raw-event=PEER_FOUND(4))

*Jun 17 21:01:02.651: %REDUNDANCY-5-PEER_MONITOR_EVENT: Active detected a standby insertion
 (raw-event=PEER_REDUNDANCY_STATE_CHANGE(5))

*Jun 17 21:01:53.686: %HA_CONFIG_SYNC-6-BULK_CFGSYNC_SUCCEED: Bulk Sync succeeded
*Jun 17 21:01:54.688: %RF-5-RF_TERMINAL_STATE: Terminal state reached for (SSO)
Switch#
Switch#
Switch#
Switch#
Switch#
Switch#
Switch#
Switch#sh switch
Switch/Stack Mac Address : c4b3.6a70.f480 - Foreign Mac Address
Mac persistency wait time: Indefinite
                                           H/W   Current
Switch#   Role    Mac Address     Priority Version  State
-------------------------------------------------------------------------------------
*1       Active   c4b3.6a71.0b00     1      V01     Ready
 2       Standby  5486.bc78.c900     1      V01     Ready
```

Reactivate Cluster from vManage

To add new CSP devices or when CSP devices are considered for RMA process, perform the following steps:

Procedure

Step 1 Log into Cisco vManage, at the URL HTTPS: //vManage-ip-address/, as a user with "admin" credentials.

Step 2 In Cisco vManage, in the **Configuration** > **Devices** screen, locate the devices that are in the deactivated cluster.

Step 3 Get new token from vManage for the devices.

Step 4 Log into NFVIS by using **admin** as the login name and **Admin123#** as the default password.

Step 5 Use the **request activate chassis-number** *chassis-serial-number* **token** *token-number* command.

Step 6 From vManage, configure the system configuration and then activate the cluster. See Create and Activate Clusters, on page 38.

If the cluster has been deleted, recreate and then activate it.

Step 7 In Cisco vManage, in the **Configuration** > **Certificates** screen, locate, and verify status of devices.

Step 8 To validate the devices, click **Valid** if it is invalid.

Step 9 In the **Configuration|Ceritificates** screen, click **Send to Controllers**.

Manage Service Groups

A service group consists of one or more service chains. You can configure a service group through vManage. A service chain is the structure of a network service, and consists of a set of linked network functions.

VNF Placement for Service Chains in vManage

The service chain placement component chooses a CSP device that hosts each VNF in service chains. The placement decision is based on available bandwidth, redundancy and computation resources (CPUs, memory, and storage) availability. The placement logic returns an error if the bandwidth, CPU, memory, and storage needs of all the VNFs in the service chains that are configured for a Cloud OnRamp for Colocation are not met. You are notified about the resources not being available and service chains are not deployed.

Create Service Chain in a Service Group

A service group consists of one or more service chains.

Table 14: Feature History

Feature Name	Release Information	Feature Description
Monitor Service Chain Health	Cisco SD-WAN Release 19.2.1	This feature lets you configure periodic checks on the service chain data path and reports the overall status. To enable service chain health monitoring, NFVIS version 3.12.1 or later should be installed on all CSP devices in a cluster.

Procedure

Step 1 Log in to Cisco vManage with admin credentials at the URL HTTPS: //vManage-ip-address/.

Step 2 Click **Configuration** > **Cloud OnRamp for Colocation**.

a) Click the **Service Group** tab, and then click the **Create Service Group** button. Provide the service group name, description, and colocation group.

The service group name can contain 128 alphanumeric characters.

The service group description can contain 2048 alphanumeric characters.

For a multitenant cluster, choose a colocation group or a tenant from the drop-down list. For a single-tenant cluster, the colocation group **admin** is chosen by default.

b) Click **Add Service Chain**.

c) In the Add Service Chain dialog box, provide the following information:

Table 15: Add Service Chain Information

Field	Description
Name	The service chain name can contain 128 alphanumeric characters.
Description	The service chain description can contain alphanumeric 2048 characters.
Bandwidth	The service chain bandwidth is in MBPS. The default bandwidth is 10 MB and you can configure a maximum bandwidth of 5G.
Input Handoff VLANS and Output Handoff VLANS	The Input VLAN handoff and output VLAN handoff can be comma-separated values (10, 20), or a range 10–20.
Monitoring	A toggle button that allows you to enable or disable service chain health monitoring. The service chain health monitoring is a periodic monitoring service that checks health of a service chain data path and reports the overall service chain health status. By default, the monitoring service is disabled. A service chain with subinterfaces such as, SCHM (Service Chain Health Monitoring Service) can only monitor the service chain including the first VLAN from the subinterface VLAN list. The service chain monitoring reports status based on end-to-end connectivity. Therefore, ensure that you take care of the routing and return traffic path, with attention to the Cisco SD-WAN service chains for better results. **Note** • Ensure that you provide input and output monitoring IP addresses from input and output handoff subnets. However, if the first and last VNF devices are VPN terminated, you don't need to provide input and output monitoring IP addresses. For example, if the network function isn't VPN terminated, the input monitoring IP can be 192.0.2.1/24 from the inbound subnet, 192.0.2.0/24. The inbound subnet connects to the first network function and the output monitoring IP can be, 203.0.113.11/24 that comes from outbound subnet, 203.0.113.0/24 of the last network function of a service chain. • If the first or last VNF firewall in a service chain is in transparent mode, you can't monitor these service chains.
Service Chain	Choose a topology from the service chain drop-down list. For a service chain topology, you can choose any of the four validated service chains such as, Router - Firewall - Router, Firewall, Firewall - Router. See You can also create a customized service chain. See Create Custom Service Chain, on page 73.

d) In the Add Service Chain definition box, click **Add**.
Based on the service chain configuration information, a graphical representation of the service group with all the service chains and the VNFs automatically appear in the design view window. A VNF or PNF appears with a "V" or "P" around its circumference specifying that it is a virtual network function or a physical network function. It shows all the configured service chains within each service group. A check mark next to the service chain indicates that all configuration information for the service chain has been completed.

After you activate a cluster is activated, attach it with the service group, and enable monitoring service for the service chain, when you bring up the CSP device where CCM is running, Cisco vManage chooses the same CSP device to start the monitoring service. The monitoring service monitors all service chains periodically in a round robin fashion by setting the monitoring interval to 30 minutes. See Monitor Cloud OnRamp Colocation Clusters from Cisco vManage , on page 113.

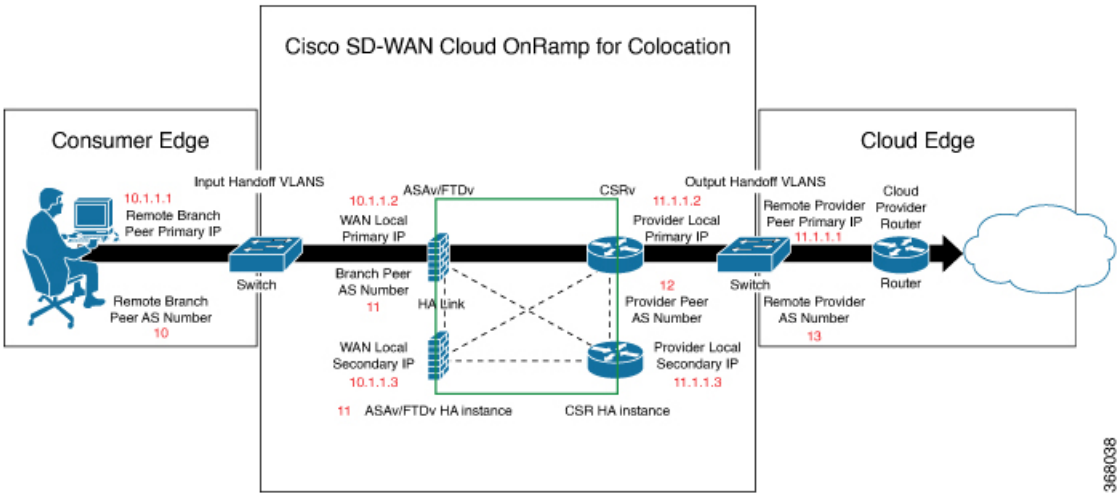
e) In the design view window, to configure a VNF, click a VNF in the service chain.
The Configure VNF dialog box appears.

f) Configure the VNF with the following information and perform the actions, as appropriate:

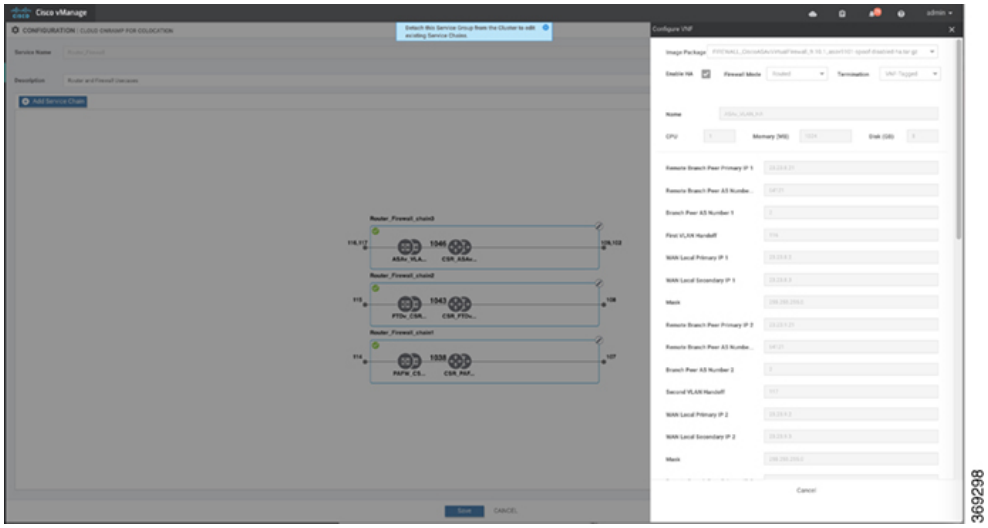
Table 16: VNF Properties of Router and Firewall

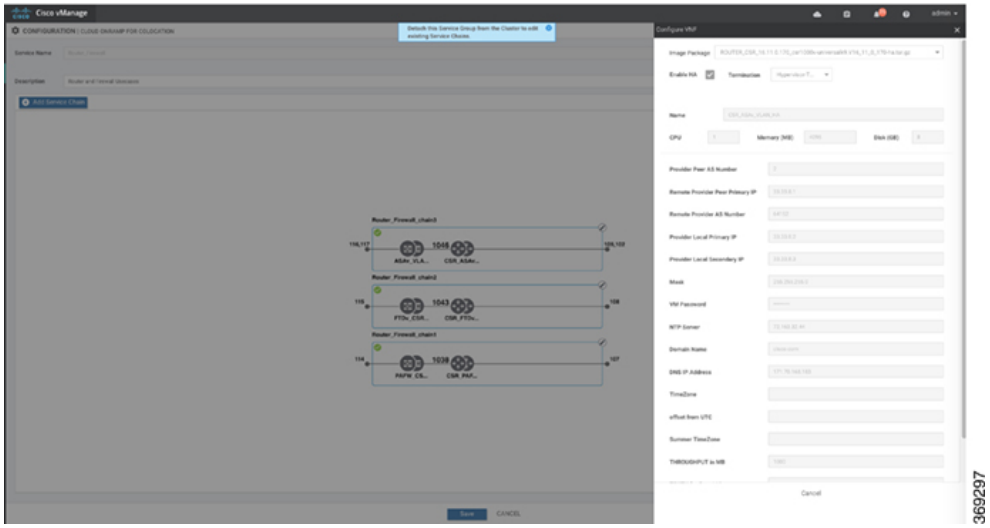
Field	Mandatory or Optional	Description
Image Package	Mandatory	Choose a router, firewall package.
Click **Fetch VNF Properties**. The available information for the image package is displayed in the **Configure VNF** dialog box.		
Name	Mandatory	VNF image name
CPU	Optional If you don't enter, the default value is considered, which is 1VCPU.	Specifies the number of virtual CPUs that are required for a VNF.
Memory	Optional If you don't enter, the default value is considered, which is 1024 MB.	Specifies the maximum primary memory in MB that the VNF can use.
Disk	Optional If you don't enter, the default value is considered, which is 8 GB.	Specifies disk in GB required for the VM.
You're prompted with any custom tokenized variables from Day-0 that requires your input. Provide the values.		

In the following image, all IP addresses, VLAN, and autonomous system within the green box are system-specific information that is generated from the VLAN, IP pools provided for the cluster. The information is automatically populated into Day-0 configurations of VMs.



The following images provide an example of the configuration for VNF IP addresses and autonomous system numbers in Cisco vManage.





If you're using a multitenant cluster and a comanged scenario, configure the Cisco SD-WAN VM by entering the values for the following fields and the remaining fields, as per the service chain design:

Note
To join the tenant overlay network, the provider should provide correct values for the following fields.

Field	Description
Serial Number	Specifies the authorized serial number of a Cisco SD-WAN device. The service provider can get the device serial number from the tenant before creating the service chain.
OTP	Specifies the OTP of the Cisco SD-WAN device that is available after authenticating it with Cisco SD-WAN Controllers. The service provider can get the OTP for the corresponding serial number from the tenant before creating the service chain.
Site Id	Specifies the identifier of the site in the tenant Cisco SD-WAN overlay network domain in which the Cisco SD-WAN device resides, such as a branch, campus, or data center. The service provider can get the site Id from the tenant before creating the service chain.
Tenant ORG Name	Specifies the tenant organization name that is included in the Certificate Signing Request (CSR). The service provider can get the organization name from the tenant before creating the service chain.
System IP connect to Tenant	Specifies the IP address to connect to the tenant overlay network. The service provider can get the IP address from the tenant before creating the service chain.
Tenant vBond IP	Specifies the IP address of the tenant Cisco vBond Orchestrator. The service provider can get the Cisco vBond Orchestrator IP address from the tenant before creating the service chain.

For edge VMs such as first and last VM in a service chain, you must provide the following addresses as they peer with a branch and the provider.

Table 17: VNF Options for First VM in Service Chain

Field	Mandatory or Optional	Description
Firewall Mode	Mandatory	Choose Routed or Transparent mode. **Note** Firewall mode is applicable to firewall VMs only.
Enable HA	Optional	Enable HA mode for the VNF.

Termination	Mandatory	Choose one of the following modes:
		• L3 mode selection with subinterfaces that are trunked
		`<type>selection</type> <val help="L3 Mode With Sub-interfaces(Trunked)" display="VNF-Tagged">vlan</val>`
		• L3 mode with IPSEC termination from a consumer-side and rerouted to the provider gateway
		`<val help="L3 Mode With IPSEC Termination From Consumer and Routed to Provider GW" display="Tunneled">vpn</val>`
		• L3 mode with access mode (nontrunked)
		`<val help="L3 Mode In Access Mode (Non-Trunked)" display="Hypervisor-Tagged">routed</val>`

g) Click **Configure**. The service chain is configured with the VNF configuration.
h) To add another service chain, repeat the procedure from Step b.
i) Click **Save**.

The new service group appears in a table under the **Service Group** tab. To view the status of the service chains that are monitored, use the task view page, which displays a list of all running tasks along with the total number of successes and failures. On the CSP device that has service chain health monotioring enabled, to determine the service chain health status, use the **show system:system status** command.

QoS on Service Chains

Table 18: Feature History

Feature Name	Release Information	Description
QoS on Service Chains	Cisco SD-WAN Release 20.1.1	This feature classifies the network traffic based on the Layer 2 virtual local-area network (VLAN) identification number. The QoS policy allows you to limit the bandwidth available for each service chain by applying traffic policing on bidirectional traffic. The bidirectional traffic is the ingress side that connects Cisco Catalyst 9500-40X switches to the consumer and egress side that connects to the provider.

Prerequisites

- Ensure that you use the Quality of Service (QoS) traffic policing on service chains that do not have shared VNF and PNF devices.

Note You cannot apply QoS policy on service chains with shared VNF devices where input and output VLANs are same for multiple service chains.

- Ensure that you use the following versions of software for QoS traffic policing:

Software	Release
Cisco NFVIS Cloud OnRamp for Colocation	4.1.1 and later
Catalyst 9500-40X	16.12.1 and later

The QoS policing policy is applied on the network traffic based on the following workflow:

1. Cisco vManage saves the bandwidth, input, or output VLAN information to VNF and PNF devices. To provide bandwidth and VLAN information, see Create Service Chain in a Service Group, on page 64.
2. CCM saves the bandwidth, input, or output VLAN values information to the Cisco Catalyst 9500-40X or Cisco Catalyst 9500-48Y4C switches.
3. CCM creates corresponding class-maps and policy-maps in Cisco Catalyst 9500-40X or Cisco Catalyst 9500-48Y4C switches based on VLAN match criteria.
4. CCM applies input service-policy on the ingress and egress ports.

Clone Service Groups

Table 19: Feature History

Feature Name	Release Information	Description
Clone Service Groups in Cisco vManage	Cisco SD-WAN Release 20.5.1 Cisco vManage Release 20.5.1	This feature allows you to create copies of service groups for different RBAC users, without having to enter the same configuration information multiple times. By cloning a service group, you can easily create service chains by leveraging the stored service chain templates.

When you clone or create copies of service chains, remember the following:

- Cisco vManage copies all configuration information of a service group to a cloned service group regardless of whether the cloned service group is attached to a cluster.
- Verify the CSV file and ensure that configuration information has a matching service group name during CSV file upload. Otherwise, an unmatched service group name can result in an error message during CSV file upload.
- To get an updated list of service group configuration values, always download service group configuration properties from the service group design view.

Procedure

Step 1 In Cisco vManage, click **Configuration** > **Cloud OnRamp for Colocation**.

Step 2 Click **Service Group**.
The service group configuration page appears and all the service groups are displayed.

Step 3 From the table, choose a service group, and then click **More Actions** to the right of the service group.

Step 4 Choose **Clone Service Group**.

A clone of the original service group appears in the service group design view. Note the following points:

- By default, the cloned service group name and VM names are suffixed with a unique string.
- To view any VM configuration, click a VM in service chains.
- Cisco vManage marks the service chains that require configuration as **Unconfigured**, next to the edit button of the service chain.

Step 5 Modify the service group name, if required. Provide a description for the service group.

Step 6 To configure a service chain, use one of the following methods:

- Click the edit button for a service chain, enter the values, and then click **Save**.
- Download the configuration values from a CSV file, modify the values, upload the file, and then click **Save**. See Steps 6, 7, 8 on how to download, modify, and upload a CSV file.

The cloned service group appears on the service group configuration page. You can now download the updated service group configuration values.

Step 7 To download the cloned service group configuration values, do one of the following:

Note
The download and upload of a CSV file is supported for creating, editing, and cloning of the service groups that aren't attached to a cluster.

- On the service group configuration page, click a cloned service group, click **More Actions** to the right of the service group, and choose **Download Properties (CSV)**.
- In the service group design view, click **Download CSV** in the upper right corner of the screen.

Cisco vManage downloads all configuration values of the service group to an Excel file in CSV format. The CSV file can consist of multiple service groups and each row represents configuration values for one service group. To add more rows to the CSV file, copy service group configuration values from existing CSV files and paste them in this file.

For example, ServiceGroup1_Clone1 that has two service chains with one VM in each of the service chains is represented in a single row.

Note
In the Excel file, the headers and their representation in the service chain design view is as follows:

- sc1/name represents the name of the first service chain.
- sc1/vm1/name represents the name of the first VNF in the first service chain.
- sc2/name represents the name of the second service chain.
- sc2/vm2/name represents the name of the second VNF in the second service chain.

Step 8 To modify service group configuration values, do one of the following:

- To modify the service group configuration in the design view, click a cloned service group from the service group configuration page.

  Click any VM in service chains to modify the configuration values, and then click **Save**.
- To modify the service group configuration using the downloaded Excel file, enter the configuration values in the Excel file manually. Save the Excel file in CSV format.

Step 9 To upload a CSV file that includes all the configuration values of a service group, click a service group in the service group configuration page, and then click **Upload CSV** from the right corner of the screen.

Click **Browse** to choose a CSV file, and then click **Upload**.

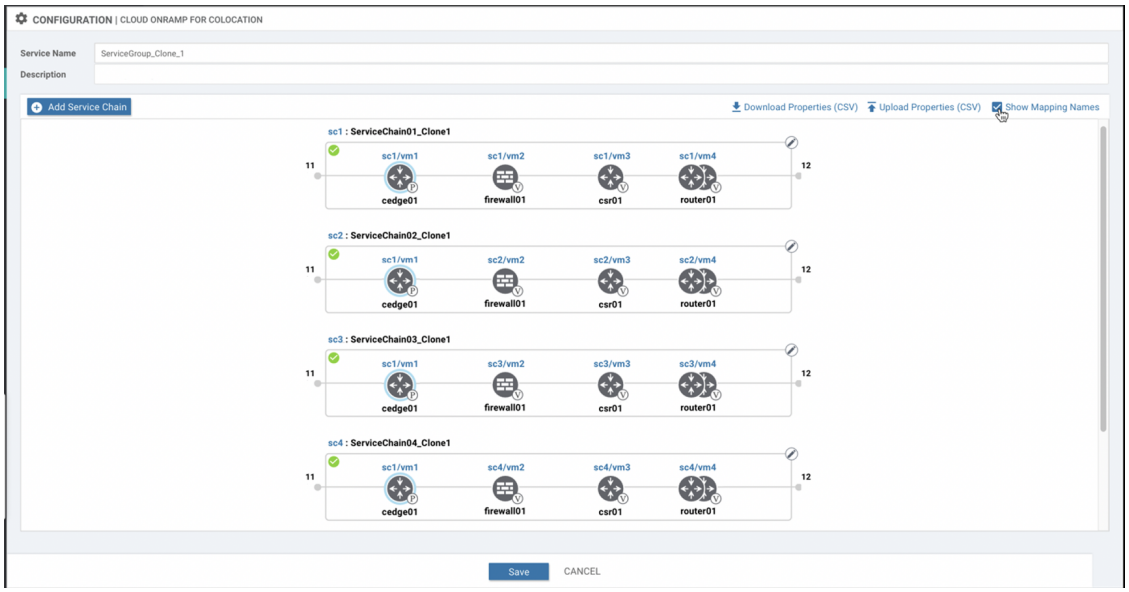
You can view the updated values displayed for the service group configuration.

Note
You can use the same CSV file to add configuration values for multiple service groups. But, you can update configuration values for a specific service group only, when uploading a CSV file using Cisco vManage.

Step 10 To know the representation of service group configuration properties in the CSV file and Cisco vManage design view, click a service group from the service group configuration page.

Click **Show Mapping Names**.

A text appears next to all the VMs in the service chains. Cisco vManage displays this text after mapping it with the configuration properties in the CSV file.



Create Custom Service Chain

You can customize service chains,

- By including extra VNFs or add other VNF types.
- By creating new VNF sequence that is not part of the predefined service chains.

Procedure

Step 1 Create a service group and service chains within the service group. See Create Service Chain in a Service Group, on page 64.

Step 2 In the **Add Service Chain** dialog box, provide the service chain name, description, bandwidth, input VLAN handoff, output VLAN handoff, monitoring health information of a service chain, and service chain configuration. Click **Add**.

For the service chain configuration, choose **Create Custom** from the drop-down. An empty service chain in the design view window is available.

Step 3 To add a VNF such as a router, load balancer, firewall, and others, click a VNF icon on the left panel, and drag the icon to its proper location within the service group box. After adding all required VNFs and forming the VNF service chain, configure each of the VNFs. Click a VNF in the service group box. The Configure VNF dialog box appears. Enter the following parameters:

a) Choose the software image to load from the **Image Package** drop-down.
b) Click **Fetch VNF Properties**.
c) Enter a name of the VNF in the **Name** field.
d) Enter the number of virtual CPUs required for the VNF in the **CPU** field.
e) Enter the amount of memory in megabytes to be allocated for the VNF in the **Memory** field.

f) Enter the amount of memory for storage in gigabytes to be allocated for the VNF in the **Disk** field.
g) Enter VNF-specific parameters, as required.

Note
These VNF details are the custom variables that are required for Day-0 operations of the VNF.

h) Click **Configure**.
i) To delete the VNF or cancel the VNF configuration, click **Delete** or **Cancel** respectively.

The customized service chains are added to a service group.



Note You can customize a VNF sequence with only up to four VNFs in a service chain.

Physical Network Function Workflow

This topic outlines the sequence of operations that you require to create shared PNF devices, configure, and monitor them. To ensure that the PNF workflow is effective, ensure that cabling is accurate, and VLAN ports are on the right ports of Cisco Catalyst 9500-40X or Cisco Catalyst 9500-48Y4C.

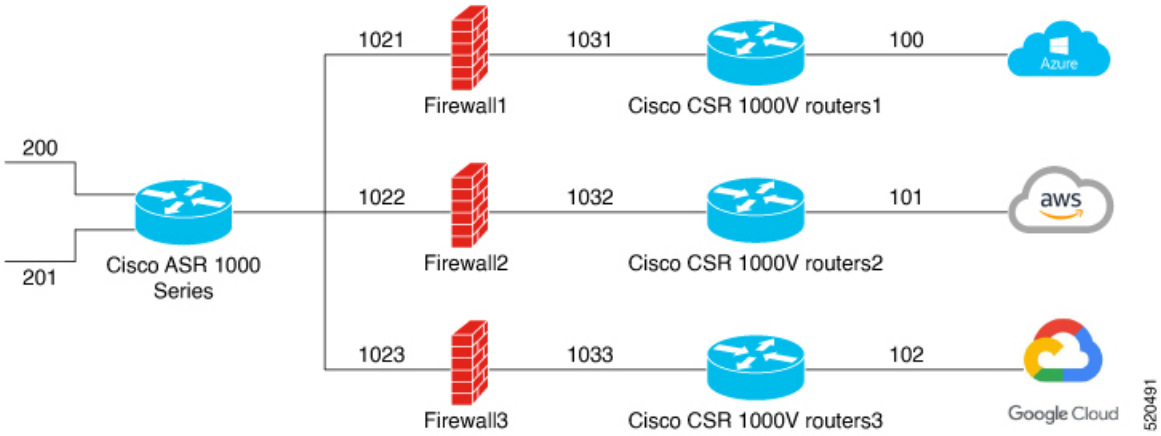
1. Connect the PNF devices to Cisco Catalyst 9500-40X or Cisco Catalyst 9500-48Y4C switch devices.
2. To make Cisco ASR 1000 Series router managed by Cisco vManage, upload WAN edge router authorized serial numbers from the Cisco Smart Account. See the "Upload WAN Edge Router Serial Numbers from Cisco Smart Account" section in the System and Interfaces Configuration Guide.
3. Create service chains by using the added PNF devices. See Custom Service Chain with Shared PNF Devices, on page 75.
4. Attach the service group to a cluster and check the configuration parameters that are generated. See Attach and Detach Service Group with Cluster, on page 87.
5. Configure the PNF and the Cisco Catalyst 9500-40X or Cisco Catalyst 9500-48Y4C switch devices according to the configuration parameters generated. See Configure PNF and Cisco Catalyst 9500 Switches, on page 78.

In the following image, the first PNF is shared with multiple service chains. These service chains access different cloud applications in Microsoft Azure, AWS, and Google Cloud. The traffic from VLAN 200 enters the Cisco ASR 1000 series PNF based on SD-WAN policy definition and fetches the next hop firewall based on VRF configuration and corresponding destination application. The return traffic should traverse the same path for each application traffic.

To configure the PNF,

1. Log into the ASR1000 Series device, and configure it based on the VLAN and IP address information available from Cisco vManage.
2. To allow specific VLANs on both inbound and outbound traffic, configure the Cisco Catalyst 9500-40X or Cisco Catalyst 9500-48Y4C switch ports where the PNF devices are connected.

Figure 16: PNF Shared with Multiple Service Chains



Custom Service Chain with Shared PNF Devices

You can customize service chains by including supported PNF devices.

Caution Ensure that you do not share PNF devices across clusters. A PNF device can be shared across service chains, or across service groups. However, a PNF device can now be shared only across a single cluster.

Table 20: Feature History

Feature Name	Release Information	Feature Description
Manage PNF Devices in Service Chains	Cisco SD-WAN Release 19.2.1	This feature lets you add Physical Network Function (PNF) devices to a network, in addition to the Virtual Network function (VNF) devices. These PNF devices can be added to service chains and shared across service chains, service groups, and a cluster. Inclusion of PNF devices in the service chain can overcome the performance and scaling issues caused by using only VNF devices in a service chain.

Before you begin

To create a customized service chain by adding a router or firewall to an existing service chain, perform the following steps:

- If a PNF device needs to be managed by vManage, ensure that the serial number is already available in the vManage, which can then be available for selection during PNF configuration.
- The FTD device can be in any position in a service chain.
- An ASR 1000 Series Aggregation Services Routers can only be in the first and last position in a service chain.
- You can add PNF devices across service chains and service groups.

- You can share PNF devices across service groups. They can be shared across service groups by entering the same serial numbers.
- You can share PNF devices across a single cluster and cannot share across multiple clusters.

Procedure

Step 1 Create a service group and service chains within the service group. See Create Service Chain in a Service Group, on page 64.

Step 2 In the **Add Service Chain** dialog box, provide the service chain name, description, bandwidth, input VLAN handoff, output VLAN handoff, monitoring health information of a service chain, and service chain configuration. Click **Add**.

For the service chain configuration, choose **Create Custom** from the drop-down. An empty service chain in the design view window is available. In the left panel, the set of VNF devices and PNF devices that you can add into the service chain appears. The 'V' in the circumference of VNF devicess represents a VNF and 'P' in the circumference of PNF devices represent a PNF.

Note
Ensure that you choose the **Create Custom** option for creating service chains by sharing PNF devices.

Step 3 To add a PNF such as physical routers, physical firewalls in a service chain, click the required PNF icon on the left panel, and drag the icon to its proper location within the service chain box.

After adding all required PNF devices, configure each of them.

a) Click a PNF device in the service chain box.

The Configure PNF dialog box appears. To configure a PNF, enter the following parameters:

b) Check **HA Enabled** if HA is enabled for the PNF device.

c) If the PNF is HA enabled, ensure that you include the HA serial number in **HA Serial**.

If the PNF device is FTD, enter the following information.

1. Enter a name of the PNF in the **Name** field.
2. Choose Routed or Transparent mode as the **Firewall Mode**.
3. Enter the serial number of the PNF device in the **PNF Serial** field.

If the PNF device is ASR 1000 Series Aggregation Services Routers, enter the following information.

1. Check **vManaged** if the device is managed by vManage.
2. Click **Fetch Properties**.
3. Enter a name of the PNF in the **Name** field.
4. Enter the serial number of the PNF device in the **PNF Serial** field.

d) Click **Configure**.

Step 4 To add service chains and share PNF devices, repeat from step 2.

Step 5 Edit an existing PNF configuration by clicking it.

Step 6 In **Share NF To**, choose the service chains with which the PNF should be shared.

After a PNF is shared, if you hover on a PNF, the respective shared PNF devices are highlighted in blue color. However, the PNFs from different service groups are not highlighted in blue color. After you choose a NF to be shared, a blue color rim appears on it. If the same PNF is shared across multiple service chains, it can be used in different positions by dragging and placing the PNF icons in a specific positon.

Figure 17: Single PNF in a Service Chain

Here a service chain consists of a single PNF, Ftd_Pnf (not shared with other service chains).

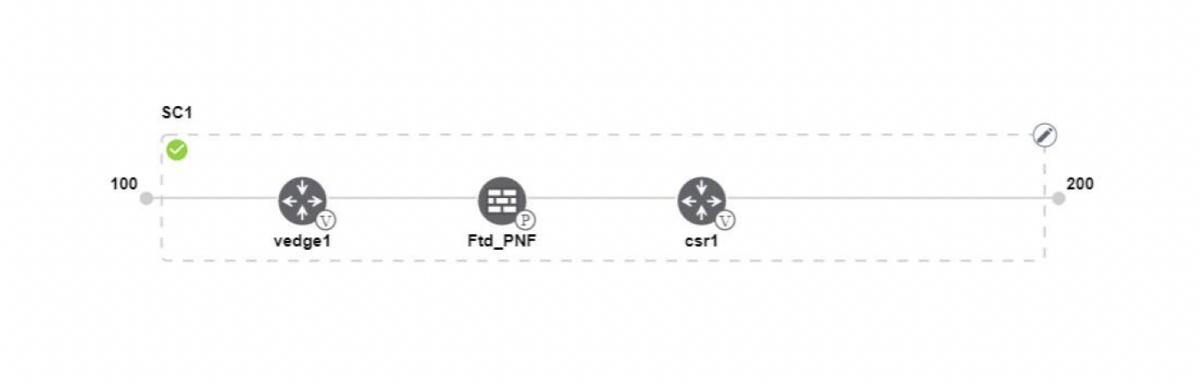


Figure 18: Two PNF Devices in Service Chains

Here service chains consist of two PNFs, FTdv_PNF shared across SC1 and SC2 and ASR_PNF (non shred).

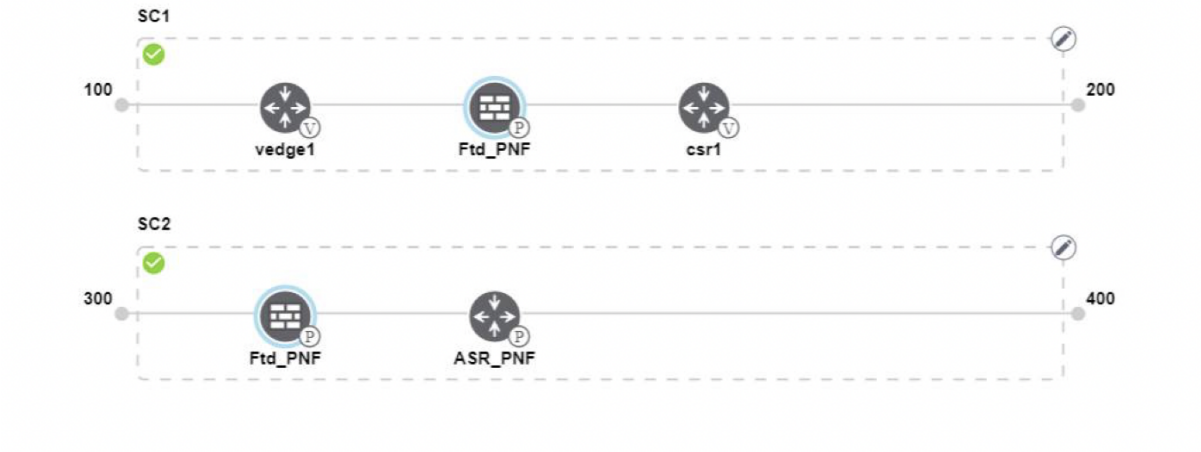
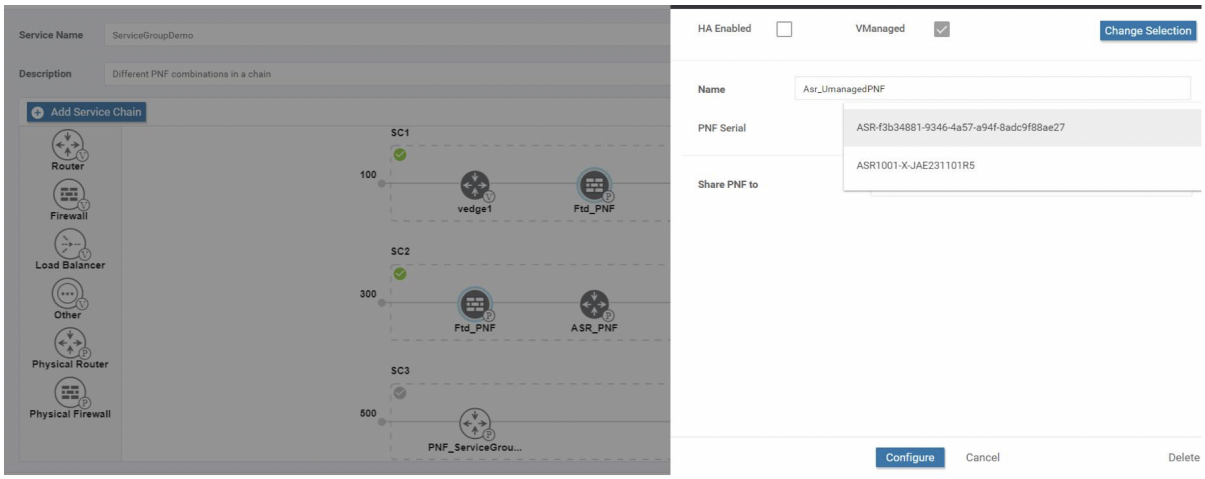


Figure 19: Three PNF Devices in Service Chains

Here service chains consist of three PNF devices in two different positions along with the vManage configuration.



Step 7 To delete a NF or cancel the NF configuration, click **Delete** or **Cancel** respectively.

You must attach the service groups to a cluster. After attaching service groups containing PNF devices with a cluster, the PNF configuration is not automatically pushed to the device unlike VNF devices. Instead, you must manually configure the PNF device by noting configuration that is generated on the Monitor Cloud OnRamp Colocation Clusters from Cisco vManage screen. The VLANs must be also configured on the Cisco Catalyst 9500-40X interfaces. See the ASR 1000 Series Aggregation Services Routers Configuration Guides and Cisco Firepower Threat Defense Configuration Guides for more information about the specific PNF configuration.

Configure PNF and Cisco Catalyst 9500 Switches

Procedure

Step 1 Identify ports from the switches where the PNF devices should be added, which are part of a service chain. To verify the availability of the ports, see .

Step 2 Connect with Cisco Catalyst 9500-40X or Cisco Catalyst 9500-48Y4C by using either the terminal server of any of the Cisco Catalyst 9500-40X or Cisco Catalyst 9500-48Y4C switches or use the **vty session** command with the IP address of the active switch.

Step 3 Configure VLANs from the generated configuration parameters on Cisco Catalyst 9500-40X or Cisco Catalyst 9500-48Y4C switches with interfaces that are connected to the PNF. See the Monitor Cloud OnRamp Colocation Clusters from Cisco vManage screen for the generated VLAN configuration.

Step 4 To configure a FTD or ASR 1000 Series device, note the configuration from the Monitor screen and then manually configure it on the device.

Custom Service Chain with Shared VNF Devices

You can customize service chains by including supported VNF devices.

Table 21: Feature History

Feature Name	Release Information	Feature Description
Share VNF Devices Across Service Chains	Cisco SD-WAN Release 19.2.1	This feature lets you share Virtual Network Function (VNF) devices across service chains to improve resource utilisation and reduce resource fragmentation.

Before you begin

Ensure that you note the following points about sharing VNF devices:

- You can share only the first, last, or both first and last VNF devices in a service chain.
- You can share a VNF with a minimum of one more service chain and maximum up to five service chains.
- Each service chain can have a maximum of up to four VNF devices in a service chain.
- You can share VNF devices only in the same service group

Procedure

Step 1 Create a service group and service chains within the service group. See Create Service Chain in a Service Group, on page 64.

Step 2 In the **Add Service Chain** dialog box, provide the service chain name, description, bandwidth, input VLAN handoff, output VLAN handoff, monitoring health information of a service chain, and service chain configuration. Click **Add**.

For the service chain configuration, choose **Create Custom** from the drop-down. An empty service chain in the design view window is available. In the left panel, the set of VNF devices and PNF devices that you can add into the service chain appears. The 'V' in the circumference of VNF devices represents a VNF and 'P' in the circumference of PNF devices represent a PNF.

Note
Ensure that you choose the **Create Custom** option for creating a shared VNF package.

Step 3 To add a VNF such as a router, load balancer, firewall, and others, click a VNF icon on the left panel, and drag the icon to its proper location within the service chain box.

After adding all required VNF devices, configure each of them.

a) Click a VNF in the service chain box.

The Configure VNF dialog box appears. To configure VNF, enter the following parameters:

b) Choose the software image to load from the **Image Package** drop-down.

To create a customized VNF package from vManage, see Create Customized VNF Image, on page 93.

c) Click **Fetch VNF Properties**.
d) Enter a name of the VNF in the **Name** field.
e) Enter the number of virtual CPUs required for the VNF in the **CPU** field.
f) Enter the amount of memory in megabytes to be allocated for the VNF in the **Memory** field.
g) Enter the amount of memory for storage in gigabytes to be allocated for the VNF in the **Disk** field.

h) Enter VNF-specific parameters, as required. See Create Service Chain in a Service Group, on page 64 for more information about VNF-specific properties.

These VNF-specific parameters are the custom user variables that are required for Day-0 operations of the VNF.

For a complete information about the list of user and system variables for different VNF types, when located at various positions, see Shared VNF Use Cases, on page 80 and Custom Packaging Details for Shared VNF, on page 159 .

Note
Ensure that you provide the values of the user variables if they are defined as mandatory, and for the system variables, vManage automatically sets the values for them.

i) Click **Configure**.

Step 4 To share VNF devices, repeat from step 2.

Step 5 Edit an existing VNF configuration by clicking it.

Step 6 Scroll down the VNF configuration slider to find the **Share NF To** field. Select the service chains from the **Share NF To** drop-down list with which the VNF should be shared.

After a VNF is shared, if you hover on a VNF, the respective shared VNF devices are highlighted in blue color. After you choose a NF to be shared, a blue rim appears on it.

Step 7 To delete a VNF or cancel the VNF configuration, click **Delete** or **Cancel** respectively.

You must attach service groups to a cluster.

Shared VNF Use Cases

The following images depict some of the shared VNF use cases and their predefined variable list:

Figure 20: Shared–Cisco vEdge Router VNF in First Position

The Cisco vEdge Router VNF in the first position is shared with the second service chain in the first position. The input to the first VNF is in access mode (hypervisor-tagged) and the neighbor (ASAv firewall) is in HA mode.

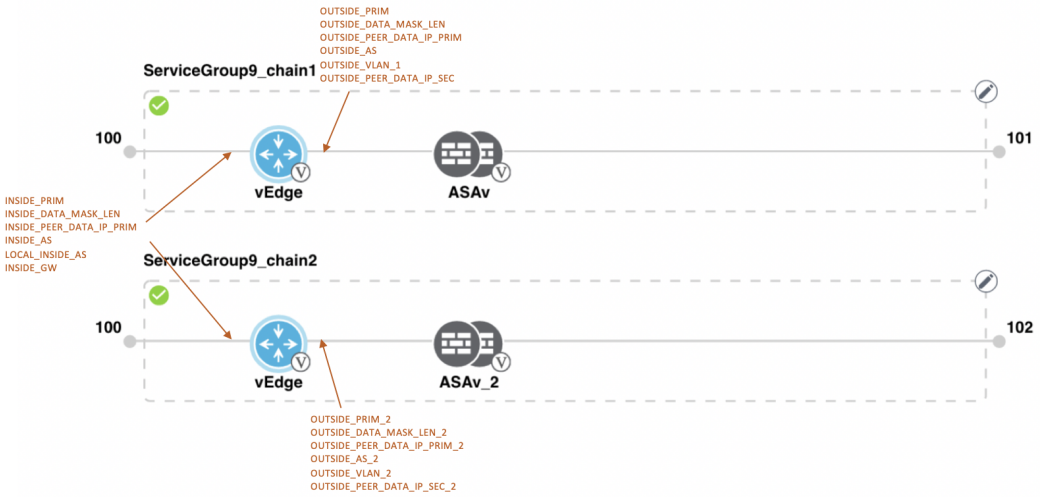


Figure 21: Shared–Cisco vEdge Router VNF in First Position

The Cisco vEdge Router VNF in the first position is shared with the second service chain in the first position. The input to the first VNF is in access mode (hypervisor-tagged) and the neighbor is in StandAlone mode.

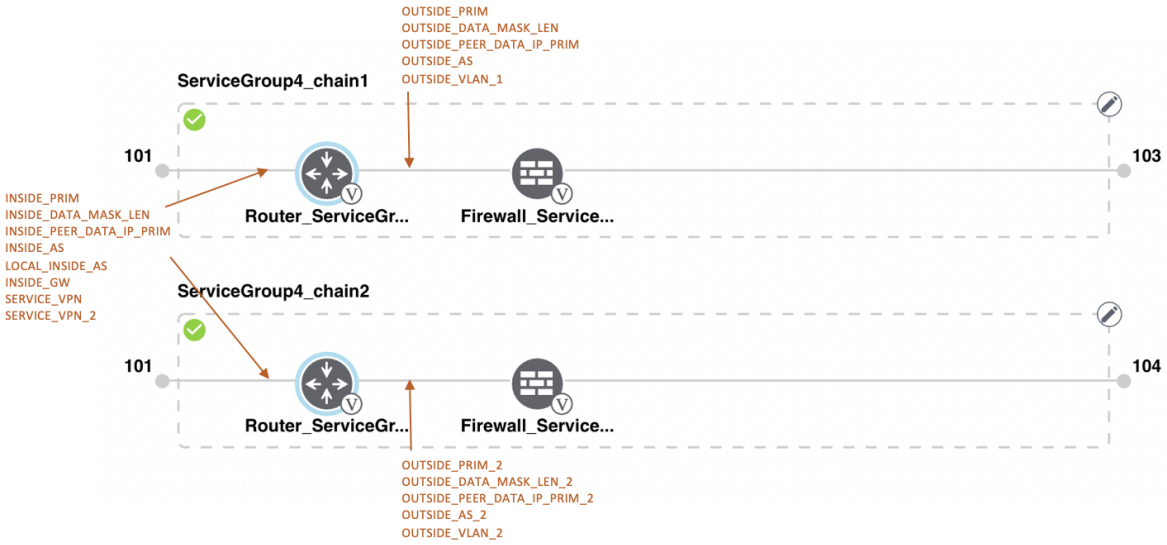


Figure 22: Shared–Cisco vEdge Router VNF in First Position

The Cisco vEdge Router VNF in the first position is shared with the second service chain in the first position. The input to the first VNF is in trunk mode (VNF-tagged) and the neighbor is in StandAlone mode.

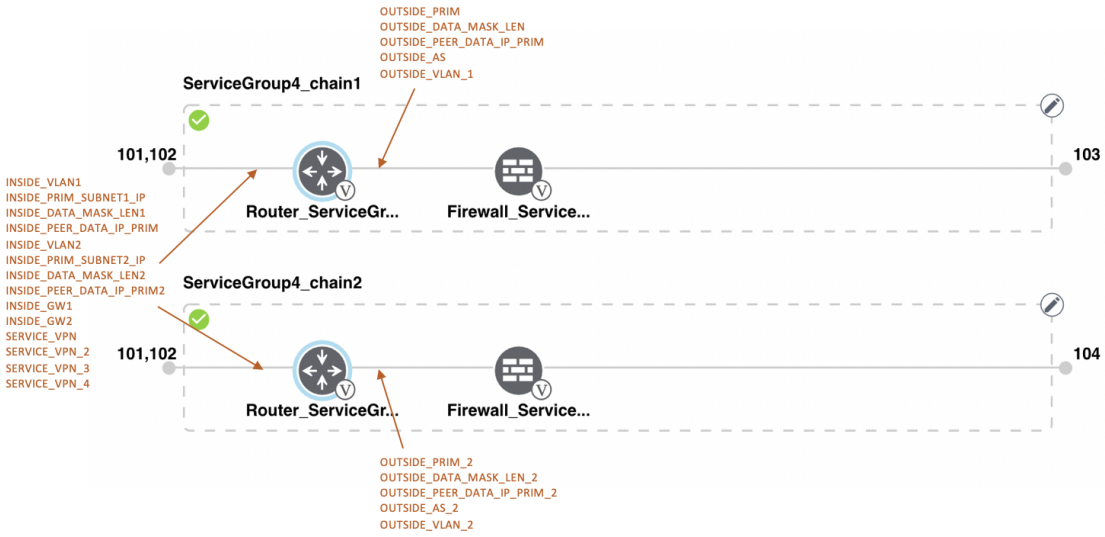


Figure 23: Shared–Cisco vEdge Router VNF in First Position

The Cisco vEdge Route VNF in the first position is shared with the second service chain in the first position. The input to the first VNF is in trunk mode (VNF-tagged) and the neighbor is in HA mode.

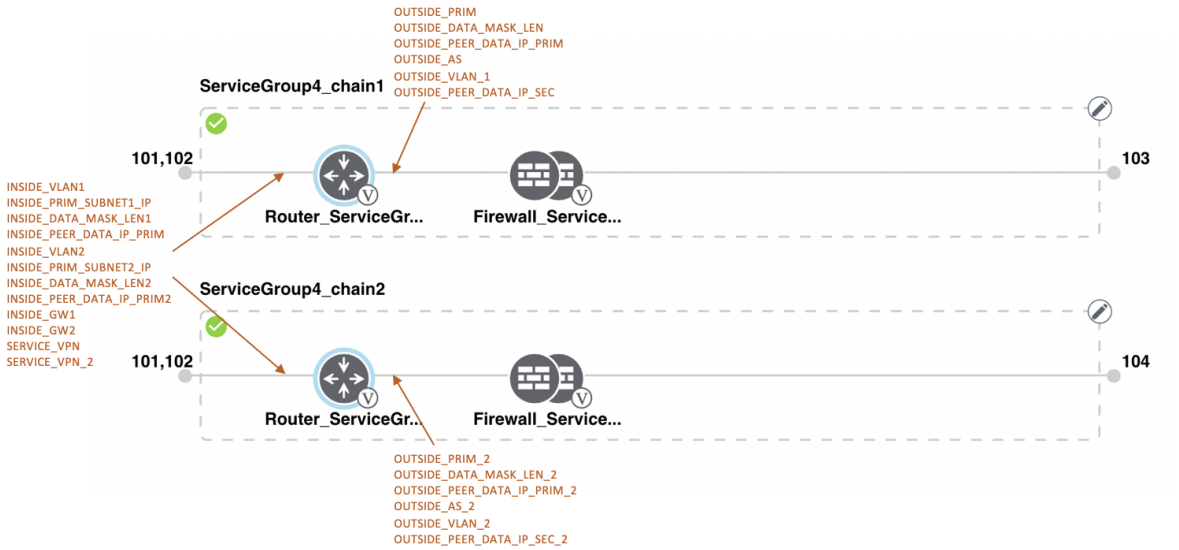


Figure 24: Shared-Cisco CSR1000V VNF in Last Position

The Cisco CSR1000V VNF in the last position is shared with the second service chain in the second position. The output from the last VNF is in access mode (hypervisor-tagged) and the neighbor (ASAv firewall) is in StandAlone mode.

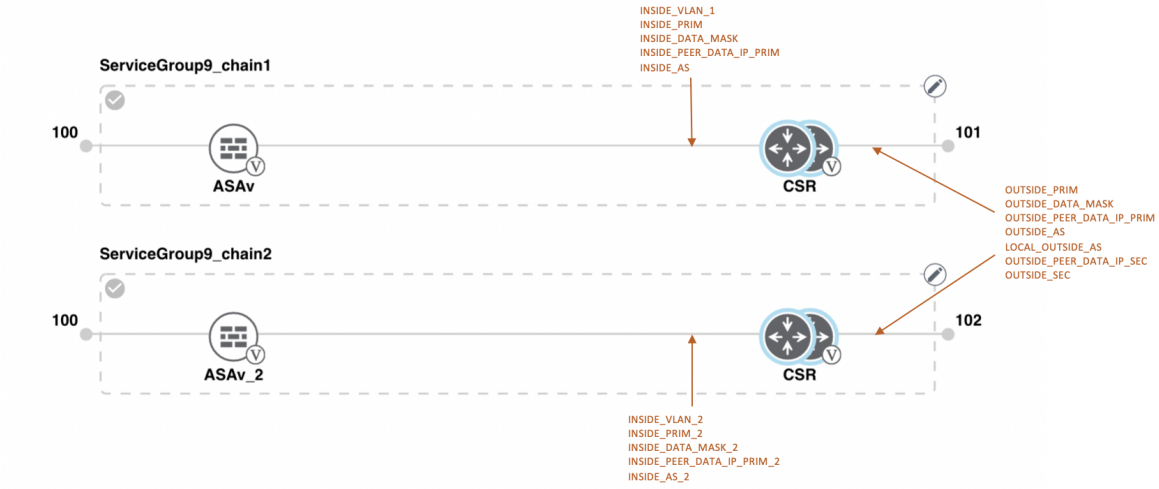


Figure 25: Shared–Cisco CSR1000V VNF in Last Position

The Cisco CSR1000V VNF in the last position is shared with the second service chain in the second position. The output from the last VNF is in access mode (hypervisor-tagged) and the neighbor is in StandAlone mode.

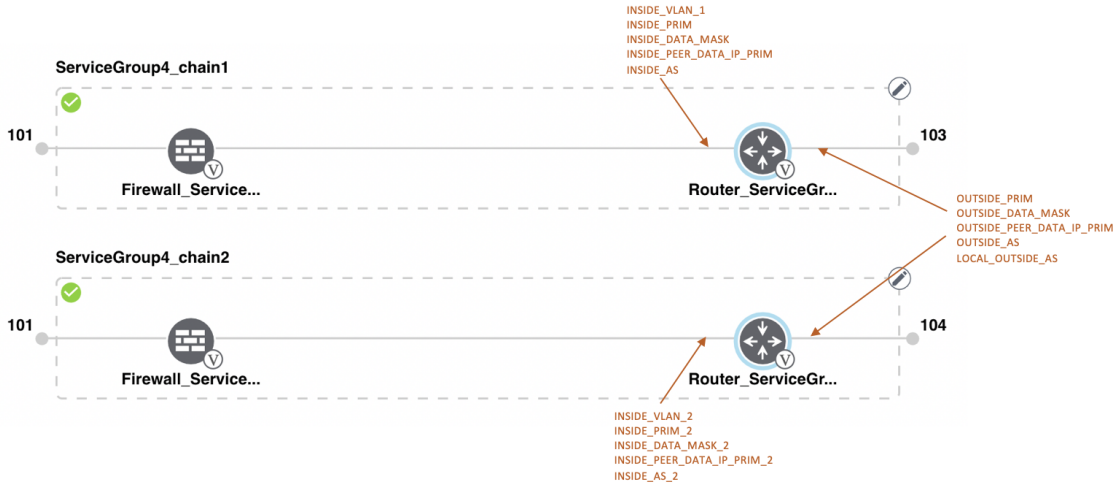


Figure 26: Shared–Cisco CSR1000V VNF in Last Position

The Cisco CSR1000V VNF in the last position is shared with the second service chain in the second position. The output from the last VNF is in access mode (hypervisor-tagged) and the neighbor (Firewall_Service) is in HA mode.

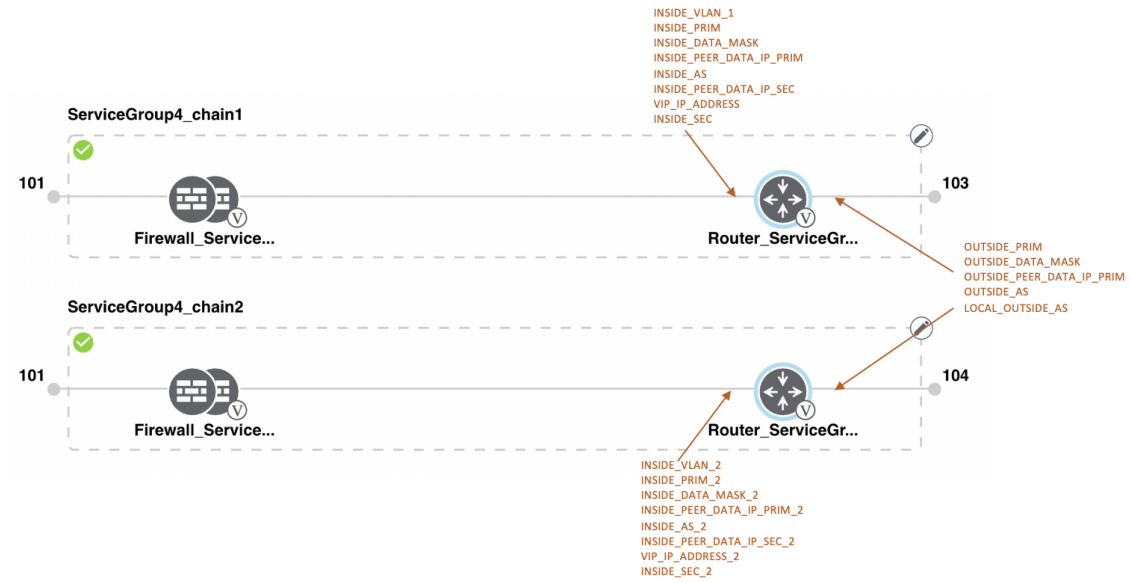


Figure 27: Shared–Cisco CSR1000V VNF in Last Position

The Cisco CSR1000V VNF in the last position is shared with the second service chain in the second position. The output from the last VNF is in access mode (hypervisor-tagged) and the neighbor (Firewall_Service) is in HA mode.

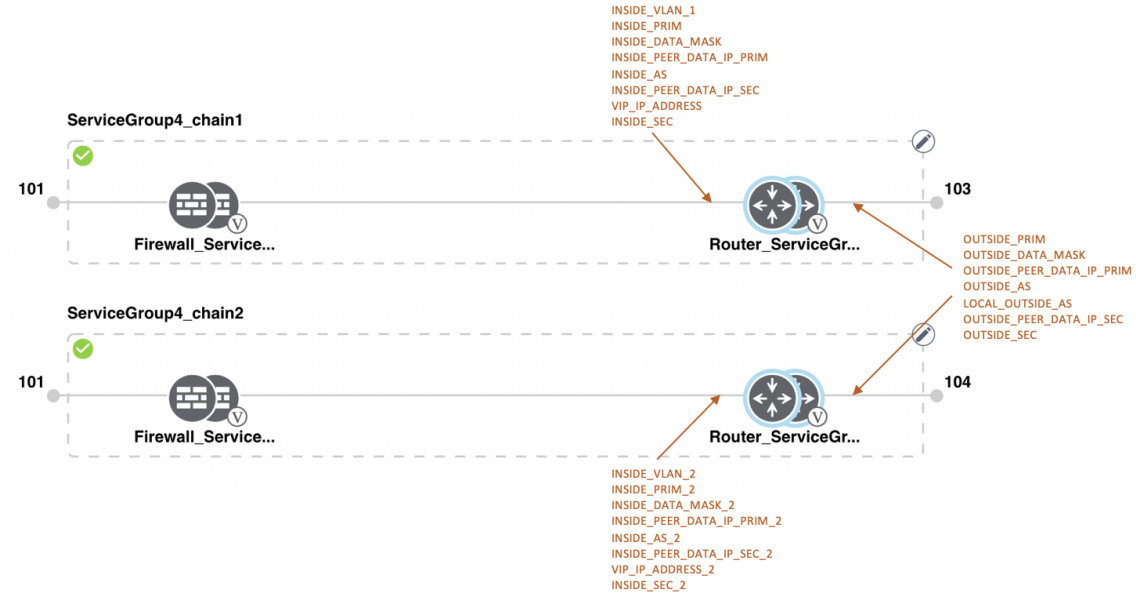


Figure 28: Shared–ASAv VNF in First Position

The ASAv VNF in the first position is shared with the second service chain in the first position. The input to the first VNF is in access mode (hypervisor-tagged) and the neighbor is in redundant mode.

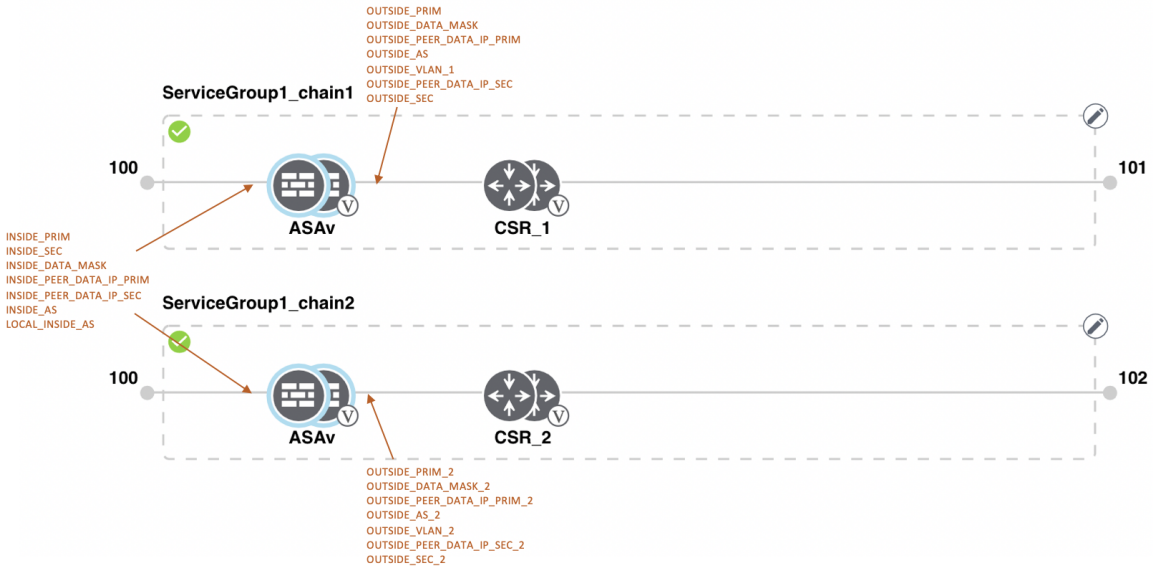


Figure 29: Shared–ASAv VNF in First Position

The ASAv (Firewall_Service) VNF in the first position is shared with the second service chain in the first position. The input to the first VNF is in access mode (hypervisor-tagged) and the neighbor is in StandAlone mode.

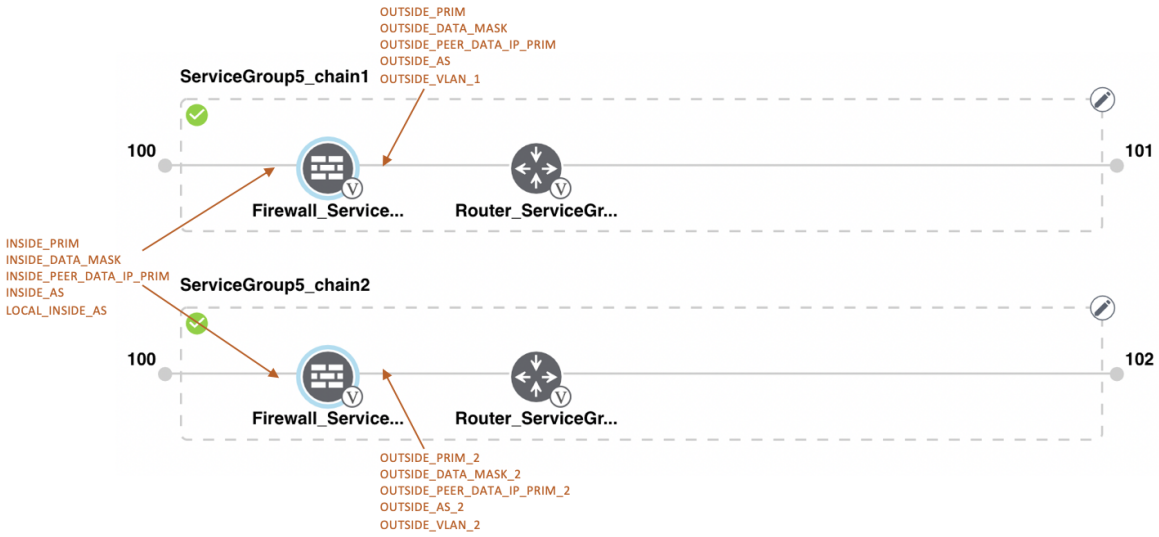


Figure 30: Shared–ASAv VNF in First Position

The ASAv (Firewall_Service) VNF in the first position is shared with the second service chain in the first position. The input to the first VNF is in access mode (hypervisor-tagged) and the neighbor, which is a router is in redundant mode.

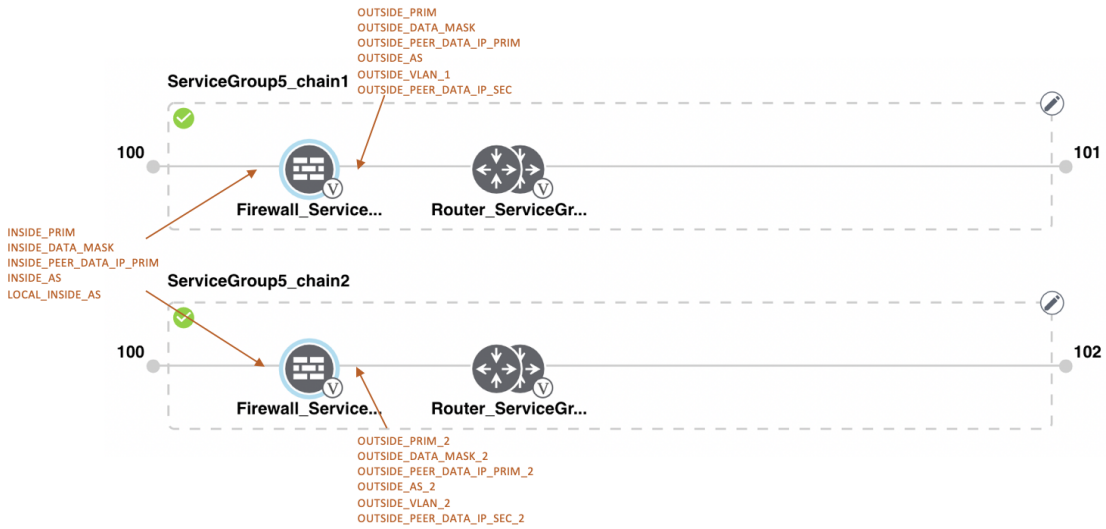
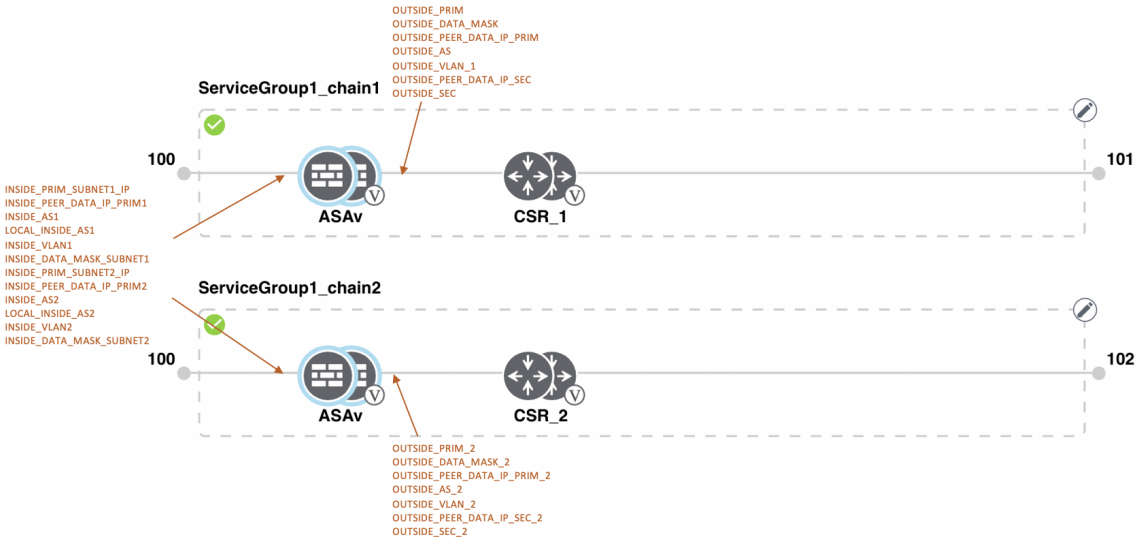


Figure 31: Shared–ASAv VNF in First Position

The ASAv VNF in the first position in HA mode is shared with the second service chain in the first position. The input to the first VNF is in trunk mode (vnf-tagged) and the neighbor is in redundant mode.



View Service Groups in vManage

To view service groups, perform the following steps:

Procedure

Step 1 Log in to Cisco vManage with admin credentials at the URL HTTPS: //vManage-ip-address/.

Step 2 In Cisco vManage, click **Configuration** > **Cloud OnRamp for Colocation**. You can view all the service groups in a table. To view service chains:

a) Click the **Service Group** tab.

b) Click a service group, click the **More Actions** icon to the right of the row, and then click **View** next to the service group.

You can view the service chains in the design window.

Edit Service Group in vManage

Before attaching a service group with a cluster, you can edit all parameters. After attaching a service group with a cluster, you can only edit monitoring configuration parameters. Also, after attaching a service group, you can only add new service chains but not edit or attach a service chain. Hence, ensure that you detach a service group from a cluster before editing an existing service chain. To edit and delete a service group, perform the following steps:

Procedure

Step 1 Log into Cisco vManage, at the URL HTTPS: //vManage-ip-address/, as a user with "admin" credentials.

Step 2 In vManage, click **Configuration** > **Cloud OnRamp for Colocation**. The CLOUD ONRAMP FOR COLOCATION Cluster screen appears, and the **Configure & Provision Cluster** button is highlighted. In the CLOUD ONRAMP FOR COLOCATION Cluster screen, perform the following tasks:

a) Click the **Service Group** tab.

b) To modify either service chain configuration or modify a VNF configuration, click a router or firewall VNF icon.

c) To add new service chains, click a service chain button.

Attach and Detach Service Group with Cluster

To complete the Cisco SD-WAN Cloud onRamp for Colocation configuration, you must attach service groups to a cluster. To attach or detach a service group from a cluster, perform the following steps:

Procedure

Step 1 Log in to Cisco vManage with admin credentials at the URL HTTPS: //vManage-ip-address/.

Step 2 In Cisco vManage, click **Configuration** > **Cloud OnRamp for Colocation**.

a) In the **Cluster** tab, click a cluster from the table, click the **More Actions** icon to the right of the row, and then click **Attach Service Groups**.

Step 3 In the **Attach Service Groups** dialog box, choose a service group from the available service groups.

Step 4 Click the right arrow to move the chosen service groups to the selected box.

Step 5 Click **Attach**.

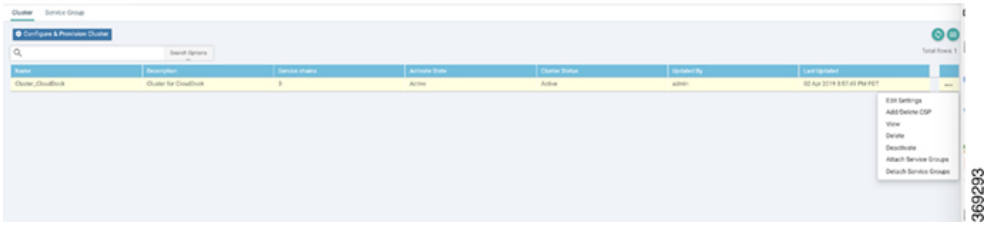
Step 6 To detach a service group from a cluster:

a) In the **Cluster** tab, click a cluster from the table, click the **More Actions** icon to the right of the row.

b) Click **Detach Service Groups**.

You can't attach or detach a single service chain within a service group.

Step 7 To verify if service groups are attached or detached, you can view from the following Cisco vManage screen:

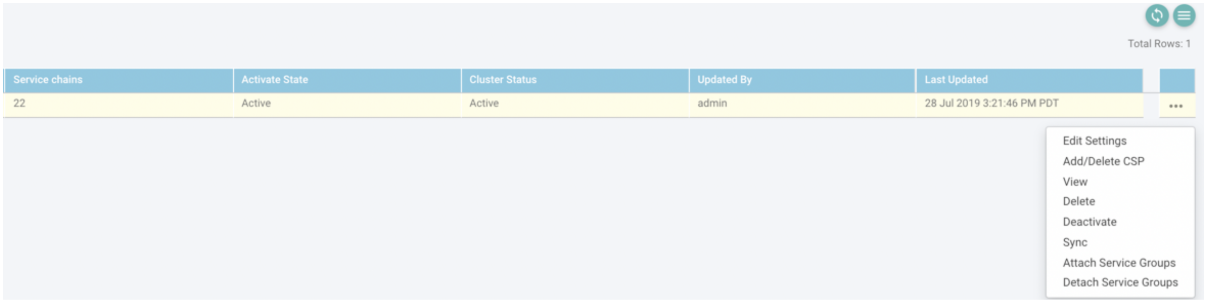


If the status of the tasks in ask view window are "FAILURE" or in "PENDING" state for long duration, see Troubleshoot Service Chain Issues, on page 155 the topic, "Troubleshoot Service Chain Issues" in the Cisco SD-WAN Cloud OnRamp for Colocation Solution Guide.

If a Cisco Colo Manager task fails, see Troubleshoot Cisco Colo Manager Issues, on page 153.

Note If a cluster goes to a "PENDING" state, click the **More Actions** icon to the right of the row, and then click the **Sync** button. This action moves the cluster back to an "ACTIVE" state. The sync button keeps the Cisco vManage in sync with the devices and is visible when a cluster is active.

Figure 32: Sync Button for a Cluster



Day-N Configuration Workflow of Cisco SD-WAN Cloud onRamp for Colocation Solution

The following is the background process for a Day-N configuration.

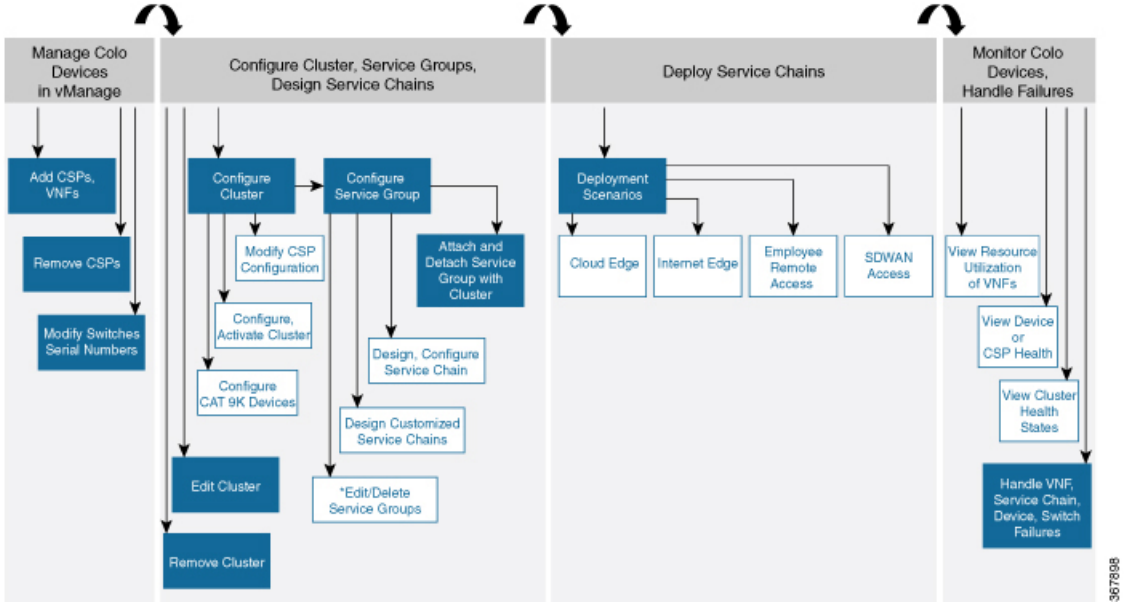
- All Day-N configuration from vManage requires clusters to be in-sync (devices have to be in synchronization with vManage) state.
- During attaching a service group, vManage runs the Placement logic to determine which VMs are placed on which CSP device.
- Switch-related Day-N configuration from vManage requires Cisco Colo Manager to be in a Healthy state.
- vManage pushes all switch-related service chain, cluster, switch configuration to Cisco Colo Manager.
- Cisco Colo Manager moves to In-progress state for any configuration that is received.
- Cisco Colo Manager translates all global and service chain configuration of Cisco Colo Manager into the device-specific configuration.
- Cisco Colo Manager reports the states to vManage whether a configuration push is a success or failure.
- All the Day-N service chain or VM configuration is sent to CSP devices.
- CSP devices send notification to vManage about the VM file download status.
- After all VMs are downloaded, vManage sends the bulk configuration to spin all VMs.
- CSP devices send notifications to vManage about VM being brought up and its states.

- If any switches return error, vManage reports error with a detailed information and the cluster moves to a FAILURE state.

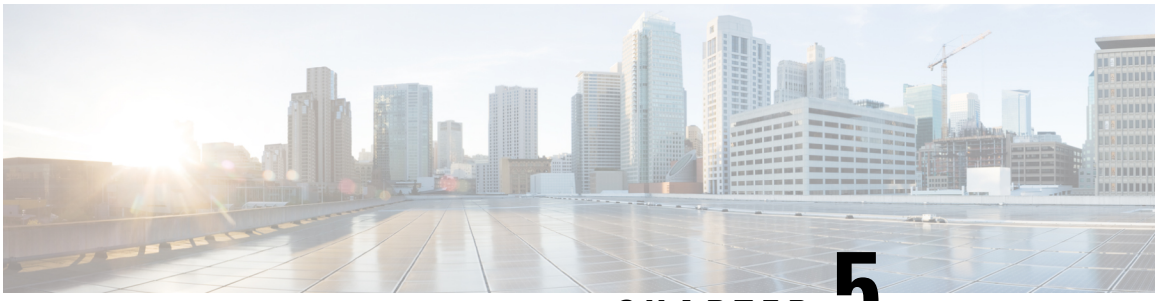
Ensure that you fix errors that are based on notifications and error messages, and then activate the Cloud OnRamp for Colocation again.

Note During the Day-N configuration, you can modify Serial Number of switches for a maximum of two switches that are allowed.

Figure 33: Day-N Workflow



Note *You can only edit service groups if they are detached from a cluster.



CHAPTER 5

Software Image Management for Cluster Components and SWIM

- Manage VM Catalog and Repository, on page 91
- Upgrade NFVIS Software Through vManage, on page 99
- Upgrade Cisco Catalyst 9500 Switches, on page 101
- Supported Upgrade Scenarios and Recommended Connections, on page 104

Manage VM Catalog and Repository

vManage supports uploading a prepackaged Cisco VM image, tar.gz in this phase. Alternatively, you can package the VM image by providing a root disk image in any of the supported formats (qcow2). Use the linux command-line NFVIS VM packaging tool, **nfvpt.py** to package the qcow2 or alternatively create a customized VM image from vManage. See Create Customized VNF Image, on page 93.

If VM is SR-IOV capable, which means sriov_supported is set to true in image_properties.xml in the vm package *.tar.gz. Also, the service chain network is automatically connected to SR-IOV network. If sriov_supported is set to false, OVS network is created on the data port channel. It is attached to VM VNICs for service chaining, which is done by using the OVS network. For the Cisco SD-WAN Cloud onRamp for Colocation solution, service chaining a VM uses homogeneous type of network. This type of network means it is either OVS or SR-IOV, and not a combination of SR-IOV and OVS.

Only two data VNICs are attached to any VM–one for inbound traffic and the other for outbound traffic. If more than two data interfaces are required, use subinterfaces configuration within the VM. The VM packages are stored in the VM catalog.

Note Each VM type such as firewall can have multiple VM images that are uploaded to vManage from same or different vendors being added to the catalog. Also, different versions that are based on the release of the same VM can be added to the catalog. However, ensure that the VM name is unique.

The Cisco VM image format can be bundled as *.tar.gz and can include:

- Root disk images to boot the VM.
- Package manifest for checksum validation of the file listing in the package.
- Image properties file in XML format that lists the VM meta data.

- (Optional) Day-0 configuration, other files that are required to bootstrap the VM.
- (Optional) HA Day-0 configuration if VM supports stateful HA.
- System generated properties file in XML format that lists the VM system properties

VM images can be hosted on both HTTP server local repository that vManage hosts or the remote server.

If VM is in NFVIS supported VM package format such as, tar.gz, vManage performs all the processing and you can provide variable key and values during VNF provisioning.



Note vManage only manages the Cisco VNFs, whereas Day-1 and Day-N configurations within VNF are not supported for other VNFs. See the NFVIS Configuration Guide, VM Image Packaging for more information about VM package format and content, and samples on image_properties.xml and manifest (package.mf).

To upload multiple packages for the same VM, same version, Communication Manager (CM) type, ensure that one of the three values (name, version, VNF type) are different. Then, you can repackage the VM *.tar.gz to be uploaded.

VNF Image Format

The service orchestrator does not distinguish between Cisco VNFs and third-party VNFs. All VNFs are categorized based on the services that are provided by the VNF such as router, firewall, load balancer, and others. The package metadata has VM_specific attributes. Based on HA NICs and management NICs specified in the package metadata file, orchestrator attaches management NIC and HA NIC. By default, management NIC is zero and HA NIC is one. The number of HA NICs that is specified is attached during VNF provisioning.

Upload VNF Images in Cisco vManage Repository

The VNF images are stored in the Cisco vManage software respository. These VNF images are referenced during service chain deployment, and then they are pushed to NFVIS during service chain attachment.

Procedure

Step 1 Log into Cisco vManage, at the URL HTTPS: //vManage-ip-address/, as a user with "admin" credentials.

Step 2 In Cisco vManage, click **Maintenance** > **Software Repository**. The Maintenance|Software Repository screen appears, and the **Add New Software** button is highlighted. To upload VNF images, use the **Virtual Images** tab. In the Maintenance|Software Repository screen, perform the following tasks:

a) To add a prepackaged VNF image, click the **Virtual Images** tab, and then click the **Upload Virtual Images** button.

b) Choose the location to store the virtual image.

- To store the virtual image on the local vManage server and then get it downloaded to CSP devices over a control plane connection, click **vManage**. The **Upload Software to vManage** dialog box appears.

  1. Drag and drop the virtual image file to the dialog box or click **Browse** to choose the virtual image from the local Cisco vManage server. For example, CSR.tar.gz, ASAv.tar.gz
  2. Click **Upload** to add the image to the virtual image repository. The virtual image repository table displays the added virtual image, and it is available for installing on the CSP devices.

- To store the image on a remote Cisco vManage server and then get it downloaded to CSP devices over an out-of-band management connection, click **Remote Server - vManage**. The **Upload Virtual Image to Remote Server - vManage** dialog box appears.

  1. In **vManage Hostname/IP Address**, enter the IP address of an interface on Cisco vManage server that is in a management VPN (typically, VPN 512).
  2. Drag and drop the virtual image file to the dialog box, or click **Browse** to choose the virtual image from the local vManage server.
  3. Click **Upload** to add the image to the virtual image repository. The virtual image repository table displays the added virtual image, and it is available for installing on the CSP devices.

c) Click **Submit**.

You can have multiple VNF entries such as a firewall from same or different vendors. Also, different versions of VNF that are based on the release of the same VNF can be added. However, ensure that the VNF name is unique.

Create Customized VNF Image

Before you begin

You can upload one or more qcow2 images in addition to a root disk image as an input file along with VM-specific properties, bootstrap configuration files (if any), and generate a compressed TAR file. Through custom packaging, you can:

- Create a custom VM package along with image properties and bootstrap files (if needed) into a TAR archive file.
- Tokenize custom variables and apply system variables that are passed with the bootstrap configuration files.

Ensure that the following custom packaging requirements are met:

- Root disk image for a VNF–qcow2
- Day-0 configuration files–system and tokenized custom variables
- VM configuration–CPU, memory, disk, NICs
- HA mode–If a VNF supports HA, specify Day-0 primary and secondary files, NICs for a HA link
- Additional Storage–If additional storage is required, specify predefined disks (qcow2), storage volumes (NFVIS layer)

Procedure

Step 1 Log into Cisco vManage, at the URL HTTPS: //vManage-ip-address/, as a user with "admin" credentials.

Step 2 In the **Maintenance** > **Software Repository** screen, click the **Add Custom VNF Package** button from the **Virtual Images** tab.

Step 3 Configure the VNF with the following VNF package properties and click **Save**.

Table 22: VNF Package Properties

Field	Mandatory or Optional	Description
Package Name	Mandatory	Specifies the filename of the target VNF package. It is the NFVIS image name with .tar or .gz extensions.
App Vendor	Mandatory	Specifies whether Cisco VNFs or third-party VNFs.
Name	Mandatory	Specifies name of the VNF image.
Version	Optional	Specifies version number of the program.
Type	Mandatory	Choose VNF type. Supported VNF types are: Router, Firewall, Load Balancer, and Other.

Step 4 To package a VM qcow2 image, click **File Upload** under **Image**, and browse to choose a qcow2 image file.

Step 5 To choose a bootstrap configuration file for VNF, if any, click the **Bootstrap Files** button under **Day 0 Configuration**, click **File Upload**, and then browse to choose a bootstrap file.

Include the following Day-0 configuration properties:

Table 23: Day-0 Configuration

Field	Mandatory or Optional	Description
Mount	Mandatory	Specifies the path where the bootstrap file gets mounted.
Parseable	Mandatory	Specifies whether a Day-0 configuration file can be parsed or not. Options are: true or false. By default, it is true.
High Availability	Mandatory	Choose high availability of a Day-0 configuration file. Supported values are: Standalone, HA Primary, HA Secondary.

Note
If any bootstrap configuration is required for a VNF, you must create *bootstrap-config* or *day0-config*.

Step 6 To add a Day-0 configuration, click **Add**, and then click **Save**. The Day-0 configuration appears in the **Day 0 Config File** table. You can tokenize the bootstrap configuration variables with system and custom variables. To tokenize

variables of a Day-0 configuration file, click **View Configuration File** against the configuration file. In the **Day 0 configuration file** dialog box, perform the following tasks:

Note
The bootstrap configuration file is an XML or a text file, and contains properties specific to a VNF and the environment. For a shared VNF, see Custom Packaging Details for Shared VNF, on page 159 for the list of system variables that must be added for different VNF types..

a) To add a system variable, in the **CLI configuration** dialog box, select and highlight a property from the text fields. Click **System Variable**. The **Create System Variable** dialog box appears.
b) Choose a system variable from the **Variable Name** drop-down, and click **Done**. The highlighted property is replaced by the system variable name.
c) To add a custom variable, in the **CLI configuration** dialog box, select and highlight a custom variable attribute from the text fields. Click **Custom Variable**. The **Create Custom Variable** dialog box appears.
d) Enter custom variable name and choose a type from **Type** drop-down.
e) To set the custom variable attribute, do the following:

- To ensure that the custom variable is mandatory when creating a service chain, check the **Type** check box against **Mandatory**.
- To ensure that a VNF includes both primary and secondary Day-0 files, check the **Type** check box against **Common**.

f) Click **Done**, and then click **Save**. The highlighted custom variable attribute is replaced by the custom variable name.

Step 7 To upload extra VM images, expand **Advance Options**, click **Upload Image**, and then browse to choose an additional qcow2 image file. Choose the root disk, Ephemeral disk 1, or Ephemeral disk 2, and click **Add**. The newly added VM image appears in the **Upload Image** table.

Note
Ensure that you do not combine ephemeral disks and storage volumes when uploading extra VM images.

Step 8 To add the storage information, expand **Add Storage**, and click **Add volume**. Provide the following storage information and click **Add**. The added storage details appear in the **Add Storage** table.

Table 24: Storage Properties

Field	Mandatory or Optional	Description
Size	Mandatory	Specifies the disk size that is required for the VM operation. The maximum disk size can be 256 if the size unit is GiB.
Size Unit	Mandatory	Choose size unit. Supported units are: MIB, GiB, TiB.
Device Type	Optional	Choose a disk or CD-ROM. Default is a disk.
Location	Optional	Specifies location of the disk or CD-ROM. By default, it is local.

Field	Mandatory or Optional	Description
Format	Optional	Choose a disk image format. Supported formats are: qcow2, raw, and vmdk. Buy default, it is raw.
Bus	Optional	Choose a value from the drop-down. Supported values for a bus are: virtio, scsi, and ide. By default, it is virtio.

Step 9 To add VNF image properties, expand **Image Properties** and provide the following image information.

Table 25: VNF Image Properties

Field	Mandatory or Optional	Description
SR-IOV Mode	Mandatory	Specifies enabling or disabling SR-IOV support. By default, it is enabled.
Monitored	Mandatory	VM health monitoring for those VMs that can be bootstrapped. Options are: enable or disable. By default, it is enabled.
Bootup Time	Mandatory	Specifies monitoring timeout period for a monitored VM. By default, it is 600 seconds.
Serial Console	Optional	Specifies serial console that is supported or not. Options are: enable or disable. By default, it is disabled.
Privileged Mode	Optional	Allows special features like promiscuous mode and snooping. Options are: enable or disable. By default, it is disabled.
Dedicate Cores	Mandatory	Facilitates allocation of a dedicated resource (CPU) to supplement a VM's low latency (for example, router and firewall). Otherwise, shared resources are used. Options are: enable or disable. By default, it is enabled.

Step 10 To add VM resource requirements, expand **Resource Requirements** and provide the following information.

Table 26: VM Resource Requirements

Field	Mandatory or Optional	Description
Default CPU	Mandatory	Specifies CPUs supported by a VM. The maximum numbers of CPUs supported are 8.
Default RAM	Mandatory	Specifies RAM supported by a VM. The RAM can range from 2–32.
Disk Size	Mandatory	Specifies disk size in GB supported by a VM. The disk size can range from 4–256.
Max number of VNICs	Optional	Specifies maximum number of VNICs allowed for the VM. The number of VNICs can range from 8–32 and the default value is 8.
Management VNIC ID	Mandatory	Specifies the management VNIC ID corresponding to the management interface. Valid range is from 0 to maximum number of VNICs.
Number of Management VNICs ID	Mandatory	Specifies number of VNICs.
High Availability VNIC ID	Mandatory	Specifies VNIC IDs where high availability is enabled. Valid range is from 0–maximum number of VNICs. It should not conflict with management VNIC Id. The default value is 1.
Number of High Availability VNICs ID	Mandatory	Specifies maximum number of VNIC IDs where high availability is enabled. Valid range is 0–(maximum number of VNICs-number of management VNICs-2) and default value is 1.

Step 11 To add Day-0 configuration drive options, expand **Day0 Configuration Drive options** and provide the following information.

Table 27: Day-0 Configuration Drive Options

Field	Mandatory or Optional	Description
Volume Label	Mandatory	Displays the volume label of the Day-0 configuration drive. Options are: V1 or V2. By default, it is V2. V2 is the config-drive label config-2. V1 is config-drive label cidata.

Field	Mandatory or Optional	Description
Init Drive	Optional	Mounts the Day-0 configuration file as a disk. The default drive is CD-ROM.
Init Bus	Optional	Choose an init bus. Supported values for a bus are: virtio, scsi, and ide. By default, it is ide.

The Software Repository table displays the customized VNF image, and it is available for choosing while creating a custom service chain.

View VNF Images in vManage Repository

Procedure

Step 1 Log into Cisco vManage, at the URL HTTPS: //vManage-ip-address/, as a user with "admin" credentials.

Step 2 In Cisco vManage, click **Maintenance** > **Software Repository**. The Maintenance|Software Repository screen appears, and the **Add New Software** button is highlighted. To view VNF images, use the **Virtual Images** tab. In the Maintenance|Software Repository screen, perform the following tasks:

a) To view VNF images, click the **Virtual Images** tab. The images in the repository are displayed in the table.

b) To filter the list, search or type a string in the Search box.

The Software Version column provides the version of the software image.

The Software Location column indicates where the software images are stored. It can be stored either in the repository on Cisco vManage server or in a repository in a remote location.

The Version Type Name column provides the type of firewall.

The Available Files column lists the names of the VNF image files.

The Update On column displays when the software image was added to the repository.

c) To view details of a VNF image, click a VNF image, click the **More Actions** icon, and click **Show Info** against the VNF image.

Delete VNF Images from vManage Repository

Procedure

Step 1 Log into Cisco vManage, at the URL HTTPS: //vManage-ip-address/, as a user with "admin" credentials.

Step 2 In Cisco vManage, click **Maintenance** > **Software Repository**. The Maintenance|Software Repository screen appears, and the **Add New Software** button is highlighted. To upload VM images, use the **Virtual Images** tab. In the Maintenance|Software Repository screen, perform the following tasks:

a) To delete a VM image, click the **Virtual Images** tab. The images in the repository are displayed in the table.
b) In the repository table, click a VM image.
c) Click the **More Actions** icon to the right of its row, and click **Delete** against the VM image.

Note
If a VNF image is being download to a router, you cannot delete the VNF image until the download process completes.



Note If the VNF image is referenced by a service chain, it cannot be deleted.

Upgrade NFVIS Software Through vManage

To upload and upgrade NFVIS, the upgrade image must be available as an archive file that can be uploaded to vManage repository through vManage. After you upload the NFVIS image, the upgraded image can be applied to a CSP device by using the Software Upgrade screen in vManage. You can perform the following tasks during upgrading NFVIS software through vManage:

- Upload NFVIS upgrade image. See Upload NFVIS Upgrade Image, on page 99.
- Upgrade a CSP device with the uploaded image. See Upgrade CSP Device with NFVIS Upgrade Image, on page 100.
- View the upgrade status in the CSP device. See the "View Log of Software Upgrade Activities" in the Cisco SD-WAN Configuration Guide.

Upload NFVIS Upgrade Image

Procedure

Step 1 Download the NFVIS upgrade image from a prescribed location to your local system. You can also download the software image to an FTP server in your network.

Step 2 Log into Cisco vManage, at the URL HTTPS: //vManage-ip-address/, as a user with "admin" credentials.

Step 3 In the **Maintenance** > **Software Repository** screen, click the **Add New Software** > **Remote Server/Remote Server - vManage** button.

You can either store the software image on a remote file server, on a remote vManage server, or on a vManage server.

Note
The vManage server is available in the current version.

vManage server–saves software images on a local vManage server.

Remote server–saves the URL pointing to the location of the software image and can be accessed through an FTP or HTTP URL.

Remote vManage server–saves software images on a remote vManager server and location of the remote vManage server is stored in the local vManage server.

Step 4 To add the image to the software repository, browse and choose the NFVIS upgrade image that you had downloaded in step1.

Step 5 Click **Add|Upload**.

The Software Repository table displays the added NFVIS upgrade image, and it is available for installing on the CSP devices. See the "Software Repository" topic in the Cisco SD-WAN Configuraion Guides.

Upgrade CSP Device with NFVIS Upgrade Image

Before you begin

Ensure that the NFVIS software versions are the files that have `.nfvispkg` extension.

Procedure

Step 1 Log into Cisco vManage, at the URL HTTPS: //vManage-ip-address/, as a user with "admin" credentials.

Step 2 In the **Maintenance** > **Software Upgrade** > **WAN Edge** screen, view the list of all CSP devices along with their current and available versions.

Step 3 Select one or more devices, and click **Upgrade**.

Step 4 Choose a CSP device on which to upgrade the NFVIS software image.

Step 5 Click the **Upgrade** button. The **Software Upgrade** dialog box appears.

Step 6 Choose the NFVIS software version to install on the CSP device. If software is located on a remote server, choose the appropriate remote version.

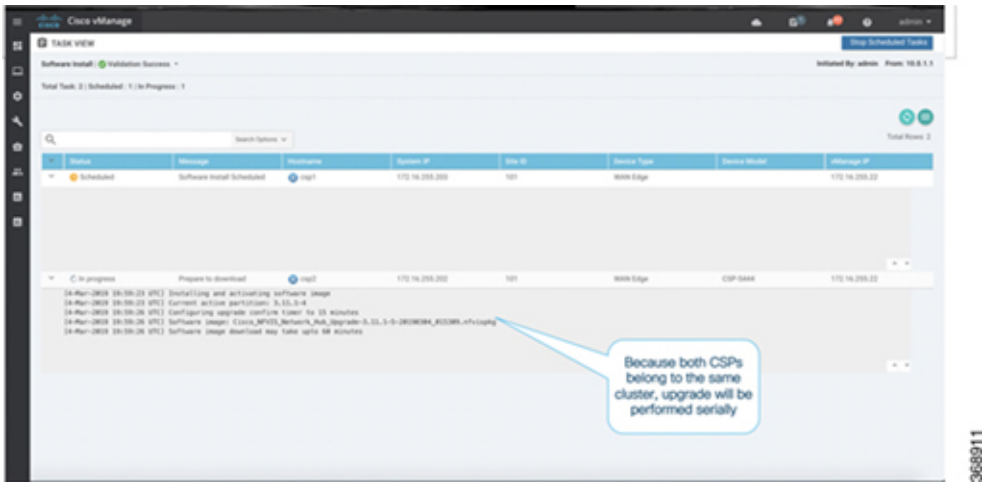
Step 7 To automatically upgrade and activate with the new NFVIS software version and reboot the CSP device, check the **Activate and Reboot** checkbox.

If you do not check the **Activate and Reboot** checkbox, the CSP device downloads and verifies the software image. However, the CSP device continues to run the old or current version of the software image. To enable the CSP device to run the new software image, you must manually activate the new NFVIS software version by selecting the device again and clicking the **Activate** button on the **Software Upgrade** page. For more information about activation, see the "Activate a New Software Image" topic in the Cisco SD-WAN Configuration Guides.

Step 8 Click **Upgrade**.

To view the status of software upgrades, the task view page displays a list of all running tasks along with total number of successes and failures. The page periodically refreshes and displays messages to indicate the progress or status of the upgrade. You can easily access the software upgrade status page by clicking the Tasks icon located in the Cisco vManage toolbar.

Note
If two or more CSP devices belonging to the same cluster are upgraded, the software upgrade for the CSP devices happen in a sequence.



Note
The **Set the Default Software Version** option is not available for NFVIS images.

The CSP device reboots and the new NFVIS version is activated on it. This reboot happens during the **Activate** phase. The activation can either happen immediately after upgrade if you check the **Activate and Reboot** check box, or by manually selecting the activate button after selecting the device again.

To verify if CSP device has rebooted and is running, Cisco vManage polls your entire network every 90 seconds up to 30 times.



Note You can delete an NFVIS software image from a CSP device if the image version is not the active version that is running on the device.

Upgrade Cisco Catalyst 9500 Switches

You can perform a software upgrade for both Cisco Catalyst 9500-40X and Cisco Catalyst 9500-48Y4C switches.

Before you begin

- Back up the running configuration in both the switches
- Ensure that you download the Cisco Catalyst 9500 upgrade software (.bin file) from cisco.com website and it is available as an archive file.

Procedure

Step 1 To copy the upgraded software from Trivial File Transfer Protocol (TFTP) to the flash of switch1, use the following commands:

a) **conf t**

Enters the configuration mode one per line. Ends with CNTL/Z.

Example:

```
c9500-1#conf t
```

b) **blocksize** *value*

Manually changes the block size in the global configuration to speed up the transfer process.

Example:

```
c9500-1(config)#ip tftp blocksize 8165
c9500-1(config)#end
```

c) **copy scp**

Securely copies switch image files to the flash of switch1.

Example:

```
c9500-1#copy scp://<cec-id>@172.16.0.151//auto/tftp-xxx-users2/yyyy/Switch_Image/
cat9k_iosxe.17.03.01.SPA.bin flash: vrf Mgmt-vrf
```

Step 2 To copy the upgraded software from one switch to another switch when they are in the SVL mode, use the following commands.

If both the switches are not in SVL mode, repeat Step 1 for switch2.

- Cisco Catalyst 9500-40X

  copy

  Copies from flash of switch1 to flash of switch2.

  ```
  c9500-1#copy flash-1:cat9k_iosxe.17.03.01.SPA.bin flash-2:
  ```

- Cisco Catalyst 9500-48Y4C

  copy

  Copies to bootflash of switch2 from switch1

  ```
  switch1#copy bootflash:cat9k_iosxe.17.03.01.SPA.bin stdby-bootflash:
  cat9k_iosxe.17.03.01.SPA.bin
  ```

Step 3 To remove the startup switch software specification, use the **no** form of the **boot system** command on Catalyst 9500 switches.

a) **config t**

Enters the configuration mode.

b) **no boot system**

Clears all startup software configuration.

Step 4 To configure the switch and reload the copied software, use the following commands:

- Cisco Catalyst 9500-40X

  a. boot system switch all flash

  Configures the boot variable to boot the switch with the newly copied software.

  ```
  c9500-1(config)#boot system switch all flash:
  cat9k_iosxe.17.03.01.SPA.bin
  ```

  b. end

  Exits global configuration mode of the switch

  ```
  c9500-1(config)#end
  ```

  c. wr mem

  Copies the switch configuration changes that you have made and save it to the configuration in flash.

  ```
  c9500-1#wr mem
  ```

- Cisco Catalyst 9500-48Y4C

  a. boot system bootflash

  Installs the upgraded software, saves the configuration, and reloads the copied software.

  ```
  switch1(config)#boot system bootflash:
  cat9k_iosxe.17.03.01.SPA.bin
  ```

  b. end

  Exits global configuration mode of the switch

  ```
  switch1(config)#end
  ```

  c. wr mem

  Copies the switch configuration changes that you have made and save it to the configuration in bootflash.

  ```
  switch1#wr mem
  ```

- Switches without SVL configuration. Configure both the switches to reload the copied software. Use the following commands on both the switches:

  a. boot system flash

  Configures the switches to boot the image from flash memory.

  ```
  Switch(config)#boot system flash:
  cat9k_iosxe.17.03.01.SPA.bin
  ```

  b. end

  Exits global configuration mode of the switch

  ```
  Switch(config)#end
  ```

c. **wr mem**

Copies the switch configuration changes that you have made and save it to the configuration in flash.

```
Switch#wr mem
```

Step 5 To verify only one boot system configuration exists in the running configuration, use the following commands:

a) **show run | i boot**

Verifies that the upgraded software is the first boot image.

Example:

```
c9500-1#show run | i boot
```

b) **license boot level**

Boots a new software license on a switch with the DNA essentials

Example:

```
c9500-1#license boot level network-advantage addon dna-advantage
```

c) **diagnostic bootup level**

Configures the bootup diagnostic level to trigger diagnostics when the switch boots up.

Example:

```
c9500-1#diagnostic bootup level minimal
```

Step 6 To reload and apply the switch configuration change, use the following command. It applies for both Cisco Catalyst 9500-40X and Cisco Catalyst 9500-48Y4C switches.:

Example:

```
c9500-1#reload
```

Supported Upgrade Scenarios and Recommended Connections

The following are the various upgrade scenarios and cluster states that determine the use of prescriptive or flexible connections.

Table 28: Supported Connections

Cisco vManage	Cisco NFVIS	Cluster State	Supported Connections
Upgrade from Releases 19.3 or 20.1.1.1 to Release 20.3.1	Upgrade from Releases 3.12 or 4.1 to Releases 4.1.1 or 4.2.1	Cluster created and active inCisco vManage Releases 19.3 or 20.1.1.1	Use prescriptive connections
Use the latest Release, 20.3.1	Use the latest Release, 4.2.1	Cluster created and active in Cisco vManage Release 20.3.1	Can use prescriptive or flexible connections
Upgrade from Release 20.1.1.1 to Release 20.3.1	Upgrade from Release 4.1 to Releases 4.1.1 or 4.2.1	Cluster created and active in Cisco vManage Release 20.1.1.1.	Use prescriptive connections

Cisco vManage	Cisco NFVIS	Cluster State	Supported Connections
Upgrade from Release 20.1.1.1 to Release 20.3.1	Upgrade from Release 4.1 to Releases 4.1.1 or 4.2.1	Cluster created and active in Cisco vManage Release 20.1.1.1. To add a new Cisco CSP device after upgrade, see Add Cisco CSP Device to Cluster After Upgrading Cisco vManage and Cisco NFVIS.	Use prescriptive connections
Upgrade from Release 20.1.1.1 to Release 20.3.1	Upgrade from Release 4.1 to Releases 4.1.1 or 4.2.1	Cluster created and active in Cisco vManage Release 20.3.1	Can use prescriptive or flexible connections

Add Cisco CSP Device to Cluster After Upgrading Cisco vManage and Cisco NFVIS

To add a Cisco CSP device to a cluster if the cluster was created before upgrading Cisco vManage to Release 20.3.1, perform the following steps:

1. Connect the cables for the newly added Cisco CSP device according to prescriptive connections.
2. Upgrade Cisco NFVIS to Release 4.2.1
3. Use the following commands on the newly added Cisco CSP device by logging into Cisco NFVIS:
   - **request csp-prescriptive-mode**

     Requests the newly added Cisco CSP device to run in prescriptive mode.
   - **request activate chassis-number** *chassis number* **token** *serial number*

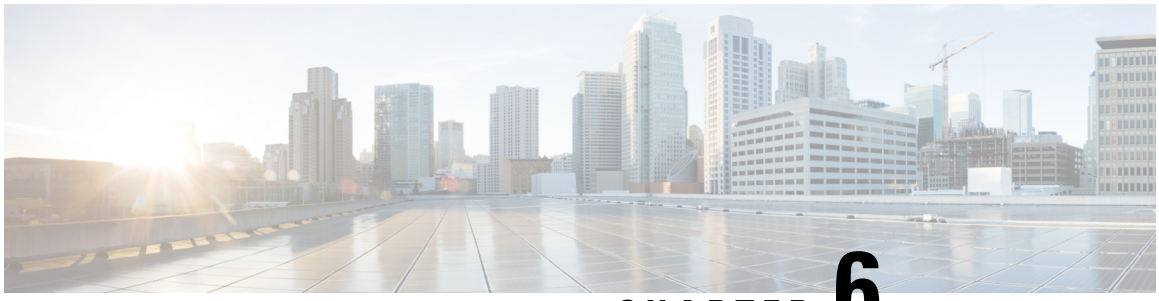
     Activates the Cisco CSP device

     Example

     ```
     request activate chassis-number 71591a3b-7d52-24d4-234b-58e5f4ad0646 token
     e0b6f073220d85ad32445e30de88a739
     ```

Recommendations Prior to Updating a Cluster

- To use an already active cluster when you upgrade to the latest release of the Cisco SD-WAN Cloud onRamp for Colocation solution, ensure that you upgrade Cisco vManage and Cisco NFVIS to the latest releases.
- To create a new cluster when you upgrade to the latest release of the Cisco SD-WAN Cloud onRamp for Colocation solution, ensure that you upgrade Cisco vManage and Cisco NFVIS to the latest releases for flexible connections.



CHAPTER 6

Monitor Cisco SD-WAN Cloud onRamp for Colocation Solution Devices

vManage displays the Cloud OnRamp for Colocation status at a cluster level that indicates the health of each device. The cluster level resources are displayed to indicate the resource availability, such as the CPU allocated and available. You can view service groups in the cluster. All the service groups under a cluster are shown in a table view that indicates the number of VMs in a service chain either up or down. Also, you can view the diagram view of a service group. This diagram view displays all service chains and VMs in a service chain that allows you to look at the resources that are allocated to a VM. Also, it displays VLANs for each VNIC attached to the VM. You can look at the VNF view, which is in tabular form that displays VNF details. You can hover over VM and get information about management IP, CPU, Memory, disk, HA, and type.

The historical and real time operational statistics such as CPU, memory, disk and VNIC utilization charts are available for each VM and CSP. The VNF view can be navigated from a device under the cluster view or from services view. See Monitor Operational Status of Cloud OnRamp for Colocation Devices from vManage, on page 107.

- Monitor Operational Status of Cloud OnRamp for Colocation Devices from vManage, on page 107
- Cisco Colo Manager States for Switch Configuration, on page 117
- Cisco Colo Manager States and Transitions from Host, on page 118
- Cisco Colo Manager Notifications, on page 118
- VM Alarms, on page 121
- VM States, on page 123
- Cloud Services Platform Real-Time Commands, on page 123

Monitor Operational Status of Cloud OnRamp for Colocation Devices from vManage

Monitoring Cloud OnRamp for Colocation devices is the process of reviewing and analyzing a Cloud OnRamp for Colocation device, such as Cloud Services Platform (CSP) devices and Cisco Colo Manager for health, inventory, availability, and other operation-related processes. You can also monitor the components of devices such as CPU, memory, fan, temperature, and so on. For more information about the monitoring dashboard, see the Cisco SD-WAN Configuration Guides configuration guides.

All notifications are sent to the vManage notification stream. To view the notification stream, see Cisco SD-WAN Command Reference.

Procedure

Step 1 Log into Cisco vManage, at the URL HTTPS: //vManage-ip-address/, as a user with "admin" credentials.

Step 2 To view a list of all devices, click **Monitor** > **Network**.

A table lists information about all devices.

If vManage cannot reach CSP devices and Cisco Colo Manager cannot reach switches, the devices are shown as unreachable.

Step 3 To monitor a device, in the **WAN Edge** tab, click a CSP device or Cisco Colo Manager from the list by clicking its hostname.

By default, the System Status screen or Application screen appears. A horizontal bar at the top of the screen with the device drop-down, device name, device IP address, device site location, device model, and more information drop-down is displayed.

Step 4 If necessary, select a different device that you want to monitor.

Step 5 From the **Select Device** bar, click the **More Info** drop-down located to the right of the bar. vManage NMS opens a drop box with a summary information about the device.

In the left pane, the following are the categories of information about the device that you can view:

- VNF Status—view performance specifications, required resources, and component network functions for each VNF See View Information About VNFs from vManage, on page 109.
- Interface—view Interface status and statistics See the "View Interfaces" topic in the Cisco SD-WAN Configuration Guides.
- Control Connections—view status and statistics for control connections See the "View Control Connections" topic in the Cisco SD-WAN Configuration Guides.
- System Status—view reboot and crash information, hardware component status, and CPU and memory usage See the "View Control Connections" topic in the Cisco SD-WAN Configuration Guides
- Colo Manager—view Cisco Colo Manager health status See View Cisco Colo Manager Health from vManage, on page 111.
- Events—view latest syslog events. See the "View Events" topic in the Cisco SD-WAN Configuration Guides.
- Troubleshooting—view information about pings and traceroute traffic connectivity tools See the "Troubleshoot a Device" topic in the Cisco SD-WAN Configuration Guides.
- Real Time—view real-time device information for feature-specific operational commands. See the "View Real-Time Data" topic in the Cisco SD-WAN Configuration Guides.

Step 6 To monitor clusters, click the **Colocation Clusters** tab from the Monitor|Network screen.

All clusters with relevant information are displayed in a tabular format. Click a cluster. See Monitor Cloud OnRamp Colocation Clusters from Cisco vManage , on page 113 for more information.

View Information About VNFs from vManage

Table 29: Feature History

Feature Name	Release Information	Description
VNF States and Color Codes	Cisco SD-WAN Release 20.1.1	This feature allows you to determine the state of a deployed VM using color codes, which you can view on the **Monitor** > **Network** page.

Table 30: Feature History

Feature Name	Release Information	Description
Network Utilization Charts for SR-IOV Enabled NICs and OVS Switch	Cisco SD-WAN Release 20.1.1	This feature allows you to view network utilization charts of VM VNICs connected to both SR-IOV enabled NICs and OVS switch.

You can view performance specifications and required resources for each VNF. Reviewing this information can help you to determine which VNF to use when you are designing a network service. To view information about VNFs, perform the following steps:

Procedure

Step 1 Log into Cisco vManage with "admin" credentials using the URL HTTPS: //vManage-ip-address/.

Step 2 In Cisco vManage, click **Monitor** > **Network**.

The right pane displays VNF information in a tabular format. The table includes information such as CPU use, memory consumption, and disk, and other core parameters that define performance of a network service.

Step 3 Click a CSP device from the table.

Step 4 From the left pane, click **VNF Status**.

Step 5 From the table, click the VNF name. The right pane displays information about the specific VNF. You can click the network utilization, CPU utilization, memory utilization, and disk utilization to monitor the resources utilization of a VNF.

The primary part of the right pane contains the following VNF information:

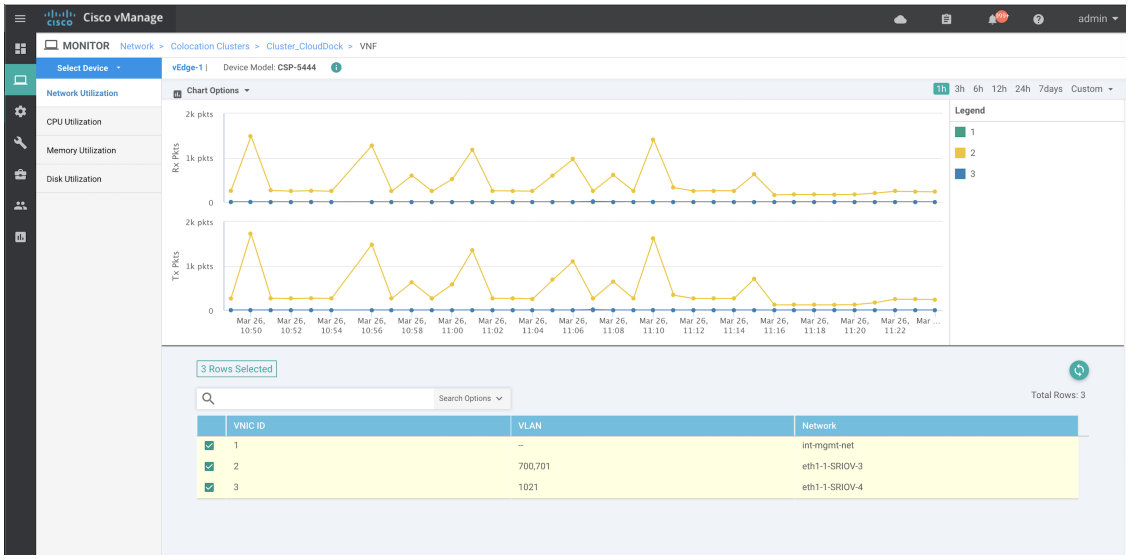
Table 31: VNF Information

Chart options bar	VNF information in graphical format	VNF information in color coded format
• Chart Options drop-down—Click Chart Options drop-down list to select the type of data to display. • Time periods—Click either a predefined time period, or a custom time period for which to display data.	Choose a VNF from the **Select Device** drop-down list to display information for the VNF.	The VNFs are shown in specific colors based on the following operational status of the VNF life cycle: • Green—VNF is healthy, deployed, and successfully booted up. • Red—VNF deployment or any other operation fails, or VNF stops. • Yellow—VNF is transitioning from one state to another.

The detail part of the right pane contains:

- Filter criteria
- VNF table that lists information about all VNFs or VMs. By default, the first six VNFs are checked. The network utilization charts for VNICs connected to SR-IOV enabled NICs and OVS switch are displayed.

Figure 34: VNF Information



The graphical display plots information for the checked VNFs

- Click the checkbox at the left to select and deselect VNFs. You can select and display information for a maximum of six VNFs at one time.
- To change the sort order of a column, click the column title.

View Cisco Colo Manager Health from vManage

You can view Cisco Colo Manager health for a device, Cisco Colo Manager host system IP, Cisco Colo Manager IP, and Cisco Colo Manager state. Reviewing this information can help you to determine which VNF to use when you are designing a network service. To view information about VNFs, perform the following steps:

Procedure

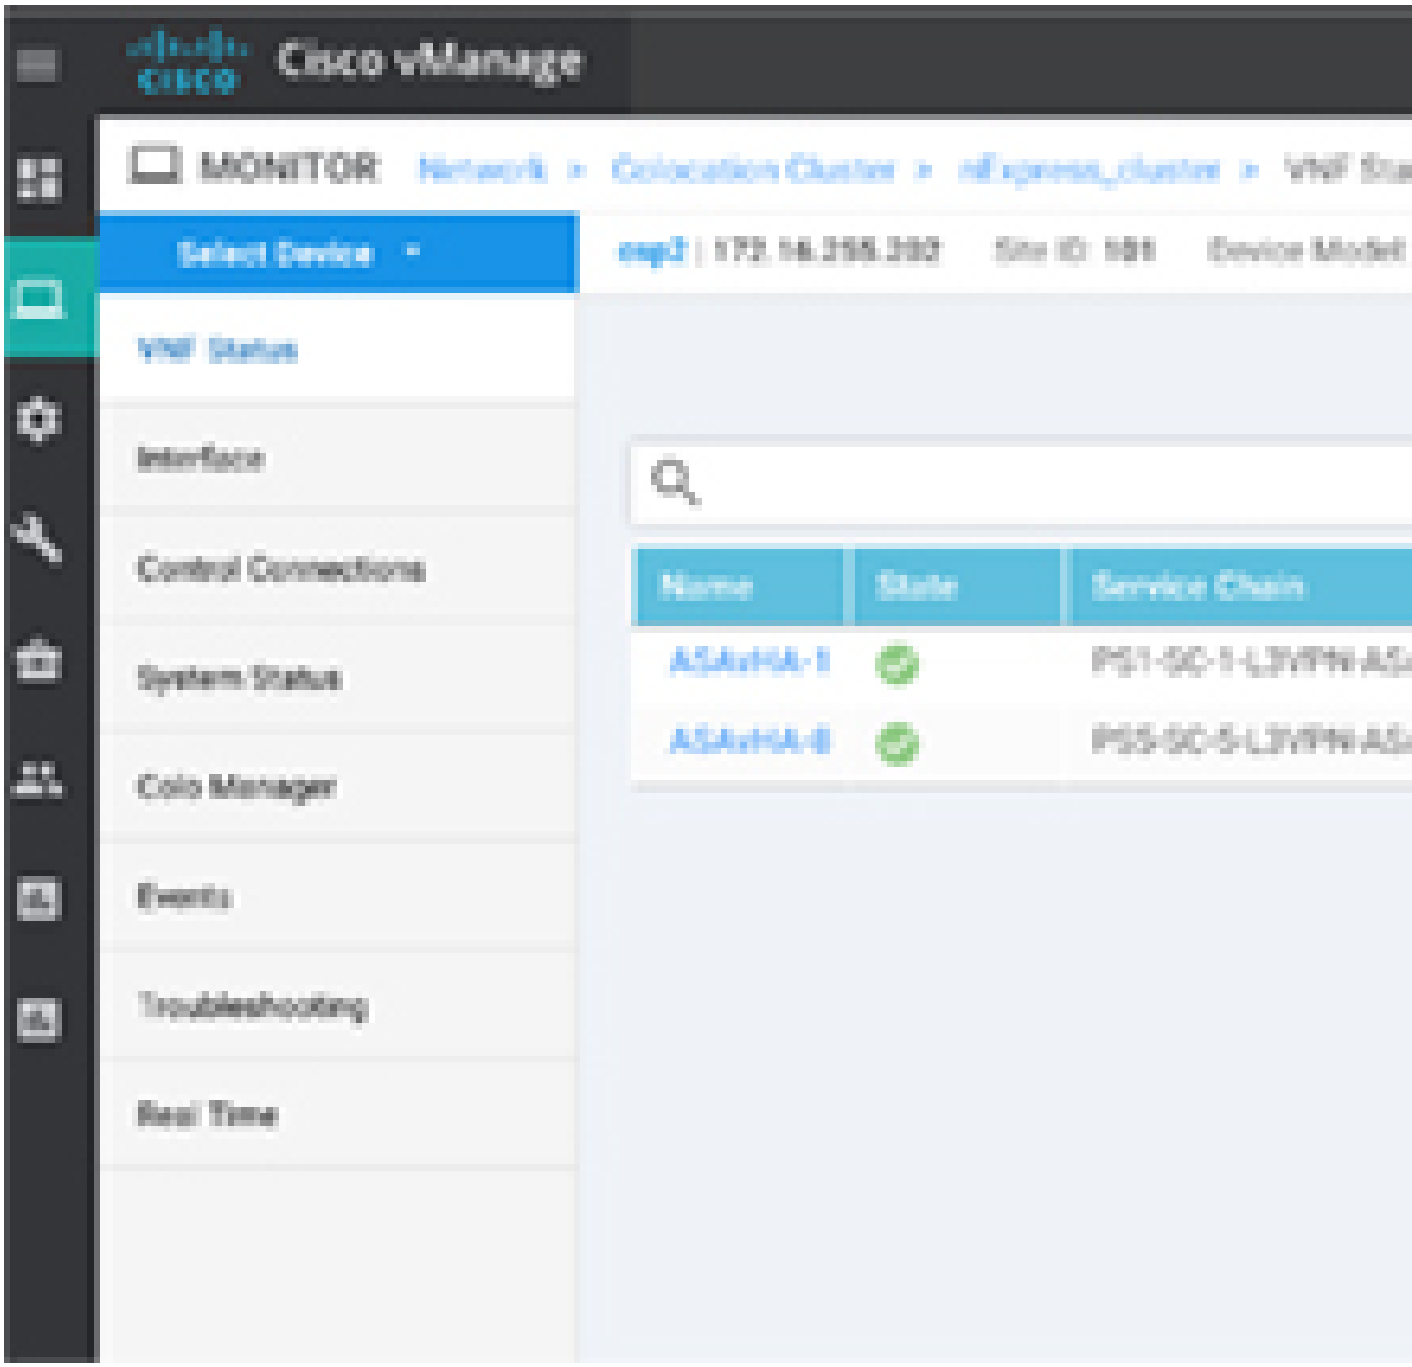
Step 1 Log into Cisco vManage, at the URL HTTPS: //vManage-ip-address/, as a user with "admin" credentials.

Step 2 In vManage, click **Monitor** > **Network**.

The right pane displays VNF information in a tabular format. The table includes information such as CPU use, memory consumption, and disk, and other core parameters that define performance of a network service.

Step 3 Click a CSP device from the table.

Step 4 From the left pane, click **Colo Manager**.



The right pane displays information about the memory usage, CPU usage, uptime, and so on, of the colo manager.

Monitor Cloud OnRamp Colocation Clusters from Cisco vManage

Table 32: Feature History

Feature Name	Release Information	Description
Network Assurance –VNFs: Stop/Start/Restart	Cisco SD-WAN Release 20.3.1 Cisco vManage Release 20.3.1	This feature provides the capability to stop, start, or restart VNFs on Cisco CSP devices from the **Colocation Clusters** tab. You can easily perform the operations on VNFs using Cisco vManage.

You can view the cluster information and their health states. Reviewing this information can help you to determine which Cisco CSP device is responsible for hosting each VNF in a service chain. To view information about a cluster, perform the following steps:

Procedure

Step 1 Log in to Cisco vManage with "admin" credentials using the URL HTTPS: //vManage-ip-address/.

Step 2 In Cisco vManage, click **Monitor** > **Network**.

Step 3 To monitor clusters, click **Colocation Clusters**.

All clusters with relevant information are displayed in a tabular format. Click a cluster name. You can monitor cluster from the **Config. View** and **Port Level View**.

- **Config. View**: The primary part of the left pane displays the CSP devices and switches that form the cluster. In the right pane, you can view the cluster information such as the available and total CPU resources, available and allocated memory, and so on, based on colocation size.

  The detail part of the left pane contains:

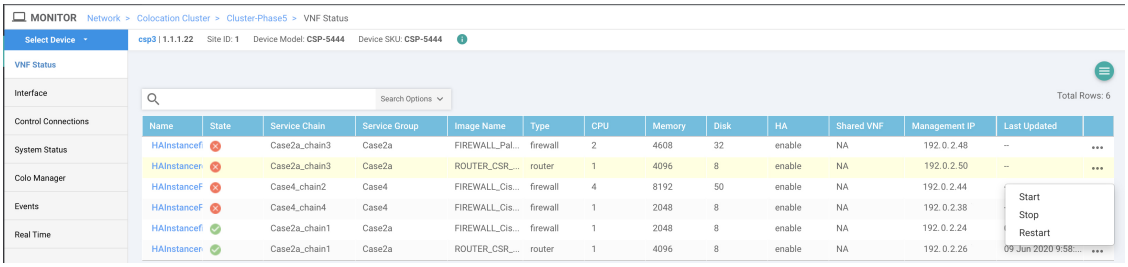
  - Filter criteria: Type a field value to filter and display the information.
  - A table that lists information about all devices in the cluster (Cisco CSP devices, PNFs, and switches).

    Click a Cisco CSP device. VNF information is displayed in tabular format. The table includes information such as VNF name, service chains, number of CPUs, memory consumption, and other core parameters that define performance of network service. See View Information About VNFs from vManage, on page 109 .

    To start, stop, or reboot a VNF, click **More Action** next to the VNF name, and choose one of the following operations:

    - **Start**.
    - **Stop**.
    - **Restart**.

Figure 35: Start, Stop, or Restart VNFs



Note
Ensure that service chain provisioning has been completed and VMs are deployed, before issuing start, stop, restart operations on any of the VNFs in that chain.

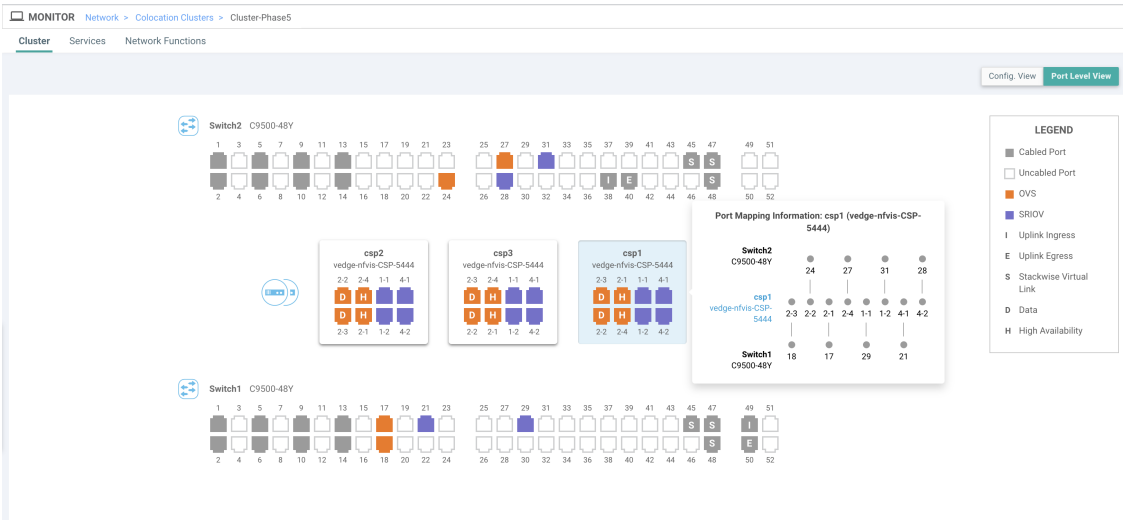
After you choose an operation on a VNF, wait until the operation is complete before you issue another operation. You can view the progress of an operation from the task view window.

- **Port Level View**: After you activate the cluster, to view the port connectivity details, click **Port Level View**.

  You can view detailed port connectivity information for the switches and CSP devices in a color coded format based on SR-IOV and OVS modes.

  To view the mapping of ports between the Catalyst 9500 switches and CSP devices, click or hover over a CSP device.

Figure 36: Monitor Port Connectivity Details of a Cluster



Step 4 Click **Services**.

In this tab, you can view:

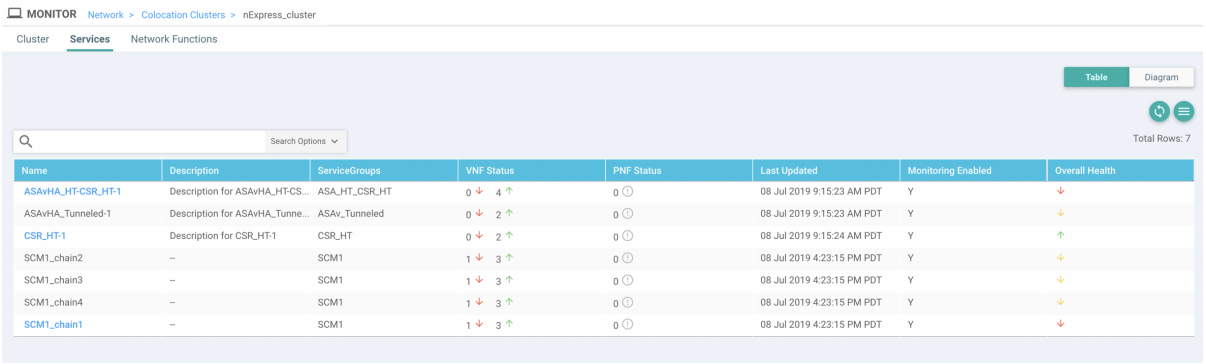
- Complete information of a service chain. The first two columns display the name and description of the service chain within the service group and the remaining columns mention about the VNF, PNF statuses, monitoring service enablemement, and the overall health of a service chain. You can also view the colocation user group associated with a service chain. The various health statuses and their representations are:

- Healthy—An up arrow in green. A service chain is in 'Healthy' status when all the VNF, PNF devices are running and are in healthy state. Ensure that you configure the routing and policy correctly.
- Unhealthy—A down arrow in red. If one of the VNFs or PNFs are in unhealthy state, the service chain is reported to be in 'Unhealthy' status. For example, after deploying a service chain, if one of the network function IP address changes on the WAN or LAN side, or the firewall policy isn't configured to let the traffic pass through, then unhealthy state is reported. This is because the network function or overall service chain is Unhealthy or both are in Unhealthy state.
- Undetermined—Down arrow in yellow. This state is reported when the health of the service chain can't be determined. This state is also reported when there's no status such as healthy or unhealthy available for the monitored service chain over a time period. You can't query or search a service chain with undetermined status.

  If a service chain consists of a single PNF and PNF is orchestrated outside of Cisco vManage, it can't be monitored. If a service chain consists of a single network function, firewall that has VPN termination on both sides which can't be monitored, then it's reported as Undetermined.

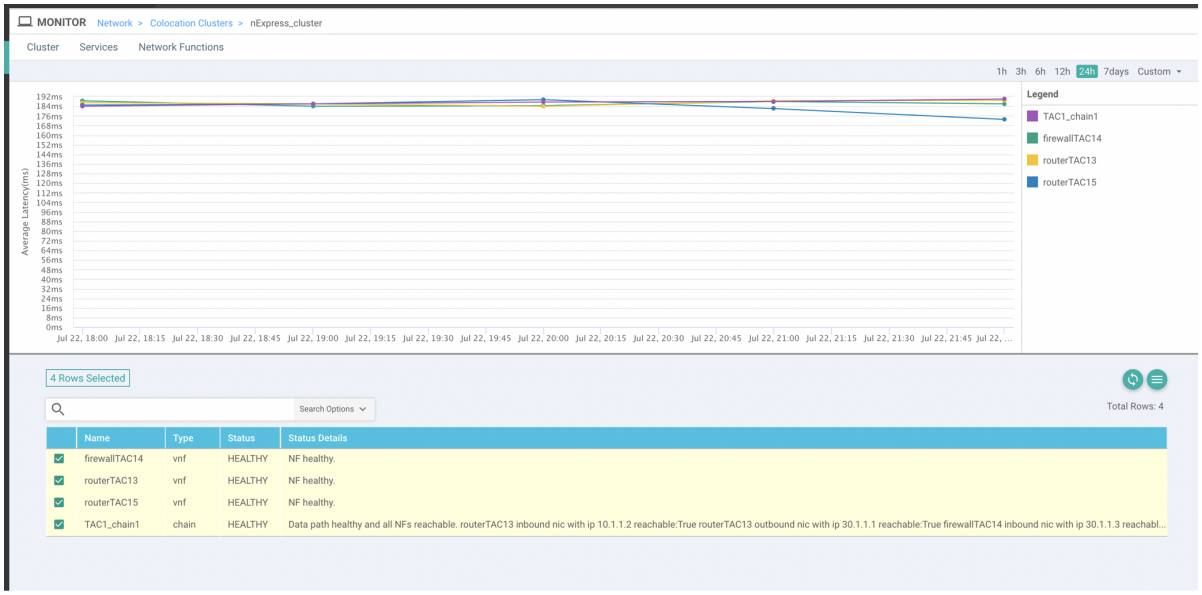
  Note
  If the status of a service chain is undetermined, you can't choose the service chain to view the detailed monitoring information.

Figure 37: Service Chain Health Monitoring Results



- If you had configured a service chain by enabling the monitoring field, then click a service group that is in Healthy or Unhealthy state. The primary part of the service chain monitoring in the right pane contains the following elements:

Figure 38: Service Chain Health Monitoring Status



Graphical display that plots the latency information of the service chain, VNFs, PNFs.

The detail part of the right pane contains:

- Filter criteria
- A table that lists information about all service chains, VNFs, PNFs, their health status, and types.
  - Check the checkbox at the left of a row to select and deselect a service chain, VNF, PNF.
  - To change the sort order of a column, click the column title.

In the following image, the status details column indicates the monitored data path and it provides the per hop analysis.

- Click the **Diagram** button and view the service group with all its service chains and VNFs in the design view window.
- Click a VNF. You can view CPU, memory, and disk allocated to the VNF in a dialog box.
- Select a service group from the **Service Groups** drop-down. The design view displays the selected service group with all its service chains and VNFs.

Step 5 Click **Network Functions**.

In this tab, you can view:

- All the virtual or physical network functions in a tabular format. On the **Show** button, you can choose to display either a VNF or PNF.

  VNF information is displayed in a tabular format. The table includes information such as VNF name, service chains, colocation user groups, CPU use, memory consumption, and other core parameters that define performance of network service. To view more information about the VNF, click a VNF name. See View Information About VNFs from vManage, on page 109 .

- PNF information is displayed in tabular format. The table includes information such as the serial number and PNF type. To view and note configuration of a specific PNF, click the desired PNF serial number. Ensure that you manually note all the configuration of the PNFs and then configure the PNF devices. For example, the following are some of the PNF configuration where you position the PNF at various locations in the service chain. see the ASR 1000 Series Aggregation Services Routers Configuration Guides and Cisco Firepower Threat Defense Configuration Guides to configure the PNFs manually.

Figure 39: PNF in the First Position with Service Chain Side Parameters

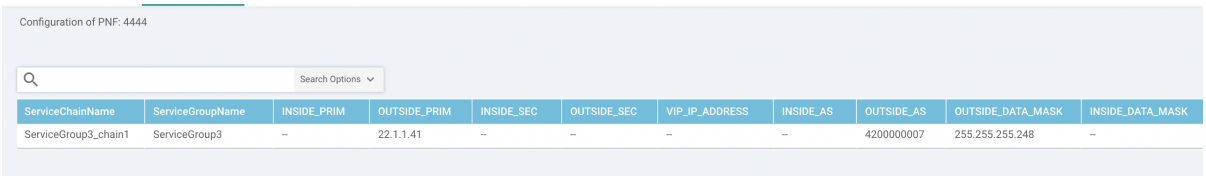


Figure 40: PNF in the First Position with Outside Neighbor Information

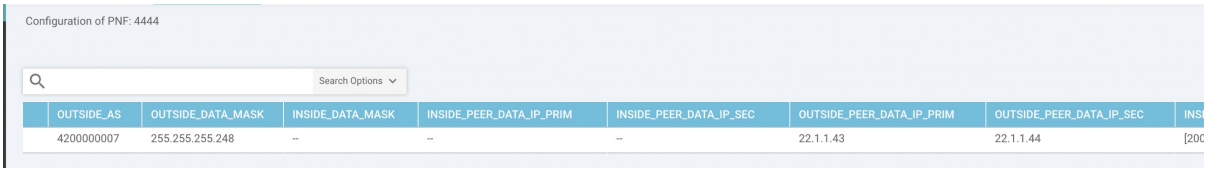


Figure 41: PNF Shared Across Two Service Chains

The ServiceGroup2_chain3 is a PNF-only service chain and therefore no configuration gets generated. The PNF is in the last position of the ServiceGroup2_chain1, so only INSIDE variables gets generated.



Figure 42: PNF Shared Across Two Service Chains with Outside Neighbor Information



Cisco Colo Manager States for Switch Configuration

The various Cisco Colo Manager states and transitions when you trigger various processes from vManage are:

- INIT state—when the Cisco Colo Manager container is successfully initialized.

- IN-PROGRESS state—when any configuration push is not possible.
- SUCCESS state—when the Cisco Colo Manager container has successfully translated and pushed the intent that is received from vManage to Cisco Catalyst 9500-40X or Cisco Catalyst 9500-48Y4C switches.
- FAILURE state—If there is any failure in processing or configuration push in Cisco Colo Manager.

When vManage pushes the Cloud OnRamp for Colocation configuration intent to the Cisco Colo Manager for the first time it moves from INIT to IN-PROGRESS state. After Cisco Colo Manager pushes the configuration, it goes back to the SUCCESS or FAILURE state. For every incremental configuration push, it goes to IN-PROGRESS state. If any of the configurations pushes fail, Cisco Colo Manager goes into FAILURE state.

Note A notification is sent when Cisco Colo Manager state changes. See Cisco Colo Manager Notifications, on page 118.

Cisco Colo Manager States and Transitions from Host

vManage depends on various CSP hosts state for the Cisco Colo Manager to be brought up, which are:

- Starting—when Cisco Colo Manager has been brought up and health check script has not been run. During this phase, vManage waits for CSP state to change to Healthy.
- Healthy—when the health check script has been run and it has passed the checks. This state implies that the operational model for configuration status can be queried or configuration can be pushed. During this phase, if Cisco Colo Manager is in INIT state, vManage pushes the device list. If Cisco Colo Manager is not in INIT state, Cloud OnRamp for Colocation may be in degraded state and recovery flow must happen.
- Unhealthy—when all the necessary packages in Network Services Orchestrator (NSO) are not up. This state can be due to various reasons such as, NSO did not come up, Cisco Colo Manager package did not come up, or other reasons. This state implies that the operational model for configuration status is not up and configuration cannot be pushed.

Cisco Colo Manager Notifications

You can view the Cisco Colo Manager notifications from Cisco Colo Manager console by using the **show notification stream viptela** command.

The various Cisco Colo Manager internal state machines are:

Table 33: CCM Notifications

Cisco Colo Manager States	Notification Trigger	Notification Output Example
INIT	Init: Cloud OnRamp for Colocation is activated and vManage brings up Cisco Colo Manager on Cisco CSP. **Note** The Cisco Colo Manager state must be in "Init" only when the docker container is initially brought up and must not be in this state unless container is deleted and brought up again.	admin@ncs# show notification stream viptela last 50 notification eventTime 2019-04-08T17:15:15.982292+00:00 ccmEvent severity-level minor host-name ccm user-id vmanage_admin config-change false transaction-id 0 status SUCCESS status-code 0 status-message init details Initializing CCM event-type CCM-STATUS ! !
IN-PROGRESS	vManage pushes intent and Cisco Colo Manager moves to in-progress state. **Note** Cisco Colo Manager generates multiple in-progress notifications for the switches that are brought up.	notification eventTime 2019-04-08T17:37:54.536953+00:00 ccmEvent severity-level minor host-name ccm user-id vmanage_admin config-change false transaction-id 0 status SUCCESS status-code 0 status-message IN-PROGRESS details Received configuration from vManage event-type CCM-STATUS !

Cisco Colo Manager States	Notification Trigger	Notification Output Example

SUCCESS	During cluster activation, after Cisco Catalyst 9500 switches have been successfully onboarded, status moves to SUCCESS. For any incremental configuration, status moves to SUCCESS only if configuration has been saved successfully in the switch devices.	`notification`
		` eventTime 2019-04-08T17:51:48.044286+00:00`
		` ccmEvent`
		`  severity-level minor`
		`  host-name ccm`
		`  user-id vmanage_admin`
		`  config-change false`
		`  transaction-id 0`
		`  status SUCCESS`
		`  status-code 0`
		`  status-message SUCCESS`
		`  details Devices done onboarding`
		`  event-type CCM-STATUS`
		` !`
		`!`
		`admin@ncs#`

Cisco Colo Manager States	Notification Trigger	Notification Output Example
FAILURE	If onboarding of switches fails during cluster activation, CCM status moves to FAILURE. If any incremental configurations are not saved, CCM status moves to FAILURE. **Note** The failure state cannot transition to another state without end-user intervention.	notification eventTime 2019-04-08T18:01:44.943198+00:00 ccmEvent severity-level critical host-name ccm user-id vmanage_admin config-change false transaction-id 0 status FAILURE status-code 0 status-message FAILURE details SVL bringup not successful. Could not sync TenGigabitEthernet2/0/* interfaces. event-type CCM-STATUS ! ! admin@ncs#
	Onboarding of switches fails during cluster activation due to wiring errors in flexible connections, and CCM status moves to FAILURE.	admin@ncs# show notification stream viptela last 100 | include Step notification details Step 5 of 7: Device switch1 : 192.168.100.21 (C9500-48Y4C-CAT2324L2HM) connected after SVL reload. details Step 6 of 7: Started sync-from for primary device switch1 : 192.168.100.21 (C9500-48Y4C-CAT2324L2HM) details Step 6 of 7: Sync-from done for primary device switch1 : 192.168.100.21 (C9500-48Y4C-CAT2324L2HM) Device list : switch1 : 192.168.100.21 (C9500-48Y4C-CAT2324L2HM), switch2 : 192.168.100.19 (C9500-48Y4C-CAT2316L2F2) details Step 6 of 7: Devices ready for LLDP query Device list : switch1 : 192.168.100.21 (C9500-48Y4C-CAT2324L2HM), switch2 : 192.168.100.19 (C9500-48Y4C-CAT2316L2F2) details Step 6.1 of 7: LLDP Query Details: csp2 has 8/8 interfaces connected, 2/4 sriov, 2/4 fortville to primary switch; 2/4 sriov, 2/4 fortville to secondary switch; Found devices with not optimum connections:- csp1 has 6/8 interfaces connected, 2/4 sriov, 2/4 fortville to primary switch; 2/4 sriov, 0/4 fortville to secondary switch; Minimum Requirement is to have 8/8 interfaces per CSP in cluster. Recommended action: Please refer to recommended topologies and minimum requirements details Step 7 of 7: Devices done onboarding Device list : switch1 : 192.168.100.21 (C9500-48Y4C-CAT2324L2HM), switch2 : 192.168.100.19 (C9500-48Y4C-CAT2316L2F2)

VM Alarms

The following are VM alarms and they can be viewed from vManage, when alarms are received.

Table 34: Alarms

Alarm	Trigger Condition	Syslog Messages
INTF_STATUS_CHANGE	interface status change	nfvis %SYS-6-INTF_STATUS_CHANGE: Interface eth0, changed state to up
VM_STOPPED	vm stopped	nfvis %SYS-6-VM_STOPPED: VM stop successful: SystemAdminTena_ROUTER_0_df6733c1- 0768-4ae6-8dce-b223ecdb036c
VM_STARTED	vm started	nfvis %SYS-6-VM_STARTED: VM start successful: SystemAdminTena_ROUTER_0_df6733c1- 0768-4ae6-8dce-b223ecdb036c
VM_REBOOTED	vm rebooted	nfvis %SYS-6-VM_REBOOTED: VM reboot successful: SystemAdminTena_ROUTER_0_df6733c1- 0768-4ae6-8dce-b223ecdb036c
VM_RECOVERY_INIT	vm recovery initiation	nfvis %SYS-6-VM_RECOVERY_INIT: VM recovery initiation successful: SystemAdminTena_ROUTER_0_df6733c1- 0768-4ae6-8dce-b223ecdb036c
VM_RECOVERY_REBOOT	vm recovery reboot	nfvis %SYS-6-VM_RECOVERY_REBOOT: VM recovery reboot successful: SystemAdminTena_ROUTER_0_df6733c1- 0768-4ae6-8dce-b223ecdb036c
VM_RECOVERY_COMPLETE	vm recovery complete	nfvis %SYS-6-VM_RECOVERY_COMPLETE: VM recovery successful: SystemAdminTena_ROUTER_0_df6733c1- 0768-4ae6-8dce-b223ecdb036c
VM_MONITOR_UNSET	vm monitoring unset	nfvis %SYS-6-VM_MONITOR_UNSET: Unsetting VM monitoring successful: SystemAdminTena_ROUTER_0_df6733c1- 0768-4ae6-8dce-b223ecdb036c
VM_MONITOR_SET	vm monitoring set	nfvis %SYS-6-VM_MONITOR_SET: Setting VM monitoring successful: SystemAdminTena_ROUTER_0_df6733c1- 0768-4ae6-8dce-b223ecdb036c

See Cisco NFVIS Configuration Guide for more information about syslog support and VM alarms.

VM States

The following are the operational status of deployed VM life cycle. In Cisco SD-WAN, you can view and monitor the VM states from Cisco vManage.

Table 35: VM States

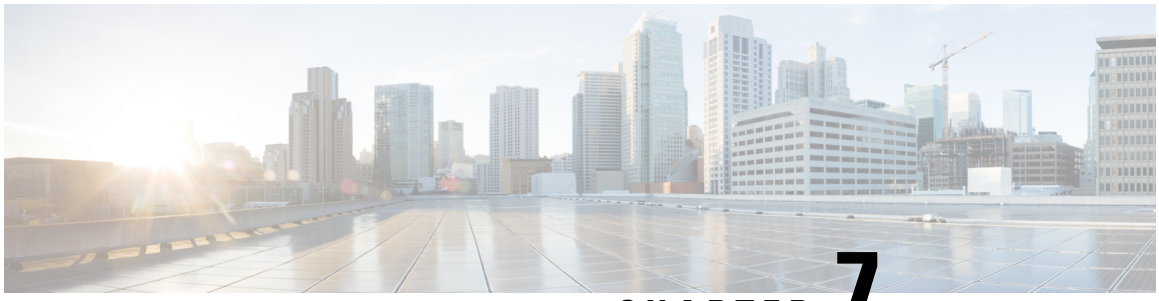
VM States	Description
VM_UNDEF_STATE	VM or VNF is transitioning from one state to another.
VM_INERT_STATE	VM or VNF is deployed but not alive.
VM_ALIVE_STATE	VM or VNF is deployed and successfully booted up or alive.
VM_ERROR_STATE	VM or VNF is in error state when deployment or any other operation fails.

Cloud Services Platform Real-Time Commands

Table 36: Real-Time Commands

System Information
Container status
show control connections
Control connection history
Control local properties
Control summary
Control statistics
Control valid vEdges
valid vManage ID
HW Alarms
HW Environments
PNICs
System Status

Host System Mgmt Info
Host System settings
Host System processes
Resource CPU allocation
RBAC Authentication
Resource CPU VNFs
Hardware Inventory
Hardware Temperature thresholds
Control affinity stats



CHAPTER 7

High Availability

The Cisco SD-WAN Cloud onRamp for Colocation solution allows various consumers to access various repetitive applications securely. The Cisco SD-WAN Cloud onRamp for Colocation solution High Availability (HA) is designed to handle several types of failure possible in a cluster deployment. The following types of failures can occur in a Cisco SD-WAN Cloud onRamp for Colocation solution deployment:

- Compute failure
- Switch failure
- Service chain failure

To resolve the failures, use the following mechanisms:

- Redundancy
- Failure detection

- Redundancy, on page 125
- Handle Various Failure Scenarios, on page 128

Redundancy

The following are the components where redundancy has been added to address failure of the component:

- x86 Compute Hardware—See Redundancy of x86 Compute Hardware, on page 126.
- Network Fabric—See Redundancy of Network Fabric, on page 126.
- Physical NIC/interface—See Redundancy of Physical NIC or Interface, on page 126.
- NFVIS Virtualization Infrastructure—See Redundancy of NFVIS, Virtualization Infrastructure, on page 126.
- Service-Chain/VNF—See Redundancy of Service Chain or VNF, on page 126.
- Cisco Colo Manager—See Recovery of Cisco Colo Manager, on page 127.

Redundancy of Network Fabric

Network Fabric—The hardware switch redundancy features are used to handle network fabric failures. In a switch failure, ensure that the standby switch takes over the traffic traversing through the failed switch.

Redundancy of x86 Compute Hardware

x86 Compute Hardware—Any hardware components such as, processor, storage, and others that are used on the x86 compute hardware can fail leading to a complete CSP system failure. The Cisco vBond orchestrator continuously monitors the health of the x86 compute platform by using ICMP ping through the management interface. In a system failure, the orchestrator shows the device status and the service chains and VMs impacted. Take desired action to bring up service chains. See Monitor Cisco SD-WAN Cloud onRamp for Colocation Solution Devices, on page 107. Depending on the operational status of the VNFs, the VMs must be brought up on a different CSP if enough are resources are available. This action allows the VNF to retain the Day-N configuration. If the VNF disk is using local storage, the entire service group must be respun on another CSP device with the Day-0 configuration that is stored in the orchestrator.

Redundancy of Physical NIC or Interface

Physical NIC or interface—If a physical NIC or interface or cable fails or gets disconnected, the VNFs that are using those interfaces are impacted. If a VNF is using an OVS network, the port channel configuration is used to achieve a link redundancy. If a VNF is using an OVS network, and if the VNF has an HA instance, that instance has been already brought up on a different CSP. The failover happens to this VNF on the second CSP. If there is no second VNF instance, the service chain with the failed VNF must be deleted and reinstantiated.

Redundancy of NFVIS, Virtualization Infrastructure

NFVIS Virtualization Infrastructure—Multiple types of failures in the NFVIS software layer can occur. One of the critical components of CSP can crash or the host Linux kernel can panic or one of the critical components fails to respond. In case of critical component failures, the NFVIS software generates netconf notifications. The orchestrator uses these notifications to show the failure on vManage dashboard. If the CSP or NFVIS crashes or control connection goes down, the orchestrator shows that device reachability is down. You can resolve a networking issue (if any), or reboot the CSP device. If device does not recover, you must proceed with removing the CSP device.

Redundancy of Service Chain or VNF

Table 37: Feature History

Feature Name	Release Information	Description
HA VNF NIC Placement for Switch Redundancy	Cisco SD-WAN Release 20.5.1 Cisco vManage Release 20.5.1	This feature provides an optimum placement of service chains and therefore maximizes the resource utilization while accounting for switch redundancy. The VNICs of the HA primary and secondary instances are placed on alternate switch interfaces to achieve redundancy at switch level.

Service Chain or VNF—Some of the VNFs in the colocation service chain such as, firewall might support stateful redundancy features by using a standby VNF, whereas VNFs such as Cisco CSR1000V might not support stateful redundancy. The Cisco SD-WAN Cloud onRamp for Colocation solution relies on the VNFs to achieve VNF high availability. The HA support at service chain level isn't available. If a VNF supports stateful HA, it detects the failure and performs a switchover. The assumption is that the previously active VNF goes down and reboots as a standby VNF if the CSP device hosting the VNF is functional, and all the NIC or interface connectivity is functional. If the VNF doesn’t come up, the HA for VNF isn’t functional from that time and your intervention is required.

If a VNF doesn’t support HA, it’s assumed that the VNF reboots if any critical process fails within the VNF and no HA support is available for such VNFs.

Starting from Cisco SD-WAN Cloud onRamp for Colocation Release 20.5.1, the service chain deployments continue seamlessly during switch failures without traffic getting dropped. The traffic movement happens due to the placement of VNICs of the HA switch instance on the CSP NICs. Unlike the VNICs of the primary switch instance, the VNICs of the HA switch instance aren’t connected to the switch in the same way.

For example, in the following image, the primary of the first VNF and second VNF are connected to the SR-IOV ports, and these ports are connected to the first switch instance. The secondary of the first VNF and second VNF are now connected to different SR-IOV ports, and these ports are connected to the second switch instance. If the first switch instance fails, the traffic continues to flow from the secondary of the first VNF and second VNF through the second switch instance.

The placement of VNICs on the CSP NICs applies only to single-tenant clusters, where the VNFs use SR-IOV interfaces. If a VNF uses OVS interface, the port channel configuration is used to achieve link redundancy.

We recommend that you consider the following for optimal utilization of resources when attaching service chains:

- Design as many service chains as possible and provision these chains so that you use all resources to the maximum capacity. This task ensures that the bandwidth isn't fragmented.
- Attach high-bandwidth service chains first followed by the low-bandwidth service chains. Similarly, attach highly available service chains first followed by the stand alone service chains for optimal resource utlization.

Recovery of Cisco Colo Manager

Cisco Colo Manager Recovery—Cisco Colo Manager is brought up on a CSP device in a Cloud OnRamp for Colocation. vManage selects a CSP with the DTLS tunnel to bring up Cisco Colo Manager. The Cisco Colo Manager recovery flow is required during the following scenario:

If a CSP hosting Cisco Colo Manager is considered for Return Material Authorization (RMA) process and there are at least two other CSP devices in the cluster after deleting this CSP, then a new Cisco Colo Manager is brought up automatically by vManage on one of the existing two CSP devices during a new configuration push.



Note You must power down the CSP device that has been considered for RMA process or perform a factory default reset on the CSP device. This task ensures that there is only one Cisco Colo Manager in the cluster.

Note A host with Cisco Colo Manager running can restart or reboot, and this action is not a recovery scenario as Cisco Colo Manager should come up intact with all the configuration and operational data.

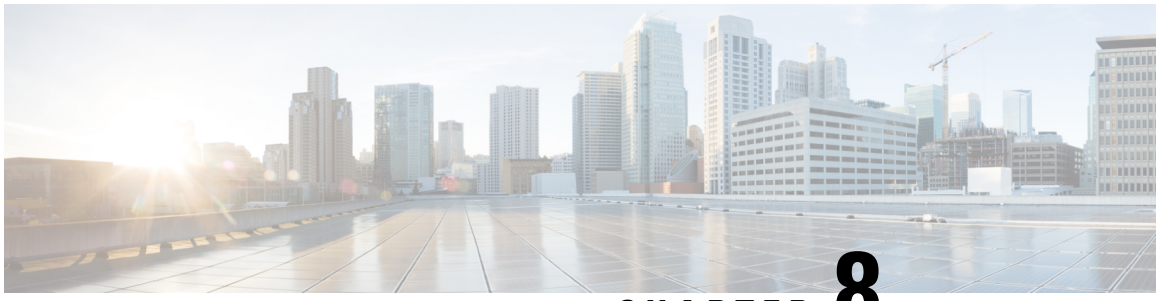
If after a cluster is successfully activated and then Cisco Colo Manager becomes unhealthy, see Troubleshoot Cisco Colo Manager Issues, on page 153.

Handle Various Failure Scenarios

- VNF failure
  - If a VM in a service chain that is HA capable goes down, the standby VM takes over. This standby service chain is functional within few seconds. NFVIS software on CSP tries to bring up the failed active VM if it is a monitored VM. If the VM recovers successfully, it switches over to active and standby modes successfully. If VM did not recover successfully and you want to bring up HA capability on this VM, tear down the service chain and bring up new service chain with HA capability. Here, VM detects that the failure is based on heartbeat and there must not be any impact on traffic (except few seconds). If an active VM recovers, this VM could become active again or stay as standby and this state varies from VM to VM.
  - If a VM is not HA capable, the service chain fails and traffic is black holed. Cisco Colo Manager detects this failure and hence vManage as it receives notification that VM is down and service chain is down, vManage sends an alert. If the VM recovers successfully, the same notification is sent and the service chain is functional without any intervention. If the VM does not recover successfully, tear down the service chain and bring up a new service chain.
- Service chain failure
  - If all VMs in a service chain support HA, service chains can have active and standby service chains. If an active service chain goes down, the standby service chain takes over and is functional within few seconds. This behavior is VM level HA and VM failover behavior takes over. NFVIS software on CSP also tries to bring up the failed active VMs (for monitored VMs) and if they recover successfully, the VMs switch over to active and standby modes successfully.
  - If VMs are not HA capable, the service chain fails and traffic is black holed. NFVIS and Cisco Colo Manager send notifications that VMs are down and vManage sends an alert. Based on the notification, bring up another active service chain. If the service chain has recovered successfully, the same notification is sent and the service chain is functional without any intervention.
- Device failure

  If a CSP is down, all the service chains and VMs running on that CSP are also down. Cisco Colo Manager sends notifications to vManage that device is not reachable and vManage itself must detect DTLS connectivity loss with the CSP device. vManage sends alert about the CSP device and you must bring up the service chains on another CSP by creating the service chains and pushing the configuration to a colocation. If there is not enough compute hardware, add another CSP to a colocation and push the service chain configuration to the CSP.
- Switch link failure

  If a link from a switch is down, the other switch takes over and service chain traffic continues.



CHAPTER 8

Cisco SD-WAN Cloud onRamp for Colocation Multitenancy

Table 38: Feature History

Feature Name	Release Information	Description
Colocation Multitenancy Using Role-Based Access Control	Cisco SD-WAN Release 20.5.1 Cisco vManage Release 20.5.1	This feature enables a service provider to manage multiple colocation clusters and share these clusters across tenants by using multiple colocation groups. In a multitenant setup, service providers don't need to deploy a unique colocation cluster for each tenant. Instead, the hardware resources of a colocation cluster are shared across multiple tenants. With multitenancy, service providers ensure that tenants view only their data by restricting access based on roles of individual tenant users.

- Overview of Colocation Multitenancy, on page 129
- Roles and Functionalities in Multitenant Environment, on page 130
- Recommended Specifications in Multitenant Environment, on page 131
- Assumptions and Restrictions in Colocation Multitenancy, on page 132
- Service Provider Functionalities, on page 133
- Tenant Functionalities, on page 137
- Monitor in Comanaged Multitenant Environment, on page 138

Overview of Colocation Multitenancy

In Cisco SD-WAN Cloud onRamp for Colocation multitenancy, a service provider can manage multiple colocation clusters using Cisco vManage in single-tenant mode. A service provider can bring up a multitenant cluster in the same way as bringing up a cluster in a single-tenant mode. A multitenant cluster can be shared across multiple tenants. See Create and Activate Clusters, on page 38.

The tenants share the hardware resources such as the Cisco Cloud Services Platform (CSP) devices and Cisco Catalyst 9500 devices of a colocation cluster. The following are the key points of this feature.

- A service provider deploys and configures the Cisco SD-WAN Controllers (Cisco vManage, Cisco vBond Orchestrator, and Cisco vSmart Controller) with valid certificates.

- A service provider sets up colocation clusters after onboarding the Cisco CSP devices and Cisco Catalyst 9500 switches.
- Cisco SD-WAN operates in a single-tenant mode and Cisco vManage dashboard appears in a single-tenant mode.
- In a colocation multitenant deployment, a service provider ensures that tenants see only their service chains by, creating roles. A service provider creates roles for each tenant in a colocation group. These tenants are permitted to access and monitor the service chains based on their roles. However, they can't configure their service chains or change the system-level settings. The roles ensure that tenants can access only the information that they are authorized to view.
- Each tenant traffic is segmented using VXLAN across the compute devices, and VLAN across the Cisco Catalyst switch fabric.
- A service provider can provision service chains on a specific cluster.

The following are the two scenarios of a colocation multitenant setup:

- Provider owned Cisco SD-WAN devices: In this scenario, the Cisco SD-WAN devices used in a service chain belong to the service provider. The CSP devices and Catalyst 9500 switches are owned, monitored, maintained by the service provider. The VM packages are owned, uploaded, and maintained by the provider. See Monitor in Comanaged Multitenant Environment, on page 138.
- Comanaged Cisco SD-WAN devices: In this scenario, the Cisco SD-WAN devices that are used in a service chain belong to a tenant overlay network. The colocation cluster devices are owned by the service provider, whereas the Cisco SD-WAN devices of a service chain are controlled by Cisco SD-WAN controllers (Cisco vManage, Cisco vBond Orchestrator, Cisco vSmart Controller) of a tenant. The CSP devices and Catalyst 9500 switches are owned, monitored, maintained by the service provider. The VM packages are owned, uploaded, and maintained by a provider. See Monitor in Comanaged Multitenant Environment, on page 138.

Roles and Functionalities in Multitenant Environment

Multitenant environments include a service provider and multiple tenants. Each role has distinct responsibilities and associated functions.

Service Provider

The service provider owns all hardware infrastructure and manages the clusters. The service provider onboards tenants by creating their roles, provisions the service chains for tenants, and can view all the service chains of all tenants.

A service provider user logs in to Cisco vManage as the **admin** user or a user who has Write permission for the Manage Users permission. The service provider user can add, edit, or delete users and user groups from the Cisco vManage server. The service provider is typically responsible for the following activities:

- Create and manage cluster for tenants.
- Upload prepackaged VM image packages and Cisco NFVIS software images on the CSP devices.
- Create custom colocation group and RBAC users.
- Create service groups and associate a colocation group to service groups.

- Upgrade CSP devices and Catalyst 9500 switches.
- Monitor service chains and VMs of all tenants.
- Start, stop, or restart operations on any of the tenant VNFs.
- Administer Cisco vManage and record system-wide logging of Cisco SD-WAN devices.

Tenants

The tenant can issue operations on the VNFs for the service chains that belong to the tenant, but they can't view, access, or issue operations on VNFs for the service chains that belong to another tenant. A tenant is responsible for the following activities:

- Monitor all service groups and the health status of service chains that belong to the tenant.
- Monitor event or alarms for VNFs that are part of the service chains that belong to the tenant.
- Issue start, stop, or restart operations on VNFs that are part of the service chains that belongs to the tenant.
- Collaborate with the service provider for any issues on cluster, service chains, or VNFs.

Recommended Specifications in Multitenant Environment

We recommend that the service providers use the following information to decide on the number of tenants, clusters, service chains per tenant, and VLANs for various colocation sizes:

Tenants	Clusters (CPUs)	Service Chains (CPUs) per Tenant	VLANs
150	2 (608)	1 (4)–Small	~300
75-150	2 (608)	2-3 (4-8)–Medium	300-450
25-50	2 (608)	4-6 (12-24)–Large	~400
300	4 (1216)	Small	~600
150-300	4 (1216)	Medium	600-900
50-100	4 (1216)	Large	~800
600	8 (2432)	Small	~1200
300-600	8 (2432)	Medium	900-1200
100-200	8 (2432)	Large	~1050
750	10 (3040)	Small	~1500
375-750	10 (3040)	Medium	600-1500
125-230	10 (3040)	Large	~1250

For example, if a service provider provisions 4 vCPU per tenant for a service chain that consists of a single VM, the provider can onboard approximately 150 tenants on two clusters with eight CSP devices. Each of these tenants or service chains requires 300 hand-off VLANs, 1 ingress, and 1 egress VLAN per service chain. For information about the number of VMs per service chain for various colocation sizes, see Ordering and Sizing of Cisco SD-WAN Cloud onRamp for Colocation Solution Devices, on page 14.

Assumptions and Restrictions in Colocation Multitenancy

Assumptions

- The wiring between Cisco CSP devices and Cisco Catalyst 9500 switches is completed as per prescriptive connections or flexible topology. To bring up multiple clusters, ensure that the wiring between the CSP devices and Catalyst 9500 switches of a cluster are in the same way as a single cluster.
- Each Cisco CSP device has two 1-GB management ports that are manually configured as port channels to the OOB management switch.
- A tenant can only monitor the event or alarms for the VNFs that are part of the service chains that they own. The tenant monitoring screens display the colocation group when viewing a service chain.

Note In a comanaged multitenant setup, the service provider provisions service chains for tenants by gathering the required information from a tenant. For example, a tenant provides the tenant organization name, tenant Cisco vBond Orchestrator IP address, tenant site Id, system ip, and so on, out of band. See Create Service Chain in a Service Group, on page 64.

- In a comanaged multitenant setup, a tenant can monitor only the VNF devices that the tenant owns.

Restrictions

- Altering a colocation cluster from a single-tenant mode to a multitenant mode and conversely isn’t supported.
- Sharing VNF devices across multiple tenants isn't supported.
- Service providers can provision multiple service groups for a tenant. But, the same service group can’t be provisioned for multiple tenants.
- Upgrading from Cisco SD-WAN Cloud onRamp for Colocation Release 20.4.1 that has a single-tenant mode to Release 20.5.1 or later that has a multitenant mode isn’t supported. This restriction means you can’t upgrade from a single-tenant mode to multitenant mode.
- Multitenancy in SR-IOV enabled PNICs isn't supported and only OVS for VNF VNICs is supported. All the PNICs in the CSP devices are in OVS mode as the current SR-IOV drivers don’t support VXLAN. The VNF VNICs are connected to OVS networks, and the performance might not be line rate.
- Managing the billing and subscription of the resources utilised by the tenants isn't supported.

Service Provider Functionalities

Provision New Tenants

The service provider provisions a new tenant by creating a colocation group, and then provides access to a tenant by creating an RBAC user for the user group associated with the colocation group. The RBAC user performs limited administrative duties within their own tenant environment.

Before you begin

A service provider should bring up clusters in the shared mode by establishing control connections with the CSP devices and activating the cluster. The service provider can create several clusters and each of these clusters can have between two to eight CSP devices and two Catalyst 9500 switches. The cluster creation operation supports an option to choose if the cluster is for a multitenant or single-tenant deployment. See Create and Activate Clusters, on page 38.

Procedure

Step 1 To onboard a tenant, a service provider creates a colocation group. This group provides access to tenants to monitor their service groups and VMs.

Step 2 A service provider adds an RBAC user and associates it with the colocation group created in Step1.

Note
The service provider shouldn't add an RBAC user if authenticating the user using the TACACS server instead of Cisco vManage. If authenticating a user using TACACS server, associate the user with the colocation group created in Step1.

Step 3 A service provider creates a service group, associates it with the colocation group, and attaches the service group to a specific cluster. See Create Service Chain in a Service Group, on page 64.

When a tenant requires a new service chain, the service provider uses handoff VLANs specific to the tenant.

Create Colocation Group

In a single-tenant Cisco vManage, a colocation cluster can be shared across multiple tenants by using colocation groups. The colocation groups are a mechanism to associate a service chain to a particular tenant. The RBAC users that are created for the tenants are the colocation groups. These users can log in to Cisco vManage using their credentials to view only their tenant-specific service chains and VNF information. If the service provider chooses to use a service group for a tenant, the colocation group needs to be created prior to creating a service group so that the colocation group can be associated with the service group.

Procedure

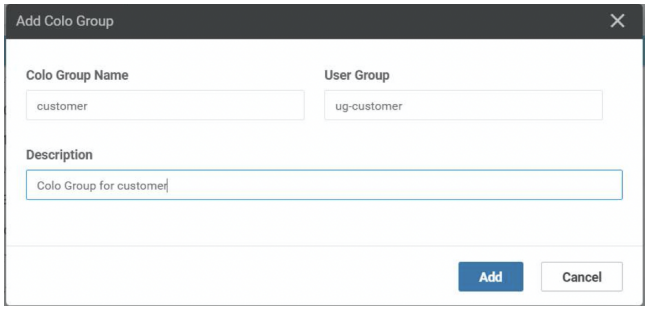
Step 1 In Cisco vManage, click **Administration** > **Colo Groups**.

Step 2 Click **Add Colo Group**.

Step 3 Provide a colocation group name, name of a user group with which the colocation group must be associated, and description.

You can view the colocation group name when creating a service group for a multitenant setup.

Step 4 Click **Add**.



View Permissions of a User Group

Procedure

Step 1 In Cisco vManage, click **Administration** > **Manage Users**.

Step 2 Click the **User Groups** tab.

Step 3 To view permissions of a user group, from the **Group Name** list, click the user group that you created.
You can view the user group and their permissions. To know about the list of user group permissions in a multitenant environment, see Manage Users Using Cisco vManage.

Create an RBAC User and Associate to Colocation Group

Procedure

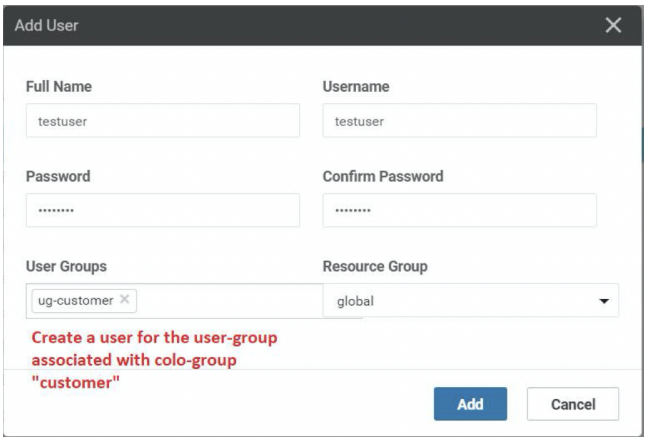
Step 1 In Cisco vManage, click **Administration** > **Manage Users**.

Step 2 Click **Add User**.

Step 3 In the **Add User** dialog box, enter the full name, username, and password for the user.

Note
You can't enter uppercase characters for usernames.

Step 4 From the **User Groups** drop-down list, choose the groups that the user must belong to. For example, choose a user group that you created for the colocation feature.



Note
The resource group **global** is chosen by default.

Step 5 Click **Add**. Cisco vManage now lists the user is in a table.

The RBAC users created for tenants or colocation groups can log in to Cisco vManage using their credentials. These users can view their tenant-specific service chains, VNF information after the service group associated with a tenant is attached to a cluster.

Delete RBAC User From Colocation User Group

To delete an RBAC user, service providers must remove the RBAC user from a colocation group if users are configured using Cisco vManage. If users are authenticated using TACACS server, service providers must disassociate the user from the user group in the TACACS server.

After deleting an RBAC user, the user can no longer access or monitor the devices of the cluster that is provisioned for the tenant. If an RBAC user is logged into Cisco vManage, deleting the user doesn't log out the RBAC user.

Procedure

Step 1 In Cisco vManage, click **Administration** > **Manage Users**.
Step 2 From the **Users** tab, click an RBAC user you want to delete.
Step 3 Click **More Actions** next to the row in the table, and click **Delete**.
Step 4 Click **OK** to confirm the deletion of the RBAC user.

Delete Tenants

To delete a tenant, service providers must remove the service groups associated with the tenant and then remove the colocation group for the tenant.

Procedure

Step 1 Locate the list of service groups associated with the tenant that you wish to delete. See View Service Groups in vManage, on page 86.

A tenant is a colocation group having one or more RBAC users associated to the same colocation group. In the service group configuration page, you can view the colocation group of the tenant.

Step 2 Detach the service group from the cluster for the tenant that you wish to delete. See Attach and Detach Service Group with Cluster, on page 87.

Step 3 To reuse the service group for another tenant, change the colocation group associated with the service group. If you delete the service group, you need to recreate it.

Step 4 Delete the colocation group for the tenant. See Manage a User Group.

Manage Tenant Colocation Clusters

A service provider manage the following tasks:

1. Activate clusters: A service provider configures devices, resource pool, system settings, and activates a cluster in the multitenant or shared mode. See Create and Activate Clusters, on page 38.
2. Create service groups and associate RBAC users to colocation groups: A service provider creates a colocation group, associates RBAC users to the colocation group, creates a service group, associates the service group with the colocation group for the multitenant mode, and attaches the service group to a specific cluster. See Create Service Chain in a Service Group, on page 64.

Note A service provider should associate specific service groups for each tenant.

3. Create VM packages: A service provider creates and uploads the VM packages into the Cisco vManage repository. The same packages can be used to provision VNFs in service chains for multiple tenants.

Note When a service group is associated with a colocation group, the SR-IOV option is ignored from the VNF package that is used for configuring the VNF. In a multitenant mode, the VNF packages only support OVS-DPDK with VXLAN.

4. Monitor service chains, VNFs of tenants: A service provider monitors all tenant service chains, and identifies service chains that are unhealthy and the tenants associated with these service chains. The service providers can collect logs from Cisco vManage or CSP devices and notify the tenants.
5. Add and remove Cisco CSP devices: To manage colocation clusters, a service provider can add or remove CSP devices.

Tenant Functionalities

Manage Colocation Clusters as Tenants



Note All tenants must monitor the service chains and VMs associated with the service chains, collaborate with service providers if any health issues arise with the service chains. The tenants can only monitor event or alarms for VNFs that are part of the service chains that belongs to the tenant.

Tenants don't have any administrative privileges and can only see the service chains that service providers create. The tenant monitoring screens display the colocation group when viewing service chains. The tenants perform the following tasks:

1. Log in to Cisco vManage as a tenant

   To log in:

   a. Log in to Cisco vManage using the URL HTTPS: //vManage-ip-address.

   b. To view the tenant monitor screen, enter the RBAC username and password.

2. View and monitor health of the tenant service chains along with the health of the VNFs. To know more about the different service chain health status, see Monitor Cloud OnRamp Colocation Clusters from Cisco vManage , on page 113.

   Click the **Diagram** button at the left to view all tenant service groups along with the service chains and VNFs in the design view.

3. To view the VNF health for a tenant, on the monitor screen, click **Network Functions**, and click a VNF name from the table.

   From the left pane, click the CPU utilization, memory utilization, and disk utilization to monitor the resources utilization of a VNF.

   You can also view the VM-specific alarms and events from the left pane.

4. To start, stop, or reboot a VNF, on the monitor screen, click a VNF name from the table. Click **More Actions** next to the VNF name, and choose one of the following operations:

   - **Start**
   - **Stop**
   - **Restart**

Monitor in Comanaged Multitenant Environment

Before you begin

- When creating a service chain using a provider Cisco vManage, the service provider should ensure that for the Cisco SD-WAN VM in a service chain, the correct UUID, and device OTP are entered. The provider has no access to the tenant overlay, and therefore a tenant should provide these information.
- When a service provider detaches a service group from a colocation cluster, the service provider should notify the tenant that the corresponding VM devices must be decommissioned using the tenant Cisco vManage.
- If a service provider needs to reattach a service group to a colocation cluster, a new OTP of the Cisco SD-WAN VM should be entered. This OTP is provided by the tenant. The service group on the provider Cisco vManage should be edited to save the new OTP of the Cisco SD-WAN VM.

Procedure

Step 1 Associate tenant Cisco SD-WAN devices with the service provider service group when creating a service chain. See Create Service Chain in a Service Group, on page 64.

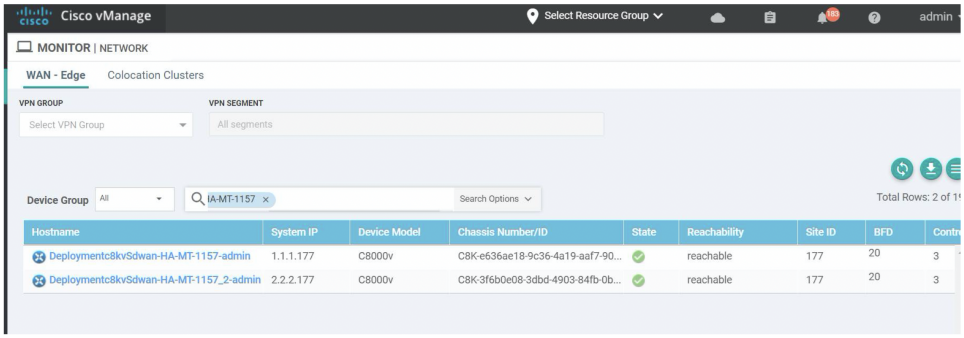
Step 2 Monitor the VNFs from the provider Cisco vManage. See Monitor Cloud OnRamp Colocation Clusters from Cisco vManage , on page 113.

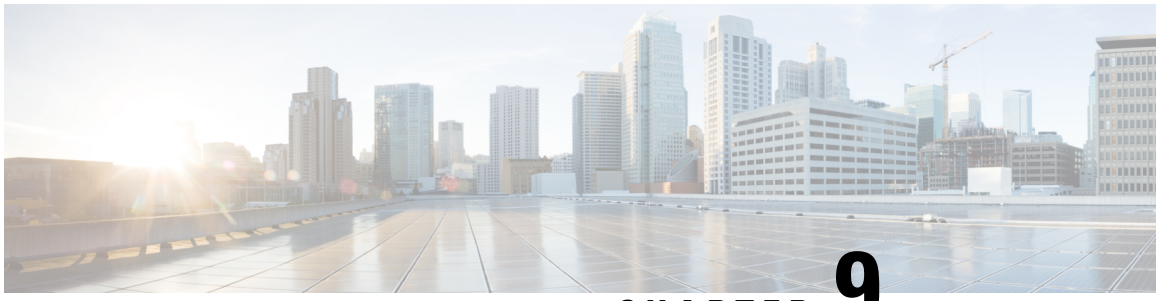
Step 3 Monitor the information about the Cisco SD-WAN devices of the VNFs from the tenant Cisco vManage.

Note
The service provider can't view information about the Cisco SD-WAN devices of the VNFs from the provider **Cisco vManage** > **Configuration** > **Devices** screen under **WAN Edge List** because these devices are controlled by the tenant.

Example:

The following devices of a VNF can be seen from the tenant Cisco vManage only.





CHAPTER 9

Troubleshoot Cisco SD-WAN Cloud onRamp for Colocation Solution

- Troubleshoot Colocation Multitenancy Issues, on page 139
- Troubleshoot Catalyst 9500 Issues, on page 140
- Troubleshoot Cisco Cloud Services Platform Issues, on page 145
- DHCP IP Address Assignment, on page 152
- Troubleshoot Cisco Colo Manager Issues, on page 153
- Troubleshoot Service Chain Issues, on page 155
- Troubleshoot Physical Network Function Management Issues, on page 157
- Log Collection from CSP, on page 157
- Troubleshoot vManage Issues, on page 157

Troubleshoot Colocation Multitenancy Issues

You can use the following commands to view the output and locate issues.

- To view an overview of the VNICs and VLANs of each VNF such as in which bridge they exist, use the `support ovs vsctl show` command.

  ```
  nfvis# support ovs vsctl show
  ```

- To verify the details of a service chain deployment with bridge, or network, or VLAN, use the `show service-chains` command.
- To view the data and HA VTEP IP addresses of a CSP device and the peer CSP devices in a colocation cluster, use the `show cluster-compute-details` command.
- To view the source and destination serial numbers of each HA bridge with the corresponding VLAN and VNID associations, use the `show vxlan tunnels` command.
- To view the data flows per tenant that you can identify by a user id with the VLAN, VNID mapping, use the `show vxlan flows` command.
- To view the VXLAN flow statistics, use the `support ovs ofctl dump-flows vxlan-br` command.
- To view the overall deployment status of VM life cycle, use the `show vm_lifecycle deployments` command.

End-to-end Ping Fails

1. Verify if the VMs are deployed by using the **show vm_lifecycle deployments all** command.
2. Verify that the service chains display the chain name attached to it by using the **show service-chains** command.
3. Verify notifications about events that have occurred on the Cisco SD-WAN device by using the **show notification stream viptela**
4. Ping the **data-vtep-ip** and **ha-vtep-ip** of the CSP peer device by using the **show cluster-compute-details** command.
5. Verify that the VLAN association per bridge, network, or VLAN is matching with the VNICs and VLANs of each VNF. Check the output from the **show service-chain** *chain-name* command matches with the output from the **support ovs vsctl show** command.
6. Contact Technical Support, if connection fails and you're unable to ping the peer CSP device.

Troubleshoot Catalyst 9500 Issues

This section covers some of the common Catalyst 9500 problems and how to troubleshoot them.

General Catalyst 9500 Issues

Switch devices are not calling home to PNP or Cisco Colo Manager

Verify the PNP list on Cisco Colo Manager to determine if the switch devices have not called home. The following are the good and bad scenarios respectively when the **show pnp list** command is used:

Devices have called home

```
admin@ncs# show pnp list
SERIAL IP ADDRESS CONFIGURED ADDED SYNCED LAST CONTACT
------------------------------------------------------------------------------
FCW2223A3VN 192.168.10.40 true true true 2018-12-18 22:53:26
FCW2223A4B3 192.168.30.42 true true true 2018-12-11 00:41:19
```

Devices have not called home

```
admin@ncs# show pnp list
SERIAL IP ADDRESS CONFIGURED ADDED SYNCED LAST CONTACT
------------------------------------------------------------------------------
```

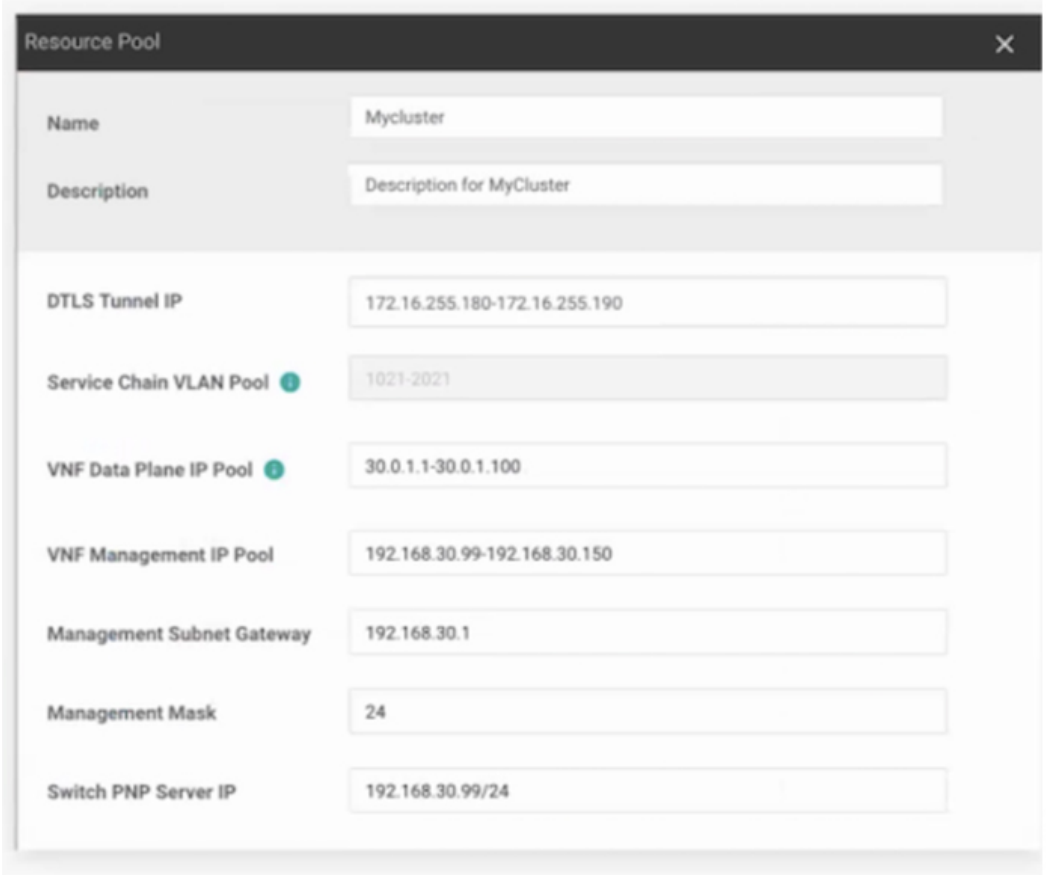
<– Empty list

Action:

1. Verify that the management interfaces on both the switches are not shut and have IP addresses.
2. Try running the **write erase** command on the switch and then reload. Verify that the IP address appears on the management interface.
3. Verify that the configuration for DHCP option 43 is valid. Here is a sample DHCP configuration where the PNP IP address is 192.168.30.99:

```
ip dhcp pool 192_NET network 192.168.30.0 255.255.255.0 dns-server 192.168.30.1
default-router 192.168.30.1 option 43 ascii "5A;B2;K4;I192.168.30.99;J9191" lease infinite
```

4. Verify that the PNP IP address provided on vManage for resource pool matches the IP address in DHCP configuration as follows:



5. Ping and determine whether both switches are reachable.

Catalyst 9500 failed to reach through DHCP option 43

Here Cisco Colo Manager is in healthy state at the host end, and Cisco Colo Manager internal state is in progress. If a cluster has already been activated, it shows that the cluster is in activation pending state. If a cluster has not been activated, it shows the cluster is not in activated state.

Action:

1. SSH into NFVIS as an admin user. Use the `ccm-console` command to log into Cisco Colo Manager. Run the `show pnp list` command.
2. If the PNP list is empty, verify the OOB status whether the Cisco Colo Manager IP address is correctly configured on the OOB switch.

Day-0 configuration push failed on both Catalyst 9500 switches

Here Cisco Colo Manager is in healthy state at the host end, and Cisco Colo Manager internal state is in progress. PnP configuration push fails with an error and Cisco Colo Manager is in-progress state.

Action:

1. Clean the Catalyst 9500 switches by using the **renumber** and **write erase** commands.
2. Deactivate and Reactivate the cluster again from vManage to repush the Day-0 configuration.

Day-0 configuration push fails on the secondary Catalyst 9K switch

Here Cisco Colo Manager is in healthy state at the host end, and Cisco Colo Manager internal state shows, "Failure." Cisco Colo Manager shows that only one switch is brought up successfully and cannot detect the secondary switch failure.

Action:

1. Clean the secondary Catalyst 9500 switch by using the **renumber** and **write erase** commands.
2. Deactivate and Reactivate the cluster again from vManage to repush the Day-0 configuration.

One of the Catalyst 9500 switches is up and running. The secondary switch is not in SVL configuration and SVL link cables are not connected

Here Cisco Colo Manager is in healthy state at the host end, and Cisco Colo Manager internal state shows, "Failure." Both switches are onboarded with an IP address. Cisco Colo Manager detects an error as both switches are connected, as the SVL link on he switches are missing. You can see both switches as "Green" in vManage.

Action:

1. Verify the SVL link cables.
2. Verify licenses of both Catalyst 9500 switches.

Day-0 configuration push fails and connectivity to switch is down

Here Cisco Colo Manager is in healthy state at the host end, and Cisco Colo Manager internal state shows, "Failure" until the next Day-0 configuration push. NSO sends notification of not being able to push configuration. You can see a switch as "Red" in vManage, which means connectivity is down.

Action:

1. Verify the health of the Catalyst 9500 switch.
2. Bring the switch back to online.
3. Start pushing Day-0 configuration again.

Unable to log into Catalyst 9500 after PNP from vManage

If vManage is not able to push more configuration to a Catalyst 9500 after PNP, you might have been locked out of the switch.

Action:

1. Log into NFVIS by using **admin** as the login name and **Admin123#** as the default password.

Note The system prompts you to change the default password at the first login attempt. Ensure that you set a strong password as per the on-screen instructions.

2. Use the **ccm console** command on NFVIS to log into Cisco Colo Manager. Run the following commands on Cisco Colo Manager to add a user to Catalyst 9500 switches.

   - ```
     config t
     cluster <cluster-name>
     system rbac users user admin password
     $9$yYkZqj7lQcrRL3$sZ23jqv5buK4lYCkt0dCbO6xYEfxRHQJiQnrlFdYHBg
     ```

Note Ensure that you set password as a scrypt string.

Now the corresponding user is added to Catalyst 9500 switches and you can SSH to the switches by using user and password.

Issues with a cluster activation, admin and password cannot be pushed to Catalyst 9500

Action:

1. If a cluster activation is in still in pending state, verify if colo-config-status is in IN-PROGRESS state. If state is In-Progress, the synchronization has not been done and no new configurations can be pushed. This process can take up to 20 minutes.
   a. If Cloud OnRamp for Colocation configuration status is In-progress state for a long time, SSH into NFVIS as an admin user. Use the **ccm-console** command to log into Cisco Colo Manager. Run the **show pnp list** command. Verify if two switches are added.
   b. If only one switch is displayed, ensure that the other switch configuration is cleaned by using the **write erase** command and reloaded. The secondary switch startup configuration must be erased and returned to its initial state.
   c. Ensure switch connectivity with PNP server in Cisco Colo Manager.
2. If a cluster has been activated successfully, verify if colo-config-status is in "SUCCESS" state. If status is displayed as Success, your admin password must have been pushed to a switch. If not, on vManage, add a new credential to the switch and then push new configurations.
3. If a cluster activation fails and colo-config-status is in "FAILED" state, use the RBAC to push a new authentication from ccm-console. In the following example, the password is encryption of "Cisco-123."

   cluster cluster system rbac users user Alpha password 9Z9Sr2VOuwjwC74$qEYAmxgoaW4m07.UjPGR9gL2ksFkcCIgIcEYOUWxDFo role administrators

Note You cannot push any RBAC configuration if a cluster is in active state. vManage does not allow out of bound change to Cisco Colo Manager.
     ```

Clean switches configuration and reset switches to factory defaults

During a cluster creation, cluster clearing, cluster deletion, the configurations of both switches must be cleaned. To clean switches configuration, perform the following steps:

Action:

1. Use the **show switch** command to determine the switch number and whether the provisioned switch exists in the switch stack. If the switch number is two, use the **switch 2 renumber 1** command.



Note The switch renumbering is essential for SVL stack mode.

2. To erase the switch startup configuration and return it to its initial state, use the **write erase** command.
3. To reload the switch with a new configuration, use the following command in privileged EXEC mode and type n for not saving the modified configuration:

```
switch(config)#reload
```

4. Perform steps 2 and 3 on the second switch device after the switch stack reloading has been completed on the first switch.

To verify addition of switch devices from Cisco Colo Manager, perform the following steps:

1. Log into Cisco Colo Manager and use the **show pnp list** command.

   The two switch devices are displayed. PNP pushes the Day-0 configuration, adds switch devices into the Cisco Colo Manager device tree, and synchronizes the device configuration with Cisco Colo Manager. If any of the switch devices cannot be viewed, the PNP of the missing switch device may be misconfigured or network may be down.

   SVL configuration that is pushed to switches issues a reboot command to switches, after the reboot. Both switch devices are up and become one stack.

2. On Cisco Colo Manager, trigger a timer for around 14 minutes to perform another synchronization on the primary device.
3. To view the device configuration and current status, use the **show cluster** *cluster-name* command.

   If status is displayed as "GREY," the switch devices are not yet added to the Cisco Colo Manager device list. If status is displayed as "RED," the switch devices are not reachable. If status is displayed as, "GREEN," the device is currently connected. Also, you can view which is the primary switch device.

4. To view the devices status in a colocation, use the **show colo-config-status** command. If status is in "In-progress," the switch devices are not yet synchronized and vManage cannot send any further configuration. See Chapter, Monitor Cisco SD-WAN Cloud onRamp for Colocation Solution Devices, on page 107 for more information about Cisco Colo Manager state transitions.

   After the timer reaches its duration (for example, 14 minutes), Cisco Colo Manager tries to synchronize again with the primary Catalyst 9500 device.

   After the second synchronization has been completed, Cisco Colo Manager state is displayed as, "SUCCESS".

Configuration on switch after QoS policy is applied

When QoS policy is applied, the following configuration appears on the switch device after you set the bandwidth for a service chain and deploy it:

```
class ASAvOnly_chain1_VLAN_210police 2000000000class ASAvOnly_chain1_VLAN_310police
2000000000policy-map
service-chain-qosclass ASAvOnly_chain1_VLAN_210police 2000000000class
ASAvOnly_chain1_VLAN_310police 2000000000
```

Troubleshoot Cisco Cloud Services Platform Issues

This section covers some of the common Cloud Services platform (CSP) problems and how to troubleshoot them.

RMA of Cisco CSP Devices

Use the **admin tech** command for the CSP device from Cisco vManage to collect the log information for the device on the **Tools** > **Operational Commands** screen. Verify the following log files:

- `nfvis_config.log`: Displays the device configuration-related logs
- `escmanager.log`: Displays VM deployment-related logs.
- `Tech-support-output`: Use the following show commands that are available from the CSP device.
  - `cat/proc/mounts`: Displays mount information
  - `show hostaction backup status`: Displays the status of the last five backups taken on the CSP device
  - `show hostaction restore-status`: Displays the status of the overall restore process and each component such as device, image and flavors, VM, and so on
  - `show vm_lifecycle deployments`: Displays the deployment name and the VM group name.

The following is an example of the mount operation on the NFS server:

```
nfvis# show running-config mount
mount nfs-mount storage sujathast/
storagetype            nfs
storage_space_total_gb 5000.0
server_ip              192.168.0.1
server_path            /NFS/colobackup
```

The following is an example of the operational status output for the last five backup operations and notifications on Cisco vManage for the last backups:

```
eventTime 2021-02-02T04:02:25.577705+00:00
viptela
severity-level minor
host-name nfvis
system-ip 10.0.0.1
user_id admin
config_change false
transaction_id 0
status SUCCESS
status_code 0
status_message Backup configuration-only to nfs:test_storage/test_config_only.bkup completed
```

```
 successfully with operational status: BACKUP-COMPLETED-PARTIALLY
details NA
event_type BACKUP_SUCCESS
severity INFO
host_name nfvis
 !
```

The following example shows that status of the device after using the `show hostaction restore-status` command:

```
nfvis# show hostaction restore-status
hostaction restore-status 2021-03-19T20:53:15-00:00
source nfs:sujathast/WZP22160NC7_2021_03_19T19_10_04.bkup
status RESTORE-ERROR
components NFVIS
status      RESTORE-ERROR
last update 2021-03-19T21:02:11-00:00
details     "Unable to load configuration Editing of storage definitions is not allowed"
components nfs:sujathast/WZP22160NC7_2021_03_19T19_10_04.bkup
status      VERIFICATION-SUCCESS
```

Clear Status of VNICs and PNICs

1. To view the PNIC stats, use the **show pnic stats** command.
2. To view the VNIC stats, use either of the following commands:
   - **show vm_lifecycle vnic_stats** for all VMs
   - **show vm_lifecycle vnic_stats** *vm-name* for a single VM
3. To clear the stats of one or more VMs, run the following commands:

   ```
   clear counters vm all
   clear counters vm vm-name vnic vnic-id
   clear counters vm vm-name vnic all
   ```

4. To clear the stats of all PNICs and VNICs, use the **clear counters all** command.

When CSP reboots, all PNIC and VNIC counters are erased and the counters are cleared. If the stats of VNICs and PNICs aren't displayed, you can use the following commands to view the stats:

```
show pnic-clear-counter
show vm_lifecycle tx_rx_clear_counters
```

Issues in Cisco CSP Device Onboarding

1. To verify that the device has established a secure control connection with the SD-WAN controllers, use the **show control connections** command.
2. To verify the device properties used to authenticate the devices, use the **show control local-properties** command.

   From the displayed output, make sure:

   - system parameters are configured to include **organization-name** and **site-id**
   - **certificate-status** and **root-ca-chain-status** are installed
   - certificate-validity is **Valid**

- **dns-name** is pointing to vBond IP address or DNS
- **system-ip** is configured, **chassis-num/unique-id**, and **serial-num/token** is available on the device

3. To view the reason for failure, if a device fails to establish connection with the Cisco SD-WAN controllers, use the**show control connections-history** command. View the **LOCAL ERROR** and **REMOTE ERROR** column to gather error details.

   The following are the reasons the Cisco CSP device fails to establish control connections with the Cisco SD-WAN controllers.

   - CRTVERFL – the error state indicates that the device authentication is failing because of a root-ca certificate mismatch between the device and the Cisco SD-WAN controller. Use the **show certificate root-ca-cert** on Cisco CSP devices to confirm that the same certificates are installed on the device and the Cisco SD-WAN controllers.
   - CTORGNMMIS - the error state indicates that the device authentication is failing because of a mismatch organization-name, compared with the organization-name configured on the Cisco SD-WAN controller. Use show **sdwan control local-properties** on CSP devices to confirm all the SD-WAN components are configured with same organization-name.
   - NOVMCFG – the error status indicates that the device hasn't been attached with a device template in Cisco vManage. This status is seen when onboarding the device using automated deployment options, which is the PnP.
   - VB_TMO, VM_TMO, VP_TMO, VS_TMO – the error indicates that the device has lost reachability to the Cisco SD-WAN controllers.

Failure in Cluster Activation

In CCM, verify if the SVL formation of switches is complete and the devices are onboarded by viewing CCM notifications status.

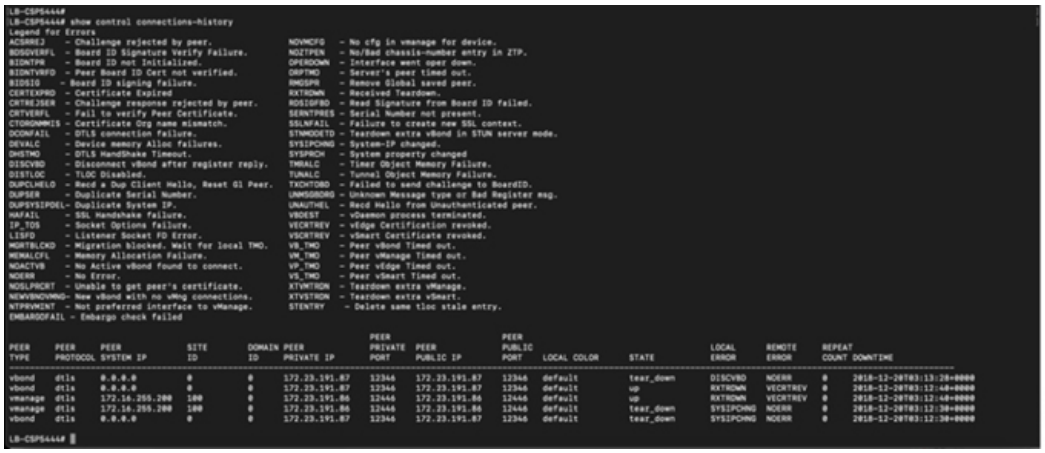
1. Ensure that all the SR-IOV and OVS ports are cabled correctly to the Catalyst 9500 switches and the interfaces are in link-up state.
2. Determine the SR-IOV and OVS ports using the **show lldp neighbors** command on a CSP device and verifying the wiring between the CSP devices and Catalyst 9500 switches.

   Ensure that the **show lldp neighbors** command displays all eight ports are powered up and reports about the neighbors.
3. Ensure that the Catalyst 9500 switches are in SVL mode and the interfaces have the description, "SVL Complete."

Failures with Certificate installation

Use the **show control connections-history** command to determine certificate installation failures.

Figure 43: Certificate Installation Failure



Action:

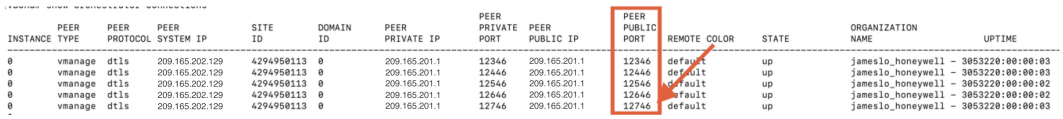
The following are the verifications that you can perform based on errors that you might encounter:

- vbond with error SERNTPRES–This error is caused, if the serial or token on device don't match with vBond serial or token. Check vManage to ensure that the device is in "valid" state and it was decommissioned properly.
- vManage with error NOVMCFG–This error is caused if the template wasn't attached to the device. Activating the cluster resolves this issue.
- On vBond, verify that the **show orchestrator valid-vedges** command shows the device correctly. This means that the device is valid with the same token that you had used.
- Ensure that the clocks on vManage and CSP devices are synchronized.

Failures with Control Connection

The **show control connections-history** displays DCONFAIL. Open the firewall to view the ports that need to be opened.

Figure 44: Failure with Control Connection, DCONFAIL



Action:

The following ports need to be opened:

Table 39: UDS and TCP Ports to be Opened

Core Number	Ports for DTLS (UDP)	Ports for TLS (TCP)
Core0	12346	23456
Core1	12446	23556

Core Number	Ports for DTLS (UDP)	Ports for TLS (TCP)
Core2	12546	23656
Core3	12646	23756
Core4	12746	23856
Core5	12846	23956
Core6	12946	24056
Core7	13046	24156

CSP doesn't have a DHCP IP address

The CSP device doesn't get displayed in vManage as a connected device.

Action:

1. Connect to a CSP through the CIMC interface.
2. Verify if the CSP has an IP address by running the **`show system:system settings`** command on the Cloud OnRamp for Colocation management port.
3. Verify if the DHCP server has IP addresses. To assign a static IP address and configure DHCP sticky IP, see DHCP IP Address Assignment, on page 152.
4. Verify that the PNP server is reachable by a ping.
5. From the PNP server, verify if the CSP device can be contacted and claimed, or redirection is successful. In the PNP portal, if it shows Pending Redirection for the device, verify if the serial number is same as CSP devices.
6. Use the **`show platform-details`** command on CSP to determine the serial number.
7. In the PNP portal, verify if it shows Connected.

CSP hasn't established connectivity with vManage

The CSP device doesn't get displayed in vManage as a connected device.

Action:

1. Verify if the CSP device has root CA installed from PNP by using the **`show certificate installed`** and **`show certificate root-ca-cert`**.
2. Verify if CSP can ping the vBond IP address. Then, attain the vBond IP by using the **`show running-config viptela-system:system`**
3. If ping to vBond fails, verify the network connectivity on the management interface.
4. If ping to vBond goes through, use the **`running-config vpn 0`** to view the configuration for control connection.
5. If the control connection configuration exists, verify vManage settings.

6. In vManage, verify if a cluster is activated and device OTP information has been included by using the **show control connections** and **show control local-properties** commands.
7. Verify if the CSP token number has been manually entered by using the **request vedge-cloud activate chassi-number token-number** command. Rerun the command with the correct OTP.

Factory reset of CSP device

To reset a CSP device to factory default, use the following command.

CSPxx# factory-default-reset all

The command deletes VMs and volumes, files including logs, notifications, images, and certificates. It erases all configuration. The connectivity is lost, admin password is changed to the factory default password. The system is rebooted automatically after reset and you must not perform any operation for 15- 20 minutes when factory reset is in progress. You can continue when prompted to proceed with the factory reset process.

CSP with a bad storage disk

The control connection is brought up and cluster is activated. The vManage monitoring screen displays all the eight CSP disks are available and one of the disks that is faulty.

Action:

Replace the faulty disk.

CSP device has less memory or CPU

The control connection is brought up and cluster is activated. The vManage monitoring screen displays that the memory threshold has reached.

Action:

Upgrade the specific CSP device that matches the minimum requirements.

I/O cards on CSP device are on wrong slots

Action:

Verify the slot details from CIMC inventory.

Colo Manager is not healthy on a CSP device

Action:

1. To verify Cisco Colo Manager state:
   a. Verify the health of the container by using the **show container ColoMgr** command. See Troubleshoot Cisco Colo Manager Issues, on page 153.
   b. View notifications about events from the Viptela device by using the **show notification stream viptela** command
2. To access Cisco Colo Manager, run the **ccm console** command on the CSP device where Cisco Colo Manager has been enabled.

This action takes you to the Cisco Colo Manager CLI. Run the **show running-config cluster** *cluster name* command.

3. Get the logs from vManage by using the **admin-tech** command. Alternatively, you can get the logs from the device directly. See Log Collection from CSP, on page 157.

Day-0 configuration push to CSP fails

The failure can be either due to CSP not having the correct hardware or Day-0 configuration of VNF has wrong input.

Action:

1. Verify the hardware configuration of CSP and ensure that it's a supported configuration.
2. Verify service chain Day-0 configuration, and then retrigger configuration push.

CSP doesn't get added to a cluster

Cluster state in the **vManage** > **Cofigurationn** > **Cloud OnRamp for Colocation** interface shows, "FAILED." The added CSP is depicted as "RED" in the Cloud OnRamp for Colocation graphical representation.

Action:

1. Verify the hardware configuration of CSP and ensure that it's supported.
2. Retry activating the cluster again

IP connectivity with CSP can't be retained

When CSP devices renew its DHCP IP, the IP connectivity to the CSP can't be retained.

Action:

For DHCP IP address allocation, ensure that the DHCP server is always on the same subnet as the CSP devices.

CSP devices aren't able to reach vManage

Action:

Perform the following steps:

1. Install NFVIS on the CSP device by using the KVM console. See the Cisco Enterprise NFV Infrastructure Software Configuration Guide for information about installing NFVIS.
2. Log in to the NFVIS system and ping gateway

   If it's not pinging or reachable, ensure OOB switch ports that are connected to the switch has port-channel configuration that is done.

   a. If port-channel configuration on a switch is missing, run the nfvis# **support ovs appctl bond-show mgmt-bond** command. The output is as follows:

```
--- mgmt-bond ----
bond_mode: balance-slb
bond may use recirculation: no, Recirc-ID : -1
bond-hash-basis: 0
updelay: 0 ms
downdelay: 0 ms
```

```
next rebalance: 3479 ms
lacp_status: configured
active slave mac: 00:00:00:00:00:00(none)
slave eth0-1: disabled
                may_enable: false
slave eth0-2: disabled
                may_enable: false
```

b. If the port channel on a switch is configured, but eth0-2 isn't connected to the switch, run the nfvis# **support ovs appctl bond-show mgmt-bond** command. The following ouput now shows that eth0-2 isn't connected to switch:

```
---- mgmt-bond ----
bond_mode: balance-slb
bond may use recirculation: no, Recirc-ID : -1
bond-hash-basis: 0
updelay: 0 ms
downdelay: 0 ms
next rebalance: 4938 ms
lacp_status: off
active slave mac: 50:2f:a8:c7:64:c2(eth0-1)


slave eth0-1: enabled
active slave
may_enable: true
hash 195: 2 kB load


slave eth0-2: disabled
may_enable: false
```



Note vManage manages the CSP devices and hence OOB configuration through NETCONF or REST API or CLI causes devices to be out of synchronization with vManage. vManage deletes this configuration when the next configuration is pushed from vManage. For any troubleshooting, to configure the CSP or NFVIS, use configuration only in shared mode and/or NETCONF target candidate followed by commit. This configuration is required as the Confd database, CDB is in a candidate mode on NFVIS for Cisco SD-WAN Cloud onRamp for Colocation solution. If the **confg t** CLI mode or NETCONF target running is used, the CDB database can be out of synchronization and cause strange behavior on CSP devices and results into an unusable cluster.

DHCP IP Address Assignment

To configure a static IP address:

1. After clean installation of the DHCP server, run **confd cli**.
2. Verify the existing configuration by using the nfvis# **show running-config vm_lifecycle** command.

   For example,

   nfvis# **show running-config vm_lifecycle networks**

   ```
   vm_lifecycle networks network int-mgmt-net
   !
   ```

3. Set up a static IPv4 address by using the nfvis# **config shared** command.

For example,

nfvis# **config shared**

```
Entering configuration mode terminal
nfvis(config)# vm_lifecycle networks network int-mgmt-net subnet int-mgmt-net-subnet
address <host-ip> gateway <host-ip-gateway> netmask <your-host-ip-netmask>  dhcp false
nfvis(config-ip-receive-acl-0.0.0.0/0)# commit
Commit complete.
nfvis(config-ip-receive-acl-0.0.0.0/0)# end
nfvis#
```

Configure DHCP Sticky IP

For sticky DHCP IP, configure the DHCP servers. Ensure that you have the serial number of the device readily available.

1. If you use CentOS 7.4 as the DHCP server, ensure that you have the following similar configuration in `/etc/dhcp/dhcpd.conf`.

   ```
   host abcxxxx175 {
   option dhcp-client-identifier <serial number>;
   }
   ```

2. If you use IOS as the DHCP server, ensure that you have the following similar configuration in an IOS DHCP server or pool.

   ```
   ip dhcp pool P_112
   host 209.165.201.12 255.255.255.0
   client-identifier 4643.4832.3xxx.3256.3xxx.48
   ```

   In this example, the IP address, 209.165.201.12 is the DHCP sticky IP for a client with identifier: 4643.4832.3xxx.3256.3xxx.48. Then, you can find out the client-identifier.

3. To find the client identifier, on an IOS DHCP server, turn on **debug ip dhcp server packet**.

   From the debug console output, you can view DHCP client-identifier of the SD-WAN Cloud OnRamp for Colocation device.

Troubleshoot Cisco Colo Manager Issues

This section covers some of the common Cisco Colo Manager problems and how to troubleshoot them.

General Cisco Colo Manager Issues

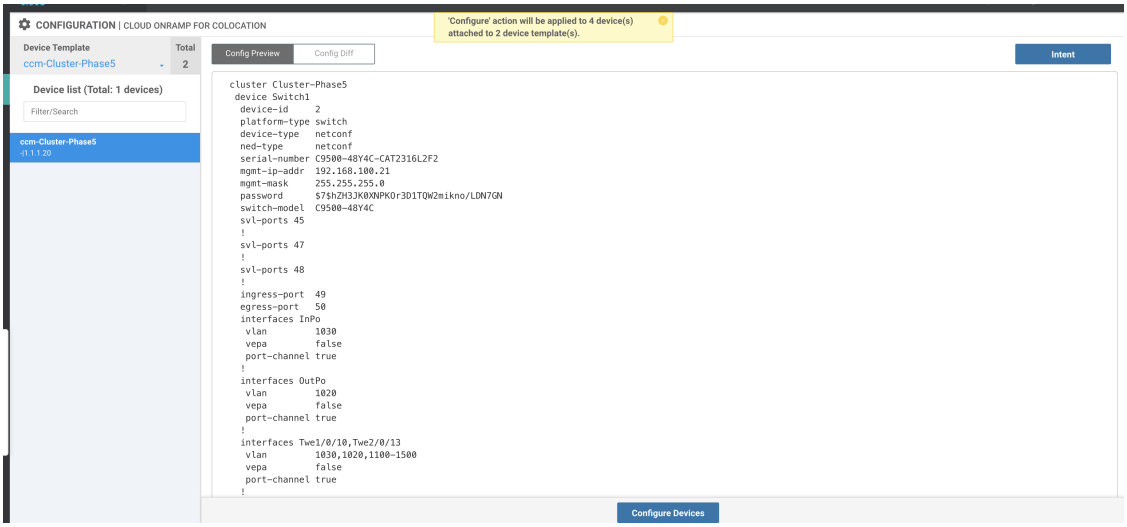
Verify Port Connectivity when SVL Formation Fails

After activating a cluster, to verify the SVL and uplink ports from CCM, perform the following steps:

1. From Cisco vManage, click **Configuration** > **Cloud OnRamp for Colocation**.
2. To verify the port connectivity of a cluster, choose the cluster from the table, click the **More Actions** icon to the right of the row, and then choose **Sync**.
3. Under **Device Template**, click the colocation cluster, and then choose the CCM cluster from the drop-down list.
4. To view the CCM configuration, click the CCM cluster.

You can now view port connectivity details of both the switch devices in the cluster and determine the connectivity issues.

Figure 45: Verification of SVL and Uplink Ports



Failure in Cisco Catalyst 9500 SVL Formation

1. Establish an SSH session with Cisco NFVIS as an admin user. Use the **ccm-console** command to log into Cisco Colo Manager and run the **show colo-config-status** command.

   ```
   admin@ncs# show colo-config-status
   ```

   Displays the recommended action.

   ```
   colo-config-status status failure
   colo-config-status description "Step 4 of 7:
   Device c9500-2 : 192.168.6.252 (CAT2324L42L)
   SVL ports specified by vmanage does not match with
   actual cabled svl ports. Recommended action: Correct
   the configured svl ports specified in cluster
   configuration by vmanage in accordance with switch
   SVL port cabling" colo-config-status severity critical
   ```

2. Ensure that the ports you choose for SVL on Cisco vManage match the physically cabled ports, and that they are detected by the Cisco Catalyst 9500 switches.

Cisco Colo Manager is unhealthy while activating a cluster for Day-0, or Cisco CSP is deleted when Cisco Colo Manager is running. Also, the new Cisco Colo Manager on the newly added Cisco CSP device fails to instantiate or becomes unhealthy

Here Cisco Colo Manager is in unhealthy state at the host end, and Cisco Colo Manager internal state shows, "FAILURE." vManage monitoring also shows Cisco Colo Manager in "UNHEALTHY" state.

Action:

1. Verify the Cisco Colo Manager state on the newly added Cisco CSP device by running the **show container ColoMgr** command.

```
CSP1# show container ColoMgr
container ColoMgr
 uuid       57b9b8646ff1066ba24707415b5449111d915664629f56221e141c1171ee283d
 ip-address 172.31.232.182
 netmask    24
 default-gw 172.31.232.2
 bridge     int-mgmt-net-br
 state      healthy
 error
CSP1#
```

2. Verify the reason for Cisco Colo Manager being in unhealthy state by looking at the error field as shown in the previous step.
3. For failures that are related to pinging the gateway, verify the Cisco Colo Manager parameters such as, IP address, mask and gateway IP address are valid. Also, verify the physical connection reachability to the gateway.
4. If any of the parameters are incorrect, fix them from vManage, and then retry activating cluster or synching.
5. If reason for Cisco Colo Manager being unhealthy are package errors, contact Technical Support.

Troubleshoot Service Chain Issues

This section covers some of the common service chain problems and how to troubleshoot them.

General Service Chain Issues

Service chain addition or deletion in to a service group fails

- Action:
- Cisco Colo Manager is in healthy state at the host end, and Cisco Colo Manager internal state shows, "FAILURE" for the configuration push. The configuration push fails, Cisco Colo Manager is in "FAILURE" state, and cluster is in "FAILURE" state.

  Action:

  1. To access Cisco Colo Manager, run the **`ccm console`** command on the CSP device where Cisco Colo Manager has been enabled.

     This action takes you to the CLI on Cisco Colo Manager. Run the following commands:

     a. **`show colo-config-status`**

        This action enables you to view the reason for failure in the description.

     b. If more information is required to debug the failure, collect logs by using the **`admin-tech`** command on CSP hosting Cisco Colo Manager. Alternatively, you can get the logs from the device directly. See Log Collection from CSP, on page 157.

  2. Verify the Day-0 configuration of VNF service chains.
  3. Provision the VNF service chain again.

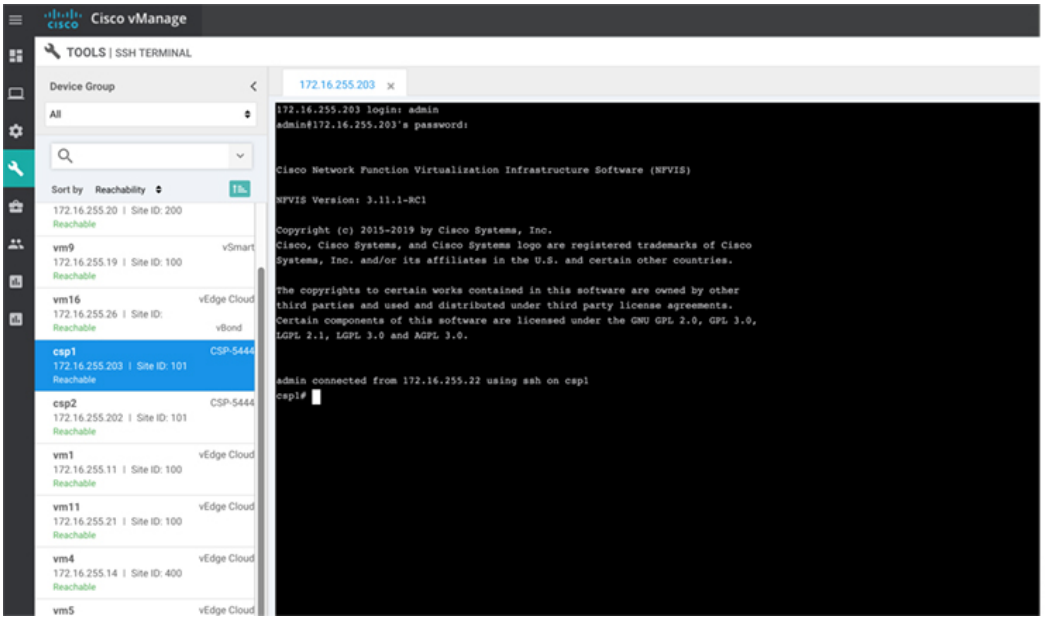
Note If service chain addition or deletion results in a failure on Cisco Colo Manager, there is an option to synchronize.

During service chain addition, VNF goes into error state

VNF is shown as down on vManage.

Action:

1. Verify the Day-0 configuration of VNF.
2. SSH from vManage to go to the CSP hosting the VNF.



3. Run the following commands:

   nfvis# show system:system deployments

   nfvis# get the VNF ID

   For example,

   NAME ID STATE

   Firewall2_SG-3 40 running

   nfvis# support show config-drive content 40

Ensure that all variables are properly replaced with key, value pairs.

Troubleshoot Physical Network Function Management Issues

To troublehsoot the sharing of PNF devices, ensure that the following are considered:

1. Cabling of PNF devices to Catalyst 9500 is correct and VLAN configurations are on the right ports of Catalyst 9500.
2. Verifying the LLDP enablement. By default, LLDP is enabled on Catalyst 9500. Ensure that you enable LLDP on PNF and check the LLDP neighbor and neighbor interface to confirm connectivity.
3. Verifying the missing configurations on PNF.

Log Collection from CSP

If CSP is not reachable from vManage, and logs need to be collected for debugging, use the **tech-support** command from CSP.

The following example shows the usage of the tech-support command:

```
nfvis# tech-support
nfvis# show system:system file-list
system:system file-list disk local 1
 name          nfvis_scp.log
 path          /data/intdatastore/logs
 size          2.1K
 typ
```

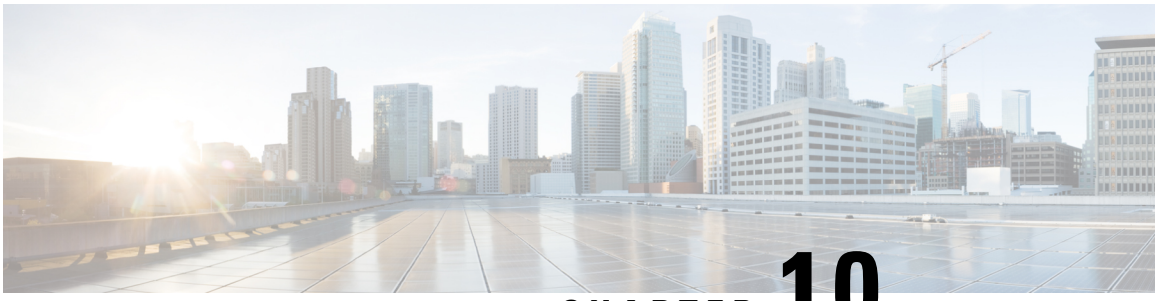
To secure copying a log file from the Cisco NFVIS to an external system or from an external system to Cisco NFVIS, the admin user can use the scp command in privileged EXEC mode. The following example shows the scp techsupport command:

```
nfvis# scp techsupport:NFVIS_nfvis_2019-04-11T15-33-09.tar.gz
cisco@172.31.232.182:/home/cisco/.
```

Troubleshoot vManage Issues

Use the following location to troubleshoot vManage issues,

SD-WAN Techzone Knowledge Base



CHAPTER 10

Custom Packaging Details for Shared VNF

- Cisco vEdge Router Variable List, on page 159
- Cisco CSR1000V Variable List, on page 163
- ASAv Variable List, on page 167

Cisco vEdge Router Variable List

In the following Cisco vEdge Router variable list, same variable names can be used for service chains five and six respectively with appropriate renumbering as mentioned for service chains.

Cisco vEdge Router Variable List

Cisco vEdge Router is in StandAlone Mode and Neighbor is in HA Mode
The input to the first VNF is in access mode (hypervisor-tagged) and the neighbor (ASAv firewall) is in HA mode.

User Variables	System Variables	
	Mandatory Variables when Service Chains 1 and 2 are Shared.	**Optional Variables wh Shared.**
DNS_SERVER	OTP	
UUID	VBOND_IP	
INSIDE_PRIM	ORG_NAME	
INSIDE_DATA_MASK_LEN	BGP_NO	
INSIDE_PEER_DATA_IP_PRIM	SYSTEM_IP	
INSIDE_AS	MGMT_PRIM	
LOCAL_INSIDE_AS	MGMT_MASK_LEN	
INSIDE_GW	MGMT_GW	
SERVICE_VPN	RCC	
SERVICE_VPN_2	VM_INSTANCE_NAME	

User Variables	System Variables	
SERVICE_VPN_3	OUTSIDE_PRIM	OUTSIDE_PRIM_3
SERVICE_VPN_4	OUTSIDE_DATA_MASK_LEN	OUTSIDE_DATA_MASK
	OUTSIDE_PEER_DATA_IP_PRIM	OUTSIDE_PEER_DATA_
	OUTSIDE_AS	OUTSIDE_AS_3
	OUTSIDE_VLAN_1	OUTSIDE_VLAN_3
	OUTSIDE_PEER_DATA_IP_SEC	OUTSIDE_PEER_DATA_
	OUTSIDE_PRIM_2	OUTSIDE_PRIM_4
	OUTSIDE_DATA_MASK_LEN_2	OUTSIDE_DATA_MASK
	OUTSIDE_PEER_DATA_IP_PRIM_2	OUTSIDE_PEER_DATA_
	OUTSIDE_AS_2	OUTSIDE_AS_4
	OUTSIDE_VLAN_2	OUTSIDE_VLAN_4
	OUTSIDE_PEER_DATA_IP_SEC_2	OUTSIDE_PEER_DATA_

Cisco vEdge Router Variable List

Cisco vEdge Router is in StandAlone Mode and Neighbor is in StandAlone Mode
The input to the first VNF is in access mode (hypervisor-tagged) and the neighbor is in StandAlone mode.

User Variables	System Variables	
	Mandatory Variables when Service Chains 1 and 2 are Shared.	**Optional Variables when Shared.**
DNS_SERVER	OTP	
UUID	VBOND_IP	
INSIDE_PRIM	ORG_NAME	
INSIDE_DATA_MASK_LEN	BGP_NO	
INSIDE_PEER_DATA_IP_PRIM	SYSTEM_IP	
INSIDE_AS	MGMT_PRIM	
LOCAL_INSIDE_AS	MGMT_MASK_LEN	
INSIDE_GW	MGMT_GW	
SERVICE_VPN	RCC	
SERVICE_VPN_2	VM_INSTANCE_NAME	

User Variables	System Variables	
SERVICE_VPN_3	OUTSIDE_PRIM	OUTSIDE_PRIM_3
SERVICE_VPN_4	OUTSIDE_DATA_MASK_LEN	OUTSIDE_DATA_MA
	OUTSIDE_PEER_DATA_IP_PRIM	OUTSIDE_PEER_DA
	OUTSIDE_AS	OUTSIDE_AS_3
	OUTSIDE_VLAN_1	OUTSIDE_VLAN_3
	OUTSIDE_PRIM_2	OUTSIDE_PRIM_4
	OUTSIDE_DATA_MASK_LEN_2	OUTSIDE_DATA_MA
	OUTSIDE_PEER_DATA_IP_PRIM_2	OUTSIDE_PEER_DA
	OUTSIDE_AS_2	OUTSIDE_AS_4
	OUTSIDE_VLAN_2	OUTSIDE_VLAN_4

Cisco vEdge Router Variable List

Cisco vEdge Router is in StandAlone Mode and Neighbor is in StandAlone Mode
The input to the first VNF is in the trunk mode (VNF-tagged) and the neighbor is in StandAlone mode.

User Variables	System Variables	
	Mandatory Variables when Service Chains 1 and 2 are Shared.	**Optional Variables wh Shared**
DNS_SERVER	OTP	
UUID	VBOND_IP	
INSIDE_VLAN1	ORG_NAME	
INSIDE_PRIM_SUBNET1_IP	BGP_NO	
INSIDE_DATA_MASK_LEN1	SYSTEM_IP	
INSIDE_VLAN2	MGMT_PRIM	
INSIDE_PRIM_SUBNET2_IP	MGMT_MASK_LEN	
INSIDE_DATA_MASK_LEN2	MGMT_GW	
INSIDE_GW1	RCC	
INSIDE_GW2	VM_INSTANCE_NAME	
SERVICE_VPN	OUTSIDE_PRIM	OUTSIDE_PRIM_3
SERVICE_VPN_2	OUTSIDE_DATA_MASK_LEN	OUTSIDE_DATA_MA

User Variables	System Variables	
SERVICE_VPN_3	OUTSIDE_PEER_DATA_IP_PRIM	OUTSIDE_PEER_DATA_
SERVICE_VPN_4	OUTSIDE_AS	OUTSIDE_AS_3
	OUTSIDE_VLAN_1	OUTSIDE_VLAN_3
	OUTSIDE_PRIM_2	OUTSIDE_PRIM_4
	OUTSIDE_DATA_MASK_LEN_2	OUTSIDE_DATA_MASK
	OUTSIDE_PEER_DATA_IP_PRIM_2	OUTSIDE_PEER_DATA_
	OUTSIDE_AS_2	OUTSIDE_AS_4
	OUTSIDE_VLAN_2	OUTSIDE_VLAN_4

Cisco vEdge Router Variable List

Cisco vEdge Router is in StandAlone Mode and Neighbor is in HA Mode
The input to the first VNF is in the trunk mode (VNF-tagged) and the neighbor is in HA mode.

User Variables	System Variables	
	Mandatory Variables when Service Chains 1 and 2 are Shared.	**Optional Variables when Shared.**
DNS_SERVER	OTP	
UUID	VBOND_IP	
INSIDE_VLAN1	ORG_NAME	
INSIDE_PRIM_SUBNET1_IP	BGP_NO	
INSIDE_DATA_MASK_LEN1	SYSTEM_IP	
INSIDE_VLAN2	MGMT_PRIM	
INSIDE_PRIM_SUBNET2_IP	MGMT_MASK_LEN	
INSIDE_DATA_MASK_LEN2	MGMT_GW	
INSIDE_GW1	RCC	
INSIDE_GW2	VM_INSTANCE_NAME	
SERVICE_VPN	OUTSIDE_PRIM	OUTSIDE_PRIM_3
SERVICE_VPN_2	OUTSIDE_DATA_MASK_LEN	OUTSIDE_DATA_MASK
SERVICE_VPN_3	OUTSIDE_PEER_DATA_IP_PRIM	OUTSIDE_PEER_DATA_
SERVICE_VPN_4	OUTSIDE_AS	OUTSIDE_AS_3

User Variables	System Variables	
	OUTSIDE_VLAN_1	OUTSIDE_VLAN_3
	OUTSIDE_PEER_DATA_IP_SEC	OUTSIDE_PEER_DA
	OUTSIDE_PRIM_2	OUTSIDE_PRIM_4
	OUTSIDE_DATA_MASK_LEN_2	OUTSIDE_DATA_MA
	OUTSIDE_PEER_DATA_IP_PRIM_2	OUTSIDE_PEER_DA
	OUTSIDE_AS_2	OUTSIDE_AS_4
	OUTSIDE_VLAN_2	OUTSIDE_VLAN_4
	OUTSIDE_PEER_DATA_IP_SEC_2	OUTSIDE_PEER_DA

Cisco CSR1000V Variable List

Cisco CSR1000V Variable List

Last Cisco CSR1000V VNF is in HA Mode and Neighbor is in StandAlone Mode
The output from the last VNF is in access mode (hypervisor-tagged) and the neighbor (ASAv firewall) is in StandAlone mode.

User Variables	System Variables	
	Mandatory Variables	**Optional Variables**
DOMAIN_NAME	VM_INSTANCE_NAME	
DNS_SERVER	TNAME	
NTP_SERVER	ORG_NAME	
TIMEZONE	BGP_NO	
OFFSET	SYSTEM_IP	
SUMMER_TIMEZONE	MGMT_PRIM	
TECH_PACKAGE	MGMT_MASK	
THROUGHPUT_IN_MB	MGMT_GW	
TOKEN_VALUE	MGMT_SEC	
PASS	INSIDE_VLAN_1	INSIDE_VLAN_3
OUTSIDE_PRIM	INSIDE_PRIM	INSIDE_PRIM_3
OUTSIDE_DATA_MASK	INSIDE_DATA_MASK	INSIDE_DATA_MASI

User Variables	System Variables	
OUTSIDE_PEER_DATA_IP_PRIM	INSIDE_PEER_DATA_IP_PRIM	INSIDE_PEER_DATA_IP
OUTSIDE_AS	INSIDE_AS	INSIDE_AS_3
LOCAL_OUTSIDE_AS	INSIDE_VLAN_2	INSIDE_VLAN_4
OUTSIDE_PEER_DATA_IP_SEC	INSIDE_PRIM_2	INSIDE_PRIM_4
OUTSIDE_SEC	INSIDE_DATA_MASK_2	INSIDE_DATA_MASK_4
	INSIDE_PEER_DATA_IP_PRIM_2	INSIDE_PEER_DATA_IP
	INSIDE_AS_2	INSIDE_AS_4

Cisco CSR1000V Variable List

Last Cisco CSR1000V VNF is in StandAlone Mode and Neighbor is in StandAlone Mode
The output from the last VNF is in access mode (hypervisor-tagged) and the neighbor is in StandAlone mode.

User Variables	System Variables	
	Mandatory Variables	**Optional Variables**
DOMAIN_NAME	VM_INSTANCE_NAME	
DNS_SERVER	TNAME	
NTP_SERVER	ORG_NAME	
TIMEZONE	BGP_NO	
OFFSET	SYSTEM_IP	
SUMMER_TIMEZONE	MGMT_PRIM	
TECH_PACKAGE	MGMT_MASK	
THROUGHPUT_IN_MB	MGMT_GW	
TOKEN_VALUE	INSIDE_VLAN_1	INSIDE_VLAN_3
PASS	INSIDE_PRIM	INSIDE_PRIM_3
OUTSIDE_PRIM	INSIDE_DATA_MASK	INSIDE_DATA_MASK_3
OUTSIDE_DATA_MASK	INSIDE_PEER_DATA_IP_PRIM	INSIDE_PEER_DATA_IP
OUTSIDE_PEER_DATA_IP_PRIM	INSIDE_AS	INSIDE_AS_3
OUTSIDE_AS	INSIDE_PEER_DATA_IP_SEC	INSIDE_PEER_DATA_IP
LOCAL_OUTSIDE_AS	VIP_IP_ADDRESS	VIP_IP_ADDRESS_3
	INSIDE_SEC	INSIDE_SEC_3

User Variables	System Variables	
	INSIDE_VLAN_2	INSIDE_VLAN_4
	INSIDE_PRIM_2	INSIDE_PRIM_4
	INSIDE_DATA_MASK_2	INSIDE_DATA_MASK
	INSIDE_PEER_DATA_IP_PRIM_2	INSIDE_PEER_DATA
	INSIDE_AS_2	INSIDE_AS_4
	INSIDE_PEER_DATA_IP_SEC_2	INSIDE_PEER_DATA
	VIP_IP_ADDRESS_2	VIP_IP_ADDRESS_4
	INSIDE_SEC_2	INSIDE_SEC_4

Cisco CSR1000V Variable List

Last Cisco CSR1000V VNF is in StandAlone Mode and Neighbor is in HA Mode
The output from the last VNF is in access mode (hypervisor-tagged) and the neighbor is in HA mode.

User Variables	System Variables	
	Mandatory Variables	**Optional Variables**
DOMAIN_NAME	VM_INSTANCE_NAME	
DNS_SERVER	TNAME	
NTP_SERVER	ORG_NAME	
TIMEZONE	BGP_NO	
OFFSET	SYSTEM_IP	
SUMMER_TIMEZONE	MGMT_PRIM	
TECH_PACKAGE	MGMT_MASK	
THROUGHPUT_IN_MB	MGMT_GW	
TOKEN_VALUE	INSIDE_VLAN_1	INSIDE_VLAN_3
PASS	INSIDE_PRIM	INSIDE_PRIM_3
OUTSIDE_PRIM	INSIDE_DATA_MASK	INSIDE_DATA_MASK
OUTSIDE_DATA_MASK	INSIDE_PEER_DATA_IP_PRIM	INSIDE_PEER_DATA
OUTSIDE_PEER_DATA_IP_PRIM	INSIDE_AS	INSIDE_AS_3
OUTSIDE_AS	INSIDE_VLAN_2	INSIDE_VLAN_4
LOCAL_OUTSIDE_AS	INSIDE_PRIM_2	INSIDE_PRIM_4

User Variables	System Variables	
	INSIDE_DATA_MASK_2	INSIDE_DATA_MASK_4
	INSIDE_PEER_DATA_IP_PRIM_2	INSIDE_PEER_DATA_IP
	INSIDE_AS_2	INSIDE_AS_4

Cisco CSR1000V Variable List

Last Cisco CSR1000V VNF is in HA Mode and Neighbor is in HA Mode
The output from the last VNF is in access mode (hypervisor-tagged) and the neighbor is in HA mode.

User Variables	System Variables	
	Mandatory Variables	**Optional Variables**
DOMAIN_NAME	VM_INSTANCE_NAME	
DNS_SERVER	TNAME	
NTP_SERVER	ORG_NAME	
TIMEZONE	BGP_NO	
OFFSET	SYSTEM_IP	
SUMMER_TIMEZONE	MGMT_PRIM	
TECH_PACKAGE	MGMT_MASK	
THROUGHPUT_IN_MB	MGMT_GW	
TOKEN_VALUE	MGMT_SEC	
PASS	INSIDE_VLAN_1	INSIDE_VLAN_3
OUTSIDE_PRIM	INSIDE_PRIM	INSIDE_PRIM_3
OUTSIDE_DATA_MASK	INSIDE_DATA_MASK	INSIDE_DATA_MASK_3
OUTSIDE_PEER_DATA_IP_PRIM	INSIDE_PEER_DATA_IP_PRIM	INSIDE_PEER_DATA_IP
OUTSIDE_AS	INSIDE_AS	INSIDE_AS_3
LOCAL_OUTSIDE_AS	INSIDE_PEER_DATA_IP_SEC	INSIDE_PEER_DATA_IP
OUTSIDE_PEER_DATA_IP_SEC	VIP_IP_ADDRESS	VIP_IP_ADDRESS_3
OUTSIDE_SEC	INSIDE_SEC	INSIDE_SEC_3
	INSIDE_VLAN_2	INSIDE_VLAN_4
	INSIDE_PRIM_2	INSIDE_PRIM_4
	INSIDE_DATA_MASK_2	INSIDE_DATA_MASK_4

User Variables	System Variables	
	INSIDE_PEER_DATA_IP_PRIM_2	INSIDE_PEER_DATA
	INSIDE_AS_2	INSIDE_AS_4
	INSIDE_PEER_DATA_IP_SEC_2	INSIDE_PEER_DATA
	VIP_IP_ADDRESS_2	VIP_IP_ADDRESS_4

ASAv Variable List

Note In the following ASAv variable list, same variable names can be used for service chains five and six respectively with appropriate renumbering as mentioned for service chains.

ASAv Variable List

First ASAv VNF is in HA Mode and Neighbor is in HA Mode
The input to the first VNF is in access mode (hypervisor-tagged) and the neighbor is in HA mode.

User Variables	System Variables	
	Mandatory Variables when Service Chains 1 and 2 are Shared.	**Optional Variables wh Shared.**
DNS_SERVER	OTP	
OFFSET	VBOND_IP	
SUMMER_TIMEZONE	ORG_NAME	
DOMAIN_NAME	BGP_NO	
NTP_SERVER_NAME	SYSTEM_IP	
LIC_LEVEL	RCC	
ID_TOKEN	VM_INSTANCE_NAME	
PASS	TNAME	
TIMEZONE	HA_PRIM_IP	
INSIDE_PRIM	HA_SEC_IP	
INSIDE_SEC	HA_MASK	
INSIDE_DATA_MASK	MGMT_PRIM	
INSIDE_PEER_DATA_IP_PRIM	MGMT_MASK	

User Variables	System Variables	
INSIDE_PEER_DATA_IP_SEC	MGMT_GW	
INSIDE_AS	MGMT_SEC	
LOCAL_INSIDE_AS	OUTSIDE_PRIM	OUTSIDE_PRIM_3
	OUTSIDE_DATA_MASK	OUTSIDE_DATA_MASK
	OUTSIDE_PEER_DATA_IP_PRIM	OUTSIDE_PEER_DATA_
	OUTSIDE_AS	OUTSIDE_AS_3
	OUTSIDE_VLAN_1	OUTSIDE_VLAN_3
	OUTSIDE_PEER_DATA_IP_SEC	OUTSIDE_PEER_DATA_
	OUTSIDE_SEC	OUTSIDE_SEC_3
	OUTSIDE_PRIM_2	OUTSIDE_PRIM_4
	OUTSIDE_DATA_MASK_2	OUTSIDE_DATA_MASK
	OUTSIDE_PEER_DATA_IP_PRIM_2	OUTSIDE_PEER_DATA_
	OUTSIDE_AS_2	OUTSIDE_AS_4
	OUTSIDE_VLAN_2	OUTSIDE_VLAN_4
	OUTSIDE_PEER_DATA_IP_SEC_2	OUTSIDE_PEER_DATA_
	OUTSIDE_SEC_2	OUTSIDE_SEC_4

ASAv Variable List

First ASAv VNF is in StandAlone Mode and Neighbor is in StandAlone Mode

The input to the first VNF is in access mode (hypervisor-tagged) and the neighbor is in StandAlone mode.

User Variables	System Variables	
	Mandatory Variables when Service Chains 1 and 2 are Shared.	**Optional Variables when Shared.**
DNS_SERVER	OTP	
OFFSET	VBOND_IP	
SUMMER_TIMEZONE	ORG_NAME	
DOMAIN_NAME	BGP_NO	
NTP_SERVER_NAME	SYSTEM_IP	
LIC_LEVEL	RCC	

User Variables	System Variables	
ID_TOKEN	VM_INSTANCE_NAME	
PASS	TNAME	
TIMEZONE	MGMT_PRIM	
INSIDE_PRIM	MGMT_MASK	
INSIDE_DATA_MASK	MGMT_GW	
INSIDE_PEER_DATA_IP_PRIM	OUTSIDE_PRIM	OUTSIDE_PRIM_3
INSIDE_AS	OUTSIDE_DATA_MASK	OUTSIDE_DATA_MA
LOCAL_INSIDE_AS	OUTSIDE_PEER_DATA_IP_PRIM	OUTSIDE_PEER_DA
	OUTSIDE_AS	OUTSIDE_AS_3
	OUTSIDE_VLAN_1	OUTSIDE_VLAN_3
	OUTSIDE_PRIM_2	OUTSIDE_PRIM_4
	OUTSIDE_DATA_MASK_2	OUTSIDE_DATA_MA
	OUTSIDE_PEER_DATA_IP_PRIM_2	OUTSIDE_PEER_DA
	OUTSIDE_AS_2	OUTSIDE_AS_4
	OUTSIDE_VLAN_2	OUTSIDE_VLAN_4

ASAv Variable List

First ASAv VNF is in StandAlone Mode and Neighbor is in HA Mode
The input to the first VNF is in access mode (hypervisor-tagged) and the neighbor is in HA mode.

User Variables	System Variables	
	Mandatory Variables when Service Chains 1 and 2 are Shared.	**Optional Variables wh Shared.**
DNS_SERVER	OTP	
OFFSET	VBOND_IP	
SUMMER_TIMEZONE	ORG_NAME	
DOMAIN_NAME	BGP_NO	
NTP_SERVER_NAME	SYSTEM_IP	
LIC_LEVEL	RCC	
ID_TOKEN	VM_INSTANCE_NAME	

User Variables	System Variables	
PASS	TNAME	
TIMEZONE	MGMT_PRIM	
INSIDE_PRIM	MGMT_MASK	
INSIDE_DATA_MASK	MGMT_GW	
INSIDE_PEER_DATA_IP_PRIM	OUTSIDE_PRIM	OUTSIDE_PRIM_3
INSIDE_AS	OUTSIDE_DATA_MASK	OUTSIDE_DATA_MASK
LOCAL_INSIDE_AS	OUTSIDE_PEER_DATA_IP_PRIM	OUTSIDE_PEER_DATA_
	OUTSIDE_AS	OUTSIDE_AS_3
	OUTSIDE_VLAN_1	OUTSIDE_VLAN_3
	OUTSIDE_PEER_DATA_IP_SEC	OUTSIDE_PEER_DATA_
	OUTSIDE_PRIM_2	OUTSIDE_PRIM_4
	OUTSIDE_DATA_MASK_2	OUTSIDE_DATA_MASK
	OUTSIDE_PEER_DATA_IP_PRIM_2	OUTSIDE_PEER_DATA_
	OUTSIDE_AS_2	OUTSIDE_AS_4
	OUTSIDE_VLAN_2	OUTSIDE_VLAN_4
	OUTSIDE_PEER_DATA_IP_SEC_2	OUTSIDE_PEER_DATA_

ASAv Variable List

First ASAv VNF is in HA Mode and Neighbor is in HA Mode

The input to the first VNF is in the trunk mode (vnf-tagged) and the neighbor is in HA mode.

User Variables	System Variables	
	Mandatory Variables when Service Chains 1 and 2 are Shared.	**Optional Variables when Shared.**
DNS_SERVER	OTP	
OFFSET	VBOND_IP	
SUMMER_TIMEZONE	ORG_NAME	
DOMAIN_NAME	BGP_NO	
NTP_SERVER_NAME	SYSTEM_IP	
LIC_LEVEL	RCC	

User Variables	System Variables	
ID_TOKEN	VM_INSTANCE_NAME	
PASS	TNAME	
TIMEZONE	HA_PRIM_IP	
INSIDE_PRIM_SUBNET1_IP	HA_SEC_IP	
INSIDE_PEER_DATA_IP_PRIM1	HA_MASK	
INSIDE_AS1	MGMT_PRIM	
LOCAL_INSIDE_AS1	MGMT_MASK	
INSIDE_VLAN1	MGMT_GW	
INSIDE_DATA_MASK_SUBNET1	MGMT_GW	
INSIDE_PRIM_SUBNET2_IP	OUTSIDE_PRIM	OUTSIDE_PRIM_3
INSIDE_PEER_DATA_IP_PRIM2	OUTSIDE_DATA_MASK	OUTSIDE_DATA_MA
INSIDE_AS2	OUTSIDE_PEER_DATA_IP_PRIM	OUTSIDE_PEER_DA
LOCAL_INSIDE_AS2	OUTSIDE_AS	OUTSIDE_AS_3
INSIDE_VLAN2	OUTSIDE_VLAN_1	OUTSIDE_VLAN_3
INSIDE_DATA_MASK_SUBNET2	OUTSIDE_PEER_DATA_IP_SEC	OUTSIDE_PEER_DA
INSIDE_PRIM_SUBNET3_IP	OUTSIDE_SEC	OUTSIDE_SEC_3
INSIDE_PEER_DATA_IP_PRIM3	OUTSIDE_PRIM_2	OUTSIDE_PRIM_4
INSIDE_AS3	OUTSIDE_DATA_MASK_2	OUTSIDE_DATA_MA
LOCAL_INSIDE_AS3	OUTSIDE_PEER_DATA_IP_PRIM_2	OUTSIDE_PEER_DA
INSIDE_VLAN3	OUTSIDE_AS_2	OUTSIDE_AS_4
INSIDE_DATA_MASK_SUBNET3	OUTSIDE_VLAN_2	OUTSIDE_VLAN_4
INSIDE_PRIM_SUBNET4_IP	OUTSIDE_PEER_DATA_IP_SEC_2	OUTSIDE_PEER_DA
INSIDE_PEER_DATA_IP_PRIM4	OUTSIDE_SEC_2	OUTSIDE_SEC_4
INSIDE_AS4		
LOCAL_INSIDE_AS4		
INSIDE_VLAN4		
INSIDE_DATA_MASK_SUBNET4		