

Cisco Enterprise Network Function Virtualization Infrastructure Software Configuration Guide, Release 3.5.x

First Published: 2017-03-16

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

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

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

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

NOTWITHSTANDING ANY OTHER WARRANTY HEREIN, ALL DOCUMENT FILES AND SOFTWARE OF THESE SUPPLIERS ARE PROVIDED “AS IS" WITH ALL FAULTS. CISCO AND THE ABOVE-NAMED SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, THOSE OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND 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)

© 2017 Cisco Systems, Inc. All rights reserved.



CONTENTS

Full Cisco Trademarks with Software License ?

PREFACE

About this Document ix

About this Document ix

Audience x

Related Documentation x

Communications, Services, and Additional Information xi

CHAPTER 1

About Cisco Enterprise NFVIS 1

Benefits of Cisco Enterprise NFVIS 1

Supported Hardware Platforms 2

Supported VMs 3

Key Tasks You can Perform Using Cisco Enterprise NFVIS 4

CHAPTER 2

Installing Cisco Enterprise NFVIS Using the KVM Console 5

Installation Prerequisites 5

Image Signing and Verification 6

RPM Signing 6

RPM Signature Verification 6

Image Integrity Verification Using sha256sum 6

Entering BIOS Setup 7

Installing Cisco Enterprise NFVIS on the Cisco UCS C220 M4 and M5 Rack Servers or Cisco CSP Devices 7

Logging Into the CIMC GUI 7

Installing Cisco Enterprise NFVIS on Cisco UCS E-Series Servers 8

Sample Configuration on the Cisco ISR Router to Bring Up a Cisco UCS E Server 10

Installing Cisco Enterprise NFVIS on a Cisco ENCS 5100 and 5400 12
Install Cisco Enterprise NFVIS on Cisco ENCS 5104 14

CHAPTER 3 Setting Up System Configuration 17
Default System Configuration on the Cisco ENCS 17
Default System Configuration on the Cisco UCS C220 M4 Server and Cisco CSP 2100 19
Default System Configuration on the Cisco UCS E-Series Servers 20
Accessing NFVIS 20
Configuring VLAN for NFVIS Management Traffic 26
Configuring System Routes 26
User Roles and Authentication 27
Rules for User Passwords 27
Creating Users and Assigning Roles 27
Configuring Minimum Length for Passwords 28
Configuring Password Lifetime 29
Deactivating Inactive User Accounts 29
Activating an Inactive User Account 29
Certification 30
Secure Copy Command 30
Configuring the IP Receive ACL 31
Port 22222 and Management Interface ACL 32
Configuring Your Banner and Message of the Day 32
Setting the System Time Manually or With NTP 33
Enable or Disable Access to NFVIS Portal 34
Configuring System Logs 35
Network File System Support 36
Secure Boot of host 37
CIMC Access Control 38
CIMC Access using NFVIS 38
BIOS-CIMC Update 39
NFVIS Password Recovery 39
DPDK Support on NFVIS 40
DPDK Support for NFVIS 3.10.x 40
DPDK VM Migration for NFVIS 3.11.x 41

Enhancements to DPDK Support in NFVIS 3.12.x 41
Import and Export NFVIS VM 43
Backup and Restore NFVIS and VM Configurations 46
Dynamic SR-IOV 55

CHAPTER 4 Cisco Network Plug-n-Play Support 59
PnP Discovery Methods 60
Configuring PnP Discovery Methods 61
PnP Action 64

CHAPTER 5 Overview of VM Life Cycle Management 65
Workflow of VM Life Cycle Management 65
Uploading VM Images to an NFVIS Server 67
VM Bootstrap Configuration Options with a VM Deployment 68
OpenStack Configuration Drive Support for Third Party VMs 69
Performing Resource Verification 70
Configuring Management IP Subnet 71
VM States 71

CHAPTER 6 VM Deployment Scenarios 73
Registering VM Images 73
Single VM Deployment 74
Steps for Deploying a VM 74
Service Chaining of VMs 77
Service Chaining with two VM Images 77
Steps for Service Chaining with Two VM Images 77
Service Chaining of Multiple VMs with Windows or Linux Servers 78
Steps for Service Chaining of Multiple VMs with Windows or Linux Servers 78

CHAPTER 7 SPAN Session or Port Mirroring 79
About SPAN Sessions 79
Configuring SPAN Sessions 79
Configuration Examples for SPAN Session Scenarios 81
Example: SPAN Session Traffic on a Physical Interface 81

Example: SPAN Session Traffic on a LAN SRIOV 82
Example: SPAN Session Traffic on a VLAN 83

CHAPTER 8 Configuring Packet Capture 85

CHAPTER 9 VM Image Packaging 87
VM Image Packaging Utility 87
Contents 87
Usage 88
NFVIS Specific Enhancements 94
VM Packaging Utility Usage Examples 94
Standard VM Image Packaging 96
Generating a VM Package 96
Appendix 96
VM Image Package Files 96
Package Manifest File 97
Bootstrap Configuration File 98
VM Image Properties File 98
Example: Package.mf 103
Example: Image Properties 103
Example: Bootstrap Configuration File 104
Image Properties Template File 105

CHAPTER 10 Upgrading Cisco Enterprise NFVIS 107

CHAPTER 11 Configuring vBranch High Availability 111
Prerequisites for vBranch HA 111
vBranch HA Design and Topology 112
Enable Virtual NIC Failure Detection with Track Feature 112
Isolating LAN and Transit Link Traffic for vBranch HA 114
Packet Flow for vBranch HA 116
Configuration Examples for vBranch HA 117
Example: Active Cisco ENCS Configuration with ISRv1 117
Example: Standby Cisco ENCS Configuration with ISRv2 119

Cisco ENCS Failure Points 120

CHAPTER 12 Cisco ENCS Single WAN IP Deployment Scenarios 125

Single WAN IP Deployment 125

Preconfiguring the Cisco ENCS for a Single WAN IP Deployment 126

Single WAN IP Deployment with Gigabit Ethernet Interface 0/0 127

Single WAN IP Deployment with the 4G Interface 128

CHAPTER 13 Resetting to Factory Default 131

CHAPTER 14 NETCONF Event Notifications 133

nfvisEvent 134

vmlcEvent 150



About this Document

This guide provides information about install and configure Cisco Enterprise Network Function Virtualization Infrastructure Software (Cisco Enterprise NFVIS) on a supported Cisco hardware device. The guide also provides details on virtual machine deployments, configuration of software features, and life cycle management using Representation State Transfer (REST) application programming interface (API).

This guide has been divided into chapters starting from installing to maintenance and troubleshooting.

- Quick Start your ENCS Hardware
- Install and Upgrade Cisco NFVIS
- Manage your Host System
- VM Lifecycle Management
- Plug n Play
- System Logs
- ENCS 5400 Switch Management
- Maintain and Monitor Cisco NFVIS
- Configure vBranch HA
- Cisco ENCS Single WAN IP Deployment Scenarios
- Troubleshooting and Debugging

- About this Document, on page ix
- Audience, on page x
- Related Documentation, on page x
- Communications, Services, and Additional Information, on page xi

About this Document

This guide provides information about install and configure Cisco Enterprise Network Function Virtualization Infrastructure Software (Cisco Enterprise NFVIS) on a supported Cisco hardware device. The guide also provides details on virtual machine deployments, configuration of software features, and life cycle management using Representation State Transfer (REST) application programming interface (API).

This guide has been divided into chapters starting from installing to maintenance and troubleshooting.

- Quick Start your ENCS Hardware
- Install and Upgrade Cisco NFVIS
- Manage your Host System
- VM Lifecycle Management
- Plug n Play
- System Logs
- ENCS 5400 Switch Management
- Maintain and Monitor Cisco NFVIS
- Configure vBranch HA
- Cisco ENCS Single WAN IP Deployment Scenarios
- Troubleshooting and Debugging

Audience

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

Related Documentation

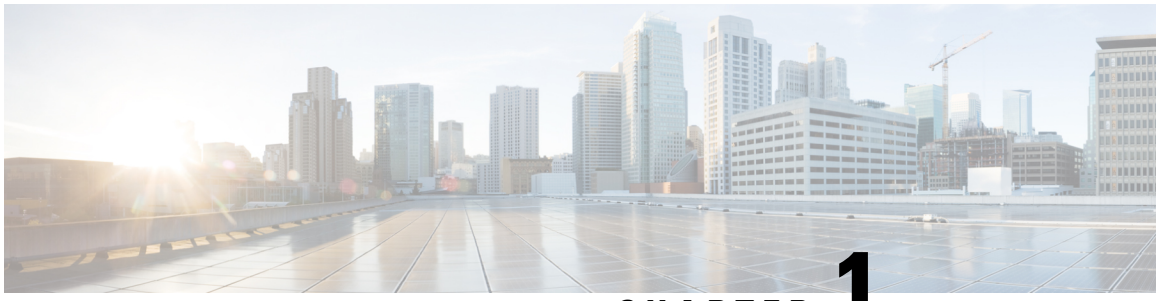
- Cisco Network Function Virtualization Infrastructure Software Getting Started Guide
- API Reference for Cisco Enterprise Network Function Virtualization Infrastructure Software
- Cisco Enterprise Network Function Virtualization Infrastructure Software Configuration Guide, Release 4.x
- Cisco Enterprise Network Function Virtualization Infrastructure Software Command Reference
- Release Notes for Cisco NFV SD-Branch features in Cisco vManage Release 20.12.x
- Design and Deployment Guide of Cisco NFVIS SD-Branch using Cisco SD-WAN Manager
- Cisco Catalyst 8200 Series Edge uCPE Data Sheet
- Cisco Cloud Services Platform 5000 Series Data Sheet
- Cisco 5400 Enterprise Network Compute System Hardware Installation Guide
- Cisco 5400 Enterprise Network Compute System Data Sheet
- Configuration Guide for Cisco Network Plug and Play on Cisco APIC-EM, Release 1.5.x
- Cisco SD-WAN Controller Compatibility Matrix and Recommended Computing Resources, Cisco SD-WAN Release 20.12.x

Communications, Services, and Additional Information

- To receive timely, relevant information from Cisco, sign up at Cisco Profile Manager.
- To get the business results 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 DevNet.
- 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

About Cisco Enterprise NFVIS

Cisco Enterprise Network Function Virtualization Infrastructure Software (Cisco Enterprise NFVIS) is a Linux-based infrastructure software designed to help service providers and enterprises to design, deploy and manage network services. Cisco Enterprise NFVIS helps dynamically deploy virtualized network functions, such as a virtual router, firewall, and WAN accelerator on supported Cisco devices. Such virtualized deployments of VNFs also leads to device consolidation. You no longer need separate devices. Automated provisioning and centralized management also eliminates costly truck rolls.

Cisco Enterprise NFVIS provides a Linux-based virtualization layer to the Cisco Enterprise Network Function Virtualization (ENFV) solution.

Cisco ENFV Solution Overview

The Cisco ENFV solution helps convert your critical network functions into a software which can deploy network services across dispersed locations in minutes. It provides a fully integrated platform that can run on top of a diverse network of both virtual and physical devices with the following primary components:

- Cisco Enterprise NFVIS
- VNFs
- Unified Computing System (UCS) and Enterprise Network Compute System (ENCS) hardware platforms
- Digital Network Architecture Center (DNAC)

- Benefits of Cisco Enterprise NFVIS, on page 1
- Supported Hardware Platforms, on page 2
- Supported VMs, on page 3
- Key Tasks You can Perform Using Cisco Enterprise NFVIS, on page 4

Benefits of Cisco Enterprise NFVIS

- Consolidates multiple physical network appliances into a single server running multiple virtual network functions.
- Deploys services quickly and in a timely manner.
- Cloud based VM life cycle management and provisioning.
- Life cycle management to deploy and chain VMs dynamically on the platform.

- Programmable APIs.

Supported Hardware Platforms

Depending on your requirement, you can install Cisco Enterprise NFVIS on the following Cisco hardware platforms:

- Cisco 5100 Series Enterprise Network Compute System (Cisco ENCS)
- Cisco 5400 Series Enterprise Network Compute System (Cisco ENCS)
- Cisco Catalyst 8200 Series Edge Universal CPE
- Cisco Catalyst 8300 Series Edge Universal CPE
- Cisco UCS C220 M4 Rack Server
- Cisco UCS C220 M5Rack Server
- Cisco UCS C M6 Rack Servers (UCSC-C220-M6S, UCSC-C240-M6SX, and UCSC-C240-M6S)
- Cisco Cloud Services Platform 2100 (CSP 2100)
- Cisco Cloud Services Platform 5228 (CSP-5228), 5436 (CSP-5436) and 5444 (CSP-5444 Beta)
- Cisco ISR4331 with UCS-E140S-M2/K9
- Cisco ISR4351 with UCS-E160D-M2/K9
- Cisco ISR4451-X with UCS-E180D-M2/K9
- Cisco UCS-E160S-M3/K9 Server
- Cisco UCS-E180D-M3/K9
- Cisco UCS-E1120D-M3/K9

Cisco ENCS

The Cisco 5100 and 5400 Series Enterprise Network Compute System combines routing, switching, storage, processing, and a host of other computing and networking activities into a compact one Rack Unit (RU) box. This high-performance unit achieves this goal by providing the infrastructure to deploy virtualized network functions and acting as a server that addresses processing, workload, and storage challenges.

Cisco Catalyst 8200 Series Edge Universal CPE

The Cisco Catalyst 8200 Edge uCPE is the next generation of Cisco Enterprise Network Compute System 5100 Series that combines routing, switching and application hosting into a compact one rack unit device for the small and Medium Virtualized Branch. These platforms are designed to allow customers to run virtualized network functions and other applications as virtual machines on the same hardware platform powered by Cisco NFVIS hypervisor software. These devices are 8 Core x86 CPUs with HW Acceleration for IPSec crypto traffic with higher number of WAN ports. They have a NIM slot and a PIM slot to choose different WAN, LAN and LTE/5G modules for the Branch.

Cisco Catalyst 8300 Series Edge Universal CPE

The Cisco Catalyst 8300 Series Edge Universal Customer Premises Equipment (uCPE) is a purpose-built x86 platform that is designed for branch virtualization. It enables device consolidation across network and security functions, improves operational flexibility and service agility, simplifies network operations, and results in reduced deployment times and fewer truck rolls for delivery of add-on services.

Cisco UCS C220 M4/M5 Rack Server

The Cisco UCS C220 M4 Rack Server is a high-density, general-purpose enterprise infrastructure and application server that delivers world class performance for a wide range of enterprise workloads, including virtualization, collaboration, and bare-metal applications.

Cisco UCS C M6 Rack Server

The Cisco UCS C220 M6 Rack Server is the most versatile general-purpose infrastructure and application server in the industry. This high-density, 1RU, 2-socket rack server delivers industry-leading performance and efficiency for a wide range of workloads, including virtualization, collaboration, and bare-metal applications.

Cisco CSP 2100-X1, 5228, 5436 and 5444 (Beta)

Cisco Cloud Services Platform is a software and hardware platform for data center network functions virtualization. This open kernel virtual machine (KVM) platform is designed to host networking virtual services. Cisco Cloud Services Platform devices enables network, security, and load balancer teams to quickly deploy any Cisco or third-party network virtual service.

Note CSP 5000 series devices support ixgbe drivers.

Caution If CSP platforms are running NFVIS, Return Material Authorization (RMA) is not supported.

Cisco UCS E-Series Server Modules

The Cisco UCS E-Series Servers (E-Series Servers) are the next generation of Cisco UCS Express servers. E-Series Servers are a family of size, weight, and power efficient blade servers that are housed within the Generation 2 Cisco Integrated Services Routers (ISR G2), Cisco 4400, and Cisco 4300 Series Integrated Services Routers. These servers provide a general-purpose compute platform for branch office applications deployed either as bare metal on operating systems, such as Microsoft Windows or Linux; or as virtual machines on hypervisors.

Supported VMs

Currently, Cisco Enterprise NFVIS supports the following Cisco VMs and third-party VMs:

- Cisco Catalyst 8000V Edge Software
- Cisco Integrated Services Virtual (ISRv)
- Cisco Adaptive Security Virtual Appliance (ASAv)

- Cisco Virtual Wide Area Application Services (vWAAS)
- Linux Server VM
- Windows Server 2012 VM
- Cisco Firepower Next-Generation Firewall Virtual (NGFWv)
- Cisco vEdge
- Cisco XE SD-WAN
- Cisco Catalyst 9800 Series Wireless Controller
- ThousandEyes
- Fortinet
- Palo Alto
- CTERA
- InfoVista

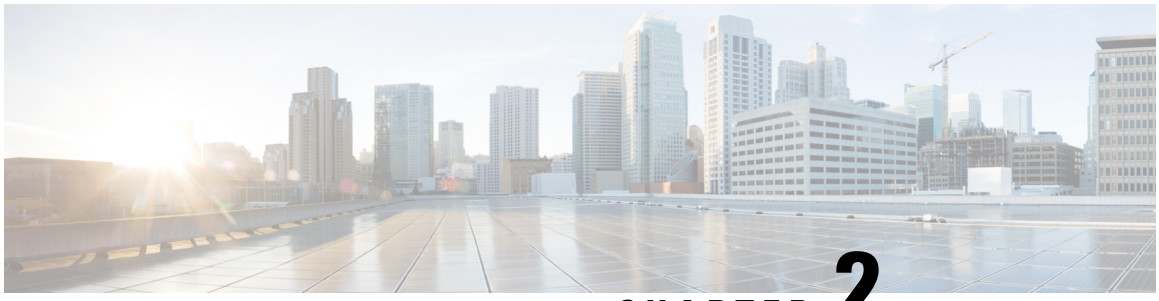
Key Tasks You can Perform Using Cisco Enterprise NFVIS

- Perform VM image registration and deployment
- Create new networks and bridges, and assign ports to bridges
- Perform service chaining of VMs
- Perform VM operations
- Verify system information including CPU, port, memory, and disk statistics
- SR-IOV support on all interfaces of all platforms, with the exception of UCS-E backplane interface

The APIs for performing these tasks are explained in the API Reference for Cisco Enterprise NFVIS.

Note NFVIS can be configured through Netconf interface, REST APIs and command-line interface as all the configurations are exposed through YANG models.

From a Cisco Enterprise NFVIS command-line interface, you can connect to another server and VMs remotely using the SSH client.



CHAPTER 2

Installing Cisco Enterprise NFVIS Using the KVM Console

- Installation Prerequisites , on page 5
- Image Signing and Verification, on page 6
- Entering BIOS Setup, on page 7
- Installing Cisco Enterprise NFVIS on the Cisco UCS C220 M4 and M5 Rack Servers or Cisco CSP Devices, on page 7
- Installing Cisco Enterprise NFVIS on Cisco UCS E-Series Servers, on page 8
- Installing Cisco Enterprise NFVIS on a Cisco ENCS 5100 and 5400, on page 12

Installation Prerequisites

Ensure that the following prerequisties are met:

- The IP address is configured for Cisco Integrated Management Controller (CIMC) as well as a login account with administrative privileges.
- The Cisco Enterprise NFVIS ISO image is available.
- The IP address of the system (required for remote access) is available.
- Hyper-threading is enabled in BIOS. By default, hyper-threading is enabled in BIOS on the UCS-C, UCS-E and ENCS platforms.

Note The installation steps are slightly different for Cisco UCS and Cisco ENCS platforms. See the following sections for details:

Installing Cisco Enterprise NFVIS on the Cisco UCS C220 M4 and M5 Rack Servers or Cisco CSP Devices, on page 7

Installing Cisco Enterprise NFVIS on Cisco UCS E-Series Servers, on page 8

Installing Cisco Enterprise NFVIS on a Cisco ENCS 5100 and 5400, on page 12

Assumptions

- The user is familiar with the supported hardware device, CIMC, Cisco Network Plug and Play, and Cisco Application Policy Infrastructure Controller Enterprise Module (APIC-EM).
- The initial setup of the hardware device is complete, and the device is ready for loading Cisco Enterprise NFVIS.
- The user is familiar with general Linux installation.

For more details on the supported hardware devices, see respective documentation available on Cisco.com.

Image Signing and Verification

Cisco Enterprise NFVIS supports RPM signing and signature verification for all RPM packages in the ISO and upgrade images. You can also verify the integrity of the Cisco Enterprise NFVIS ISO and upgrade images.

RPM Signing

All RPM packages in the Cisco Enterprise NFVIS ISO and upgrade images are signed to ensure cryptographic integrity and authenticity. This guarantees that the RPM packages have not been tampered with and the RPM packages are from Cisco Enterprise NFVIS. The private key, used for signing the RPM packages, is created and securely maintained by Cisco.

RPM Signature Verification

Cisco Enterprise NFVIS verifies all RPM packages during installation or upgrade. The following table describes the Cisco Enterprise NFVIS behavior when the signature verification fails during installation or upgrade.

Scenario	Description
Cisco Enterprise NFVIS 3.7.1 installation	If the signature verification fails while installing Cisco Enterprise NFVIS, the installation is aborted.
Cisco Enterprise NFVIS upgrade from 3.6.x to Release 3.7.1	The RPM signatures are verified when the upgrade is being performed. If the signature verification fails, an error is logged but the upgrade is completed.
Cisco Enterprise NFVIS upgrade from Release 3.7.1 to later releases	The RPM signatures are verified when the upgrade image is registered. If the signature verification fails, the upgrade is aborted.

Image Integrity Verification Using sha256sum

RPM signing and signature verification can be done only for the RPM packages available in the Cisco NFVIS ISO and upgrade images. To ensure the integrity of all additional non-RPM files available in the Cisco NFVIS ISO image, a hash of the Cisco NFVIS ISO image is published along with the image. Similarly, a hash of the Cisco NFVIS upgrade image is published along with the image. To verify that the hash of Cisco NFVIS ISO image or upgrade image matches the hash published by Cisco, run the following command and compare the hash with the published hash:

```
% /usr/bin/sha256sum ImageFile
4db533d96d8705db8af904ab754349151adea504b81337155cc591c6203e3295 ImageFile
```

Entering BIOS Setup



Note This section applies only to ENCS 5400 and UCS-E series routers.

When you enter the BIOS setup for the first time, ensure that you secure the BIOS by setting up an admin-level and a user-level password. You have to set up the admin password when you access the BIOS menu for the first time. The user password (which only gives access to a small subset of BIOS options) must be set inside the BIOS setup menu.

To set up the admin password, press F2 when the system boots up. You will be prompted to set the password.

To set up the user password, after you log in, go to the 'Security' tab and set the password.

Installing Cisco Enterprise NFVIS on the Cisco UCS C220 M4 and M5 Rack Servers or Cisco CSP Devices

This section provides information about a series of tasks you need to perform to install Cisco Enterprise NFVIS on a Cisco UCS C220 M4 and M5 Rack Server or Cisco CSP 2100-X1, CSP 5228 and CSP 5436.

Logging Into the CIMC GUI

Before you begin

- Ensure that you have configured the IP address to access CIMC as well as a login account with administrative privileges.
- The recommended CIMC version is 3.0(3c) or newer.
- Ensure that a Virtual Drive is created with RAID and Boot Drive is set to true.

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.

Procedure

Step 1 Log into CIMC.

Step 2 Launch KVM Console:

From theCIMC Homepage, click Launch KVM on the upper right corner.You can choose from Java based KVM or HTML based KVM. It is recommnded to use HTML based KVM. Ensure that the pop-up blocker is disabled. KVM Console will open in a separate window.

Step 3 Map virtual media:

a) From KVM Console, click on **Virtual Media** and then click on **Activate Virtual Devices**.
b) Click on **Virtual Media** again and then click on **Map CD/DVD**, browse, and select the Cisco Enterprise NFVIS ISO image, click **Open** and Map Drive to mount the image.
c) Click on **Virtual Media** again to ensure the NFVIS ISO image is now mapped to CD/DVD.

Step 4 Configure Boot Order:

a) From the **CIMC Compute**, select **BIOS**.
b) Select **Configure Boot Order** and click on **Configure Boot Order**.

The **Configure Boot Order** dialog box appears.

c) Select **Advanced**.
d) From **Add Boot Device**, select **Add Virtual Media**, and the **Add Virtual Media** dialog box appears.
e) Enter a name, select **KVM Mapped DVD**, set state to **Enabled** and order to 1, click on **Save Changes**.
f) From **Add Boot Device**, select **Add Local HDD**, the **Add Virtual Media** dialog box appears.
g) Enter a name, set Sstate to Enabled and order to 2, click on **Save Changes**.
h) Click **Close**.

Step 5 Power cycle server to start the installation:

From CIMC Home page, click on **Host Power** on the upper right corner. Reboot the server by selecting the **Power Off** option. After the server is down, select the **Power On** option.

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

Step 6 Initial login: After the installation is complete, the system automatically reboots from the hard drive. Log into the system when the command prompt **nfvis login** is displayed after the reboot.

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 attempt. 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 7 Verification: You can verify the installation using the System API, CLI, or by viewing the system information from the Cisco Enterprise NFV portal.

Installing Cisco Enterprise NFVIS on Cisco UCS E-Series Servers

Before you begin

- Configure the Gigabit Ethernet interface on the Cisco ISR router.
- Configure the UCS E interface on the Cisco ISR router.
- Ensure that you have the IP address configured for CIMC access as well as a login account with administrative privileges.

- The recommneded CIMC version is 3.2(7) or newer.
- Ensure that a Virtual Drive is created with RAID and Boot Drive is set to true.
- Ensure that the Cisco UCS E server has one of the following supported firmware versions or above:
  - BIOS UCSED.2.5.0.3 or later for UCS-E160D-M2/K9 and UCS-E180D-M2/K9
  - BIOS UCSES.1.5.0.5 or later for UCS-E140S-M2/K9

For more details on how to perform the basic configuration on the Cisco ISR routers, see the following guides:

- Sample Configuration on the Cisco ISR Router to Bring Up a Cisco UCS E Server, on page 10
- Getting Started Guide for Cisco UCS E-Series Servers, Release 1.0(2) Installed in the Cisco ISR 4451-X

For details on how to configure an IP address for CIMC, see the Getting Started Guide for Cisco UCS E-Series Servers, Release 1.0 on cisco.com.

Procedure

Step 1 Log into CIMC.

Step 2 Launch KVM Console:

From theCIMC Homepage, click Launch KVM on the upper right corner.You can choose from Java based KVM or HTML based KVM. It is recommnded to use HTML based KVM. Ensure that the pop-up blocker is disabled. KVM Console will open in a separate window.

Step 3 Map virtual media:

a) From KVM Console, click on **Virtual Media** and then click on **Activate Virtual Devices**.
b) Click on **Virtual Media** again and then click on **Map CD/DVD**, browse, and select the Cisco Enterprise NFVIS ISO image, click **Open** and Map Drive to mount the image.
c) Click on **Virtual Media** again to ensure the NFVIS ISO image is now mapped to CD/DVD.

Step 4 Configure Boot Order:

a) From the **CIMC Compute**, select **BIOS**.
b) Select **Configure Boot Order** and click on **Configure Boot Order**.

   The **Configure Boot Order** dialog box appears.

c) From the **CD/DVD** area, select **Cisco vKVM-Mapped vDVD**, and then click **Add**.
d) From **HDD**, select **RAID Adapter**, and then click **Add**.
e) Set the boot order sequence using the **Up** and **Down** options. The Cisco vKVM-Mapped vDVD boot order option must be the first choice. Click **Save Changes** to complete the boot order setup.

Step 5 Power cycle server to start the installation:

From CIMC Home page, click on **Host Power** on the upper right corner. Reboot the server by selecting the **Power Off** option. After the server is down, select the **Power On** option.

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

Step 6 Initial login: After the installation is complete, the system automatically reboots from the hard drive. Log into the system when the command prompt **nfvis login** is displayed after the reboot.

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 attempt. 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 7 Verification: You can verify the installation using the System API, CLI, or by viewing the system information from the Cisco Enterprise NFV portal.

What to do next

You can verify the default configuration, and set up initial IP configuration to launch the Cisco Enterprise NFV portal. For details, see Performing Initial System Configuration.

Sample Configuration on the Cisco ISR Router to Bring Up a Cisco UCS E Server

The following sample configuration shows the basic configuration performed on the Cisco ISR 4451 router with DHCP enabled.

```
Last configuration change at 02:36:37 UTC Thu Feb 18 2016
!
version 15.5
service timestamps debug datetime msec
service timestamps log datetime msec
no platform punt-keepalive disable-kernel-core
!
hostname NFVIS-ISR4451
!
boot-start-marker
boot system bootflash:isr4300-universalk9.03.16.01a.S.155-3.S1a-ext.SPA.bin
boot-end-marker
!
!
vrf definition Mgmt-intf
 !
 address-family ipv4
 exit-address-family
 !
 address-family ipv6
 exit-address-family
!
!
no aaa new-model
!
!
!
ip domain name cisco.com
!
!
!
```

```
subscriber templating
!
multilink bundle-name authenticated
!
!
!
license udi pid ISR4331/K9 sn FDO192207MN
!
!
ucse subslot 1/0
 imc access-port shared-lom console
 imc ip address 172.19.183.172 255.255.255.0 default-gateway 172.19.183.1
!
spanning-tree extend system-id
!
!
redundancy
 mode none
!
!
!
vlan internal allocation policy ascending
!
!
!
interface GigabitEthernet0/0/0
 ip address 172.19.183.171 255.255.255.0
 media-type rj45
 negotiation auto
!
interface GigabitEthernet0/0/1
 no ip address
 shutdown
 negotiation auto
!
interface GigabitEthernet0/0/2
 no ip address
 shutdown
 negotiation auto
!
interface ucse1/0/0
 ip unnumbered GigabitEthernet0/0/0
 negotiation auto
 switchport mode trunk
 no mop enabled
 no mop sysid
!
interface ucse1/0/1
 no ip address
 no negotiation auto
 switchport mode trunk
 no mop enabled
 no mop sysid
!
interface GigabitEthernet0
 vrf forwarding Mgmt-intf
 no ip address
 shutdown
 negotiation auto
!
interface Vlan1
 no ip address
 shutdown
!
```

```
ip default-gateway 172.19.183.1
ip forward-protocol nd
no ip http server
no ip http secure-server
ip tftp source-interface GigabitEthernet0
ip route 0.0.0.0 0.0.0.0 172.19.183.1
ip route 172.19.183.172 255.255.255.255 ucse1/0/0
ip ssh version 2
!
!
!

control-plane
!
!
line con 0
 stopbits 1
line aux 0
 stopbits 1
line vty 0 4
 password lab
 login local
 transport input all
 transport output all
!
!
end
```

Installing Cisco Enterprise NFVIS on a Cisco ENCS 5100 and 5400



Note Software or hardware RAID controller setup is not supported with Cisco ENCS in Cisco Enterprise NFVIS Release 3.5.1.

Before you begin

- Make sure that you have configured the IP address to access CIMC.
- The recommended CIMC version is 3.2(7) or newer.

For details on how to configure an IP address for CIMC, see the Set up CIMC for UCS C-Series Server and Getting Started Guide for Cisco UCS E-Series Servers and the Cisco UCS E-Series Network Compute Engine on cisco.com.

Procedure

Step 1 Log into CIMC.

Step 2 Launch KVM Console:

From theCIMC Homepage, click Launch KVM on the upper right corner.You can choose from Java based KVM or HTML based KVM. It is recommnded to use HTML based KVM. Ensure that the pop-up blocker is disabled. KVM Console will open in a separate window.

Step 3 Map virtual media:

a) From KVM Console, click on **Virtual Media** and then click on **Activate Virtual Devices**.
b) Click on **Virtual Media** again and then click on **Map CD/DVD**, browse, and select the Cisco Enterprise NFVIS ISO image, click **Open** and Map Drive to mount the image.
c) Click on **Virtual Media** again to ensure the NFVIS ISO image is now mapped to CD/DVD.

Step 4 Configure Boot Order:

a) From the **CIMC Compute**, select **BIOS**.
b) Select **Configure Boot Order** and click on **Configure Boot Order**.

   The **Configure Boot Order** dialog box appears.

c) From the **CD/DVD** area, select **Cisco vKVM-Mapped vDVD**, and then click **Add**.
d) From **HDD**, select **RAID Adapter**, and then click **Add**.
e) Set the boot order sequence using the **Up** and **Down** options. The Cisco vKVM-Mapped vDVD boot order option must be the first choice. Click **Save Changes** to complete the boot order setup.

Step 5 Power cycle server to start the installation:

From CIMC Home page, click on **Host Power** on the upper right corner. Reboot the server by selecting the **Power Off** option. After the server is down, select the **Power On** option.

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

Step 6 Firmware Auto-Upgrade.

Starting from NFVIS 3.8.x release, firmware auto-upgrade is supported. After the NFVIS installation is complete, BIOS and/or CIMC is upgraded to the corresponding versions automatically. CIMC and NFVIS is rebooted multiple times. The firmware upgrade might take 30 minutes to one hour to complete. Do not use the system during the firmware upgrade.

Step 7 Initial login: After the installation is complete, the system automatically reboots from the hard drive. Log into the system when the command prompt **nfvis login** is displayed after the reboot.

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 attempt. 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 8 Verification: You can verify the installation using the System API, CLI, or by viewing the system information from the Cisco Enterprise NFV portal.

Install Cisco Enterprise NFVIS on Cisco ENCS 5104

Procedure

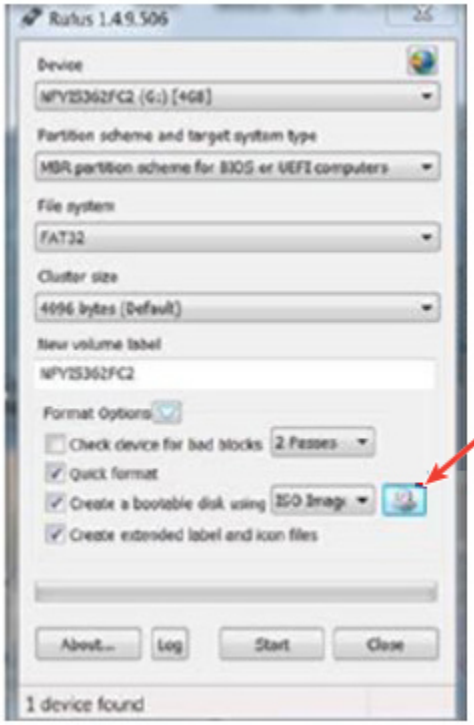
Step 1 Create bootable usb with NFVIS image.

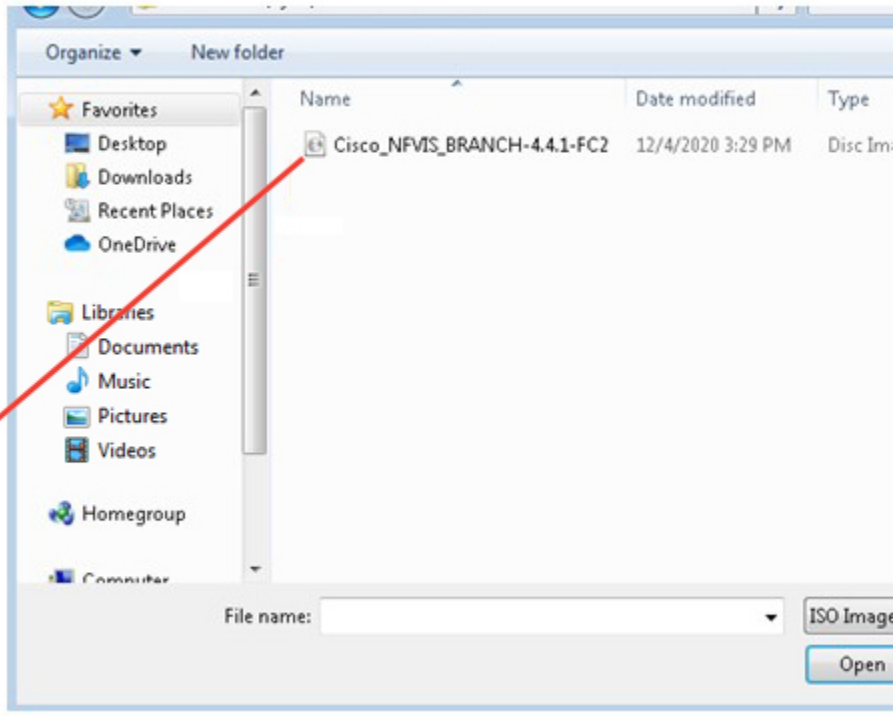
In this example, we used rufus utility in Windows environment. Rufus utility can be downloaded https://rufus.akeo.ie/. For this example, following parameters were used to burn bootable NFVIS USB device:

- Device: USB stick
- Partition scheme: MBR
- Filesystem: FAT32
- Cluster size: use default
- Volume label: use default
- Quick format: checked
- Create bootable: select "ISO Image" and click next icon then choose NFVIS image.
- Create extended label: checked

Press **Start** and wait for completion.

Eject USB thumb drive





Step 2 Insert USB device in one of USB slot in ENCS5104.

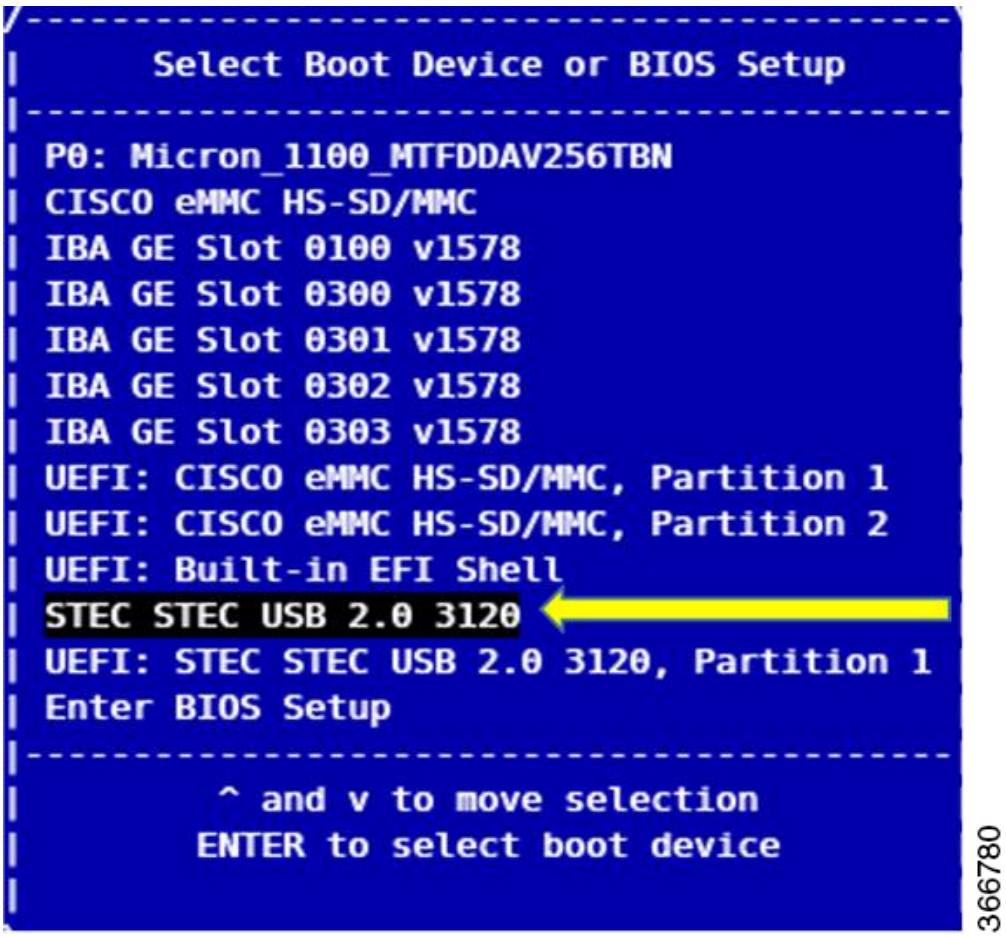
Step 3 Power on system.

Step 4 During system boot up, press F6 key.

```
Press <DEL> or <F2> to enter setup, <F6> Boot Menu, <F12> Network Boot in 5 seconds or press any key
 to continue.
```

Step 5 Once you press F6, you will see the following screenshot to select which device you want to boot from. Select your USB device.

In the following screenshot example, there is STEC USB being used. That display will vary depending on your usb device vendor. Use the arrow key to select that device.



Step 6 Wait until installation is completed. System will be rebooted once installation is done.

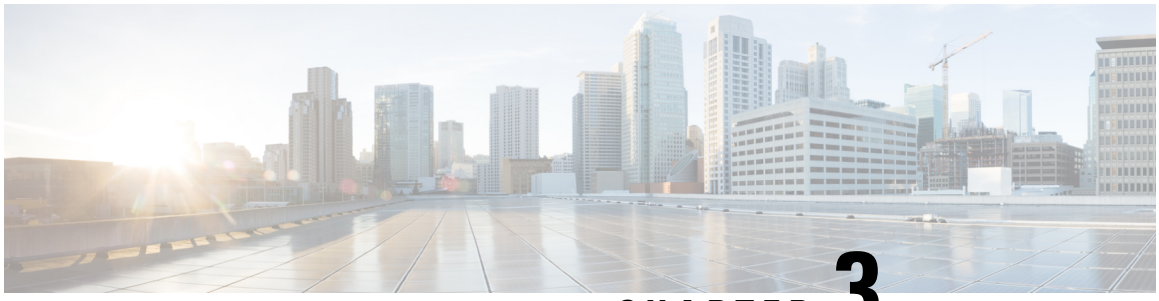
Step 7 Log into the system with username **admin** and **Admin123#** as a default password

Step 8 You will be prompted and asked to change password at the first login. You must set a strong password per the on-screen instruction to proceed.

Step 9 You can verify the installation status using the System API or command line interface per the NFVIS user guide.

What to do next

You can verify the default configuration, and set up initial IP configuration to launch the Cisco Enterprise NFV portal. For details, see Performing Initial System Configuration.



CHAPTER 3

Setting Up System Configuration

- Default System Configuration on the Cisco ENCS, on page 17
- Default System Configuration on the Cisco UCS C220 M4 Server and Cisco CSP 2100, on page 19
- Default System Configuration on the Cisco UCS E-Series Servers , on page 20
- Accessing NFVIS, on page 20
- User Roles and Authentication, on page 27
- Configuring the IP Receive ACL, on page 31
- Configuring Your Banner and Message of the Day, on page 32
- Setting the System Time Manually or With NTP, on page 33
- Enable or Disable Access to NFVIS Portal, on page 34
- Configuring System Logs, on page 35
- Network File System Support, on page 36
- Secure Boot of host, on page 37
- CIMC Access Control, on page 38
- DPDK Support on NFVIS, on page 40
- Import and Export NFVIS VM, on page 43
- Backup and Restore NFVIS and VM Configurations, on page 46
- Dynamic SR-IOV, on page 55

Default System Configuration on the Cisco ENCS

The diagram below illustrates the default network configuration of Cisco Enterprise NFVIS with the Cisco ENCS.

Figure 1: Default Network Configuration of Cisco Enterprise NFVIS with the Cisco ENCS 5400

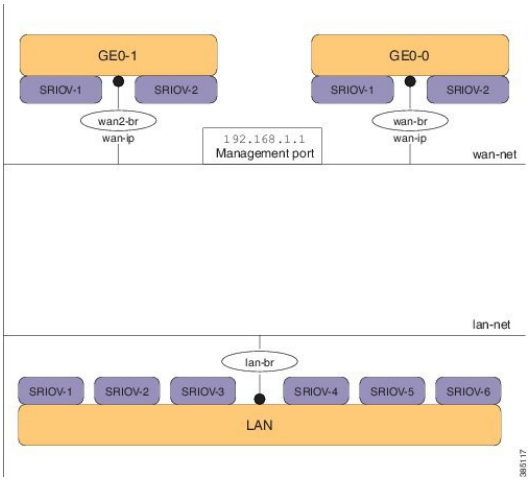
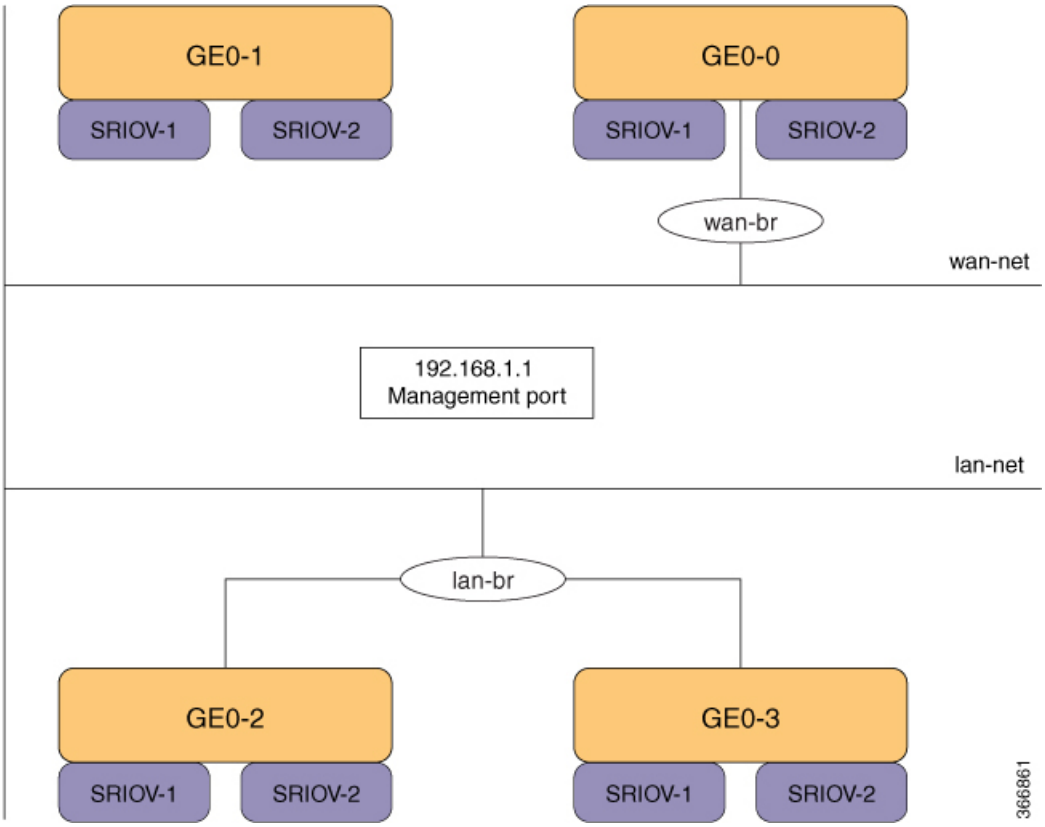


Figure 2: Default Network Configuration of Cisco Enterprise NFVIS with the Cisco ENCS 5100



- LAN ports—Eight physical Gigabit Ethernet ports for inbound and outbound traffic.
- WAN port—You can use one of the dual media Ethernet ports (wan-br and wan2-br) for DHCP connection.
- Bridges—They form a Layer 2 domain between virtual network interface controllers (vNICs) of VMs. A vNIC is used by a virtual machine to provide virtual network interfaces by defining a range of MAC

addresses. The default management IP address (192.168.1.1) for the NFVIS host is configured on the management port. Multiple VMs can use the same LAN port for local connectivity.

- Network—It is a segment Layer 2 bridge domain where only the specific VLAN traffic is allowed.
- Reserved VLANs in the LAN network on the ENCS 5400 platform—The VLAN range 2350-2449 is reserved for internal use and should not be used on the external switch ports and for virtual machines in the LAN ports". Note that this limitation doesn't apply to the WAN ports.
- Internal 192.168.10.00/24 and 192.168.50.0/24 networks—The IP subnet 192.168.10.0/24 and 192.168.50.0/24 are used for the ENCS-5400 internal networks. A user should not use this IP subnet on the NFVIS management network. In the future NFVIS releases, this internal subnet will be isolated so that users can use this for NFVIS management.

Note The following networks and bridges are automatically configured. You can configure more as required.

- A LAN network (lan-net) and a LAN bridge (lan-br)
- A WAN network (wan-net) and a WAN bridge (wan-br)

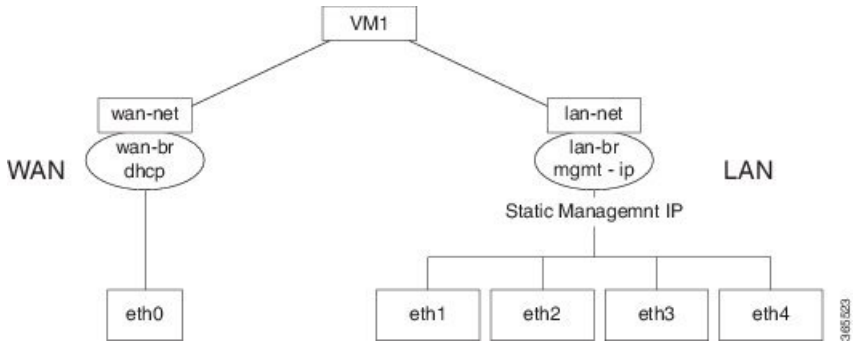
wan2-net and wan2-br are the default configurations for ENCS 5400 and ENCS 5100.

The default networks and bridges cannot be deleted.

Default System Configuration on the Cisco UCS C220 M4 Server and Cisco CSP 2100

Configuring the networks in Cisco Enterprise NFVIS allows inbound and outbound traffic and VMs to be service chained. The following diagram illustrates the default network configuration:

Figure 3: Default Network Configuration with Cisco UCS C220 M4 and Cisco CSP 2100



The following networks and bridges are created by default, and cannot be deleted. You can configure more as required.

- A LAN network (lan-net) and a LAN bridge (lan-br)—The default static management IP address (192.168.1.1) for the NFVIS host is configured on the LAN bridge. One of the ports for inbound and

outbound traffic are associated with the LAN bridge. Any LAN port can be used to access the default static IP address. By default, the hostname is set to "nfvis".

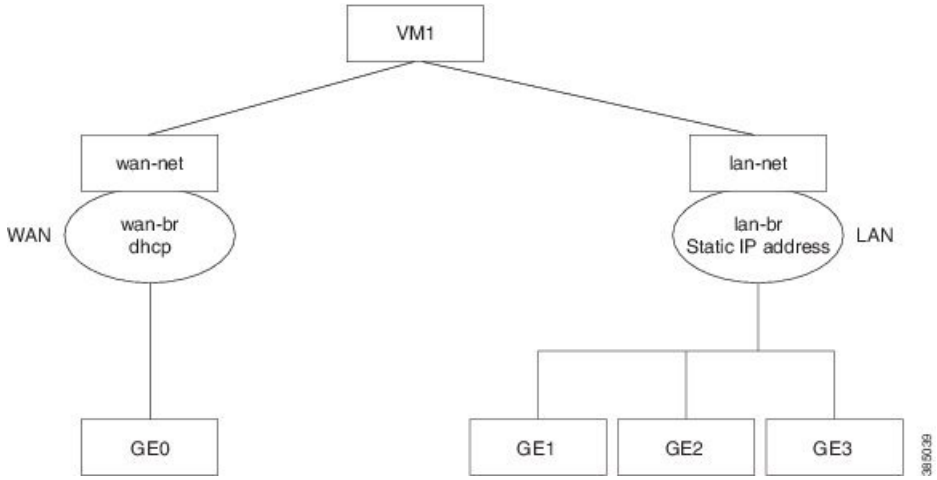
- A WAN network (wan-net) and a WAN bridge (wan-br)—This is created with the "eth0" port, and is configured to enable the DHCP connection.

By default, the first port on the device is associated with the WAN bridge. One of the other ports on the device are associated with the LAN bridge.

For more details about the initial setup, see the Installing the Server chapter in the *Cisco UCS C220 M4 Server Installation and Service Guide* or *Cisco Cloud Services Platform 2100 Hardware Installation Guide*.

Default System Configuration on the Cisco UCS E-Series Servers

Figure 4: Default Network Configuration with a Cisco UCS E-Series Server



The following networks and bridges are created by default, and cannot be deleted. You can configure more as required.

- A LAN network (lan-net) and a LAN bridge (lan-br)—The default static management IP address (192.168.1.1) for the NFVIS host is configured on the LAN bridge. All other ports for inbound and outbound traffic are associated with the LAN bridge. By default, the hostname is set to "nfvis".
- A WAN network (wan-net) and a WAN bridge (wan-br)— The physical WAN ports are on the Cisco ISR module. They are not externally available on the Cisco UCS E server. The WAN traffic comes from the ISR WAN ports, and goes through the backplane to the Cisco UCS-E server. The backplane has one internal WAN interface (GE0) to establish connection with the Cisco UCS-E server. By default, the "GE0" interface is enabled for the DHCP connection.

For more details on the initial setup, see the Getting Started Guide for Cisco UCS E-Series Servers and the Cisco UCS E-Series Network Compute Engine.

Accessing NFVIS

For initial login, use **`admin`** as the default user name, and **`Admin123#`** as the default password. Immediately after the initial login, the system prompts you to change the default password. You must set a strong password

as per the on-screen instructions to proceed with the application. All other operations are blocked until default password is changed. API returns 401 unauthorized error if the default password is not reset.

If wan-br or wan2-br have not obtained IP addresses through DHCP, the zero touch deployment is terminated. To manually apply the IP configurations answer 'y' and the system proceeds with DHCP assignment on wan-br until the configurations are changed. For DHCP assignment to continue to request IP address for PnP flow on both WAN interfaces answer 'n'.

You must adhere to the following rules to create a strong password:

- Must contain at least one upper case and one lower case letter.
- Must contain at least one number and one special character (# _ - * ?).
- Must contain seven characters or greater. Length should be between 7 and 128 characters.

You can change the default password in three ways:

- Using the Cisco Enterprise NFVIS portal.
- Using the CLI (When you first log into Cisco Enterprise NFVIS through SSH, the system will prompt you to change the password).
- Using PnP (for details, see the Cisco Network Plug-n-Play Support , on page 59).
- Using console (After the initial login using the default password, you are prompted to change the default password).

```
NFVIS Version: 3.10.0-9

Copyright (c) 2015-2018 by Cisco Systems, Inc.
Cisco, Cisco Systems, and Cisco Systems logo are registered trademarks of Cisco
Systems, Inc. and/or its affiliates in the U.S. and certain other countries.

The copyrights to certain works contained in this software are owned by other
third parties and used and distributed under third party license agreements.
Certain components of this software are licensed under the GNU GPL 2.0, GPL 3.0,
LGPL 2.1, LGPL 3.0 and AGPL 3.0.

nfvis login: console (automatic login)

login:
login:
login:
login:
login: admin

Cisco Network Function Virtualization Infrastructure Software (NFVIS)

NFVIS Version: 3.10.0-9

Copyright (c) 2015-2018 by Cisco Systems, Inc.
Cisco, Cisco Systems, and Cisco Systems logo are registered trademarks of Cisco
Systems, Inc. and/or its affiliates in the U.S. and certain other countries.

The copyrights to certain works contained in this software are owned by other
third parties and used and distributed under third party license agreements.
Certain components of this software are licensed under the GNU GPL 2.0, GPL 3.0,
LGPL 2.1, LGPL 3.0 and AGPL 3.0.

admin@localhost's password:
```

```
admin connected from ::1 using ssh on nfvis
nfvis# show version


NFVIS Version: 3.12.3

Copyright (c) 2015-2020 by Cisco Systems, Inc.
Cisco, Cisco Systems, and Cisco Systems logo are registered trademarks of Cisco
Systems, Inc. and/or its affiliates in the U.S. and certain other countries.

The copyrights to certain works contained in this software are owned by other
third parties and used and distributed under third party license agreements.
Certain components of this software are licensed under the GNU GPL 2.0, GPL 3.0,
LGPL 2.1, LGPL 3.0 and AGPL 3.0.

login: admin
NFVIS service is OK
Warning: Permanently added 'localhost' (RSA) to the list of known hosts.
admin@localhost's password:


Cisco Network Function Virtualization Infrastructure Software (NFVIS)

NFVIS Version: 3.12.3-RC8

Copyright (c) 2015-2020 by Cisco Systems, Inc.
Cisco, Cisco Systems, and Cisco Systems logo are registered trademarks of Cisco
Systems, Inc. and/or its affiliates in the U.S. and certain other countries.

The copyrights to certain works contained in this software are owned by other
third parties and used and distributed under third party license agreements.
Certain components of this software are licensed under the GNU GPL 2.0, GPL 3.0,
LGPL 2.1, LGPL 3.0 and AGPL 3.0.


admin connected from ::1 using ssh on nfvis
admin logged with default credentials
Setting admin password will disable zero touch deployment behaviors.
Do you wish to proceed? [y or n]y
Please provide a password which satisfies the following criteria:
        1.At least one lowercase character
        2.At least one uppercase character
        3.At least one number
        4.At least one special character from # _ - * ?
        5.Length should be between 7 and 128 characters
Please reset the password :
Please reenter the password :


 Resetting admin password


New admin password is set

nfvis#
System message at 2020-01-08 03:10:10...
Commit performed by system via system using system.
nfvis#
```

Note To commit the target configuration to the active (running) configuration, use the **commit** command in any configuration mode. Changes made during a configuration session are inactive until the **commit** command is entered. By default, the commit operation is pseudo-atomic, meaning that all changes must succeed for the entire commit operation to succeed.

Connect to the System

Using IPv4

The three interfaces that connect the user to the system are the WAN and WAN2 interfaces and the management interface. By default, the WAN interface has DHCP configuration and the management interface is configured with a static IP address of 192.168.1.1. If the system has a DHCP server connected to the WAN interface, the WAN interface is assigned an IP address from this server. You can use this IP address to connect to the system.

You can connect to the server locally (with an Ethernet cable) using the static management IP address. However, to be able to use a static IP address to remotely connect to a server, the default gateway needs to be configured first.

You can connect to the system in the following ways:

- Using the local portal—After the initial login, you are prompted to change the default password.
- Using the KVM console—After the initial login using the default password, you are prompted to change the default password.
- Using PnP—After the initial provisioning through PnP, the configuration file pushed by the PNP server must include the new password for the default user (admin).

Using IPv6

IPv6 can be configured in static, DHCP stateful and Stateless Autoconfiguration (SLAAC) mode. By default, DHCP IPv6 stateful is configured on the WAN interface. If DHCP stateful is not enabled on the network, the router advertisement (RA) flag decides which state the network stays in. If the RA shows the Managed (M) flag, then the network stays in DHCP mode, even if there is no DHCP server in the network. If the RA shows the Other (O) flag, then the network switches from DHCP server to SLAAC mode.

SLAAC provides IPv6 address and default gateway. Stateless DHCP is enabled in the SLAAC mode. If the server has DNS and domain configured, then SLAAC also provides those values via stateless DHCP.

Perform Static Configuration without DHCP

Note Starting from NFVIS 3.10.1 release, for ENCS 5400 and ENCS 5100, wan2-br obtains an IP address from DHCP. To configure default gateway, first use the **no bridges bridge wan2-br dhcp** command.

If you want to disable DHCP and use static configuration, you need to perform the initial configuration by setting the WAN IP address and/or management IP address, and the default gateway. You can also configure a static IP on a created bridge.

To perform initial configuration on the system without using DHCP:

```
configure terminal
```

```
system settings mgmt ip address 192.168.1.2 255.255.255.0
bridges bridge wan-br ip address 209.165.201.22 255.255.255.0
system settings default-gw 209.165.201.1
commit
```

Note When an interface is configured with a static IP address, DHCP is automatically disabled on that interface.

Now you can either use the management IP or WAN IP to access the portal.

To configure static IPv6 on the WAN interface:

```
configure terminal
system settings mgmt ipv6 address 2001:DB8:1:1::72/64
bridges bridge wan-br ipv6 address 2001:DB8:1:1::75/64
system settings default-gw-ipv6 2001:DB8:1:1::76
commit
```

Note When an interface is configured with a static IPv6 address, DHCP IPv6 is automatically disabled on that interface. There are three options for IPv6 - static, DHCP and SLAAC, out of which only one can be enabled at a time.

Secure overlay is not supported when WAN interface is configured with IPv6.

To configure DHCP on the WAN interface:

```
configure terminal
no system settings default-gw
system settings wan dhcp
commit
exit
hostaction wan-dhcp-renew
```

Note Starting from NFVIS 3.10.1, you can configure DHCP IPv6 on any bridge. You can only have one DHCP IPv6 bridge or management interface active at a time. You cannot have DHCP IPv6 and default IPv6 gateway or SLAAC IPv6 configured at the same time.

To configure DHCP IPv6 on the WAN interface:

```
configure terminal
no system settings default-gw-ipv6
system settings wan dhcp-ipv6
commit
exit
hostaction wan-dhcp-renew
```

Verify Initial Configuration

Use the **show system settings-native** command to verify initial configuration. Use **show bridge-settings** and **show bridge-settings** *bridge_name* commands to verify the configuration for any bridge on the system.

Here is an extract from the output of the **show system settings-native** command when both WAN and management interfaces have a static configuration:

```
system settings-native mgmt ip-info interface lan-br
system settings-native mgmt ip-info ipv4_address 192.168.1.2
system settings-native mgmt ip-info netmask 255.255.255.0
!
!
!
system settings-native mgmt dhcp disabled
system settings-native wan ip-info interface wan-br
system settings-native wan ip-info ipv4_address 209.165.201.22
system settings-native wan ip-info netmask 255.255.255.0
!
!
!
system settings-native wan dhcp disabled
!
!
system settings-native gateway ipv4_address 209.165.201.1
system settings-native gateway interface wan-br
```

Here is an extract from the output of the **show system settings-native** command when the management interface has a DHCP configuration and the WAN interface has a static configuration:

```
system settings-native mgmt ip-info interface MGMT
system settings-native mgmt ip-info ipv4_address 192.168.1.2
system settings-native mgmt ip-info netmask 255.255.255.0
!
!
!
system settings-native mgmt dhcp enabled
system settings-native wan ip-info interface wan-br
system settings-native wan ip-info ipv4_address 209.165.201.22
system settings-native wan ip-info netmask 255.255.255.0
!
!
!
system settings-native wan dhcp disabled
```

Here is an extract from the output of the **show system settings-native** command when the WAN interface has a DHCP configuration and the management interface has a static configuration:

```
system settings-native mgmt ip-info interface lan-br
system settings-native mgmt ip-info ipv4_address 209.165.201.2
system settings-native mgmt ip-info netmask 255.255.255.0
!
!
!
system settings-native mgmt dhcp disabled
system settings-native wan ip-info interface wan-br
system settings-native wan ip-info ipv4_address 209.165.201.22
system settings-native wan ip-info netmask 255.255.255.0
!
!
!
system settings-native wan dhcp enabled
```

Configuring VLAN for NFVIS Management Traffic

A VLAN creates independent logical networks within a physical network. VLAN tagging is the practice of inserting a VLAN ID into a packet header in order to identify which VLAN the packet belongs to.

You can configure a VLAN tag on the WAN bridge (wan-br) interface to isolate Cisco Enterprise NFVIS management traffic from VM traffic. You can also configure VLAN on any bridge on the system (wan2-br for ENCS5400 or ENCS 5100, and user-br for all systems)

By default, WAN bridges and LAN bridges are in trunk mode and allows all VLANs. When you configure native VLAN, you must also configure all the allowed VLANs at the same time. The native VLAN becomes the only allowed VLAN if you do not configure all the VLANs. If you want a network that allows only one VLAN, then create another network on top of wan-net and lan-net and make it access network.

Note You cannot have the same VLAN configured for the NFVIS management and VM traffc.

For more details on the VLAN configuration, see the Understanding and Configuring VLANs module in the Catalyst 4500 Series Switch Cisco IOS Software Configuration Guide.

Configuring System Routes

In addition to the default routes in the system, you can configure additional system routes. This configuration is specifically useful when certain destinations are not reachable through the default routes.

While you can create a route just by providing the destination and prefix length, a valid route requires that you specify either a device or a gateway or both.

To configure additional system routes:

```
configure terminal
system routes route  209.165.201.1 dev lan-br
commit
```

Verifying the System Routes Configuration

To verify the system routes configuration, use the **show system routes** command as shown below:

```
nfvis# show system routes
DESTINATION  PREFIXLEN  STATUS
----------------------------------
209.165.201.1  12 -
209.165.201.2  12 -
209.165.201.3  24 -
```

System Routes APIs and Commands

System Routes APIs	System Routes Commands
• /api/config/system/routes • /api/config/system/routes/route/<host destination,netmask>	• system routes route • show system routes

User Roles and Authentication

Role based access enables the administrator to manage different levles of access to the system's compute, storage, database, and application services. It uses the access control concepts such as users, groups, and rules, which you can apply to individual API calls. You can also keep a log of all user activities.

Table 1: Supported User Roles and Privileges

User Role	Privilege
Administrators	Owns everything, can perform all tasks including changing of user roles, but cannot delete basic infrastructure. Admin's role cannot be changed; it is always "administrators".
Operators	Start and stop a VM, and view all information
Auditors	Read-only permission

Rules for User Passwords

The user passwords must meet the following requirements:

- Must have at least seven characters length or the minimum required length configured by the admin user.
- Must not have more than 128 characters.
- Must contain a digit.
- Must contain one of the following special characters: hash (#), underscore (_), hyphen (-), asterisk (*), and question mark (?).
- Must contain an uppercase character and a lowercase character.
- Must not be same as last five passwords.

Creating Users and Assigning Roles

The administrator can create users and define user roles as required. You can assign a user to a particular user group. For example, the user "test1" can be added to the user group "administrators".

Note All user groups are created by the system. You cannot create or modify a user group.

To create a user:

```
configure terminal
rbac authentication users create-user name test1 password Test1_pass role administrators
commit
```

To delete a user:

```
configure terminal
rbac authentication users delete-user name test1
commit
```

Note To change the password, use the **rbac authentication users user change-password** command in global configuartion mode. To change the user role, use the **rbac authentication users user change-role** command in global configuration mode.

User Management APIs and Commands

User Management APIs	User Management Commands
• /api/operations/rbac/authentication/users /user/<user-name>/change-password • /api/operations/rbac/authentication/users/user /oper/change-role • /api/config/rbac/authentication/users/user?deep	• rbac authentication users • rbac authentication users user change-password • rbac authentication users user change-role

Configuring Minimum Length for Passwords

The admin user can configure the minimum length required for passwords of all users. The minimum length must be between 7 to 128 characters. By default, the minimum length required for passwords is set to 7 characters.

```
configure terminal
rbac authentication min-pwd-length 10
commit
```

Minimum Password Length APIs and Commands

APIs	Commands
/api/config/rbac/authentication/	rbac authentication min-pwd-length

Configuring Password Lifetime

The admin user can configure minimum and maximum lifetime values for passwords of all users and enforce a rule to check these values. The default minimum lifetime value is set to 1 day and the default maximum lifetime value is set to 60 days.

When a minimum lifetime value is configured, the user cannot change the password until the specified number of days have passed. Similarly, when a maximum lifetime value is configured, a user must change the password before the specified number of days pass. If a user does not change the password and the specified number of days have passed, a notification is sent to the user.



Note The minimum and maximum lifetime values and the rule to check for these values are not applied to the admin user.

```
configure terminal
rbac authentication password-lifetime enforce true min-days 2 max-days 30
commit
```

Password Lifetime APIs and Commands

APIs	Commands
/api/config/rbac/authentication/password-lifetime/	rbac authentication password-lifetime

Deactivating Inactive User Accounts

The admin user can configure the number of days after which an unused user account is marked as inactive and enforce a rule to check the configured inactivity period. When marked as inactive, the user cannot login to the system. To allow the user to login to the system, the admin user can activate the user account by using the **rbac authentication users user** *username* **activate** command.



Note The inactivity period and the rule to check the inactivity period are not applied to the admin user.

```
configure terminal
rbac authentication account-inactivity enforce true inactivity-days 2
commit
```

Deactivate Inactive User Accounts APIs and Commands

APIs	Commands
/api/config/rbac/authentication/account-inactivity/	rbac authentication account-inactivity

Activating an Inactive User Account

The admin user can activate the account of an inactive user.

```
configure terminal
rbac authentication users user guest_user activate
commit
```

Activate Inactive User Account APIs and Commands

APIs	Commands
/api/operations/rbac/authentication/users/user/username/activate	rbac authentication users user activate

Certification

Generate Sign-Request

```
nfvis(config)# system certificate signing-request ?

Possible completions:

  common-name             country-code

  locality                organization

  organization-unit-name  state
```

The .csr file will be saved in /data/intdatastore/download/nfvis.csr

Use the scp command to download the file.

Install CA Sign Certificate

After CA sign in, the user needs to use the scp command to upload the file into nfvis.

```
nfvis(config)# system certificate install-cert path file:///<full path of the file>
```

The path needs to start with "file://"

Switch Certificate

```
nfvis(config)# system certificate use-cert cert-type ca-signed
```

nginx process restarts after the switch.

The users cannot access the log files. The log files are added to all the user actions and the user can download and view some of the logs from portal. A notification is generated when the log files reach 75% capacity.

Secure Copy Command

The secure copy (**scp**) command allows only the admin user to secure copy a file from the Cisco NFVIS to an external system or from an external system to Cisco NFVIS. The **scp** command is:

scp *source destination*

Note For detailed information about how to use the **scp** command to copy to or from supported locations, see the **scp** section in Cisco Enterprise Network Function Virtualization Infrastructure Software Command Reference.

Examples

The following example copies the sample.txt file from intdatastore to an external system.

```
nfvis# scp intdatastore:sample.txt user@203.0.113.2:/Users/user/Desktop/sample.txt
```

The following example copies the test.txt file from an external system to intdatastore.

```
nfvis# scp user@203.0.113.2:/Users/user/Desktop/test.txt intdatastore:test_file.txt
```

The following example copies the test.txt file from an external system to USB.

```
nfvis# scp user@203.0.113.2:/user/Desktop/my_test.txt usb:usb1/test.txt
```

The following example copies the sample.txt file to an NFS location.

```
nfvis# scp user@203.0.113.2:/user/Desktop/sample.txt nfs:nfs_test/sample.txt
```

The following example copies the sample.txt file from an external system with IPv6 address.

```
nfvis# scp user@[2001:DB8:0:ABCD::1]:/user/Desktop/sample.txt intdatastore:sample.txt
```

The following example copies the nfvis_scp.log file to an external system.

```
nfvis# scp logs:nfvis_scp.log user@203.0.113.2:/Users/user/Desktop/copied_nfvis_scp.log
```

Configuring the IP Receive ACL

To filter out unwanted traffic, you can configure ip-receive-acl to block or allow certain traffic based on the IP address and service ports.

To configure the source network for Access Control List (ACL) access:

```
configure terminal
system settings ip-receive-acl 198.0.2.0/24
action accept priority 10
commit
```

Verifying the Trusted IP Connection

Use the **show running-config system settings ip-receive-acl** command to display the configured source network for ACL access to the management interface

```
nfvis# show running-config system settings ip-receive-acl
system settings ip-receive-acl 198.51.100.11/24
service
[ ssh https scpd]
action accept
priority 100
```

Port 22222 and Management Interface ACL

Port 22222 is used for SCP server and is closed by default on an NFVIS system. You cannot SCP a file into NFVIS from an external server. If you need to SCP file from an external server, you must first open the port.

To open port 22222:

```
config terminal
system settings ip-receive-acl address/mask_len service scpd priority 2 action accept
commit
```

The Access Control List (ACL) is identify by address. If this ACL is removed, all ACLs sharing the same address are also removed. Ensure that you configure the ACLs that share the same address once again.

Note From 3.8.1 release, only users with administrator priviledges can use the SCP command on port 22222 to upload or download only from restricted folders like /data/intdatastore/. For more information, see Host System Operations.

Caution SCP command cannot be used to copy files from one NFVIS device to another NFVIS device.

Use the **show running-config system settings ip-receive-acl** command to verify the interface configuration:

```
nfvis# show running-config system settings ip-receive-acl

system settings ip-receive-acl 10.156.0.0/16

 service  [ ssh https scpd ]

 action   accept

 priority 100

!
```

Configuring Your Banner and Message of the Day

Cisco Enterprise NFVIS supports two types of banners: system-defined and user-defined banners. You cannot edit or delete the system-defined banner, which provides copyright information about the application. Banners are displayed on the login page of the portal.

You can post messages using the Message of the Day option. The message is displayed on the portal's home page when you log into the portal.

To configure your banner and message:

```
configure terminal
banner-motd banner "This is a banner" motd "This is the message of the day"
commit
```

Note Currently, you can create banners and messages in English only. You can view the system-defined banner using the **show banner-motd** command. This command does not display the user-defined banner or message.

Banner and Message APIs and Commands

Banner and Message APIs	Banner and Message Commands
• /api/config/banner-motd • /api/operational/banner-motd	• banner-motd • show banner-motd

Setting the System Time Manually or With NTP

You can configure the Cisco Enterprise NFVIS system time manually or synchronise with an external time server using Network Time Protocol (NTP).

To set the system time manually:

```
configure terminal
system set-manual-time 2017-01-01T00:00:00
commit
```

Note NTP is automatically disabled when the time clock is set manually.

To set the system time using NTP IPv4:

```
configure terminal
system time  ntp preferred_server 209.165.201.20 backup_server 1.ntp.esl.cisco.com
commit
```

To set the system time using NTP IPv6:

```
configure terminal
system time  ntp-ipv6 2001:420:30d:201:ffff:ffff:fff4:35
commit
```

Verifying the System Time Configuration

To verify all system time configuration details, use the **show system time** command in privileged EXEC mode as shown below:

```
nfvis# show system time

system time current-time 2017-01-01T17:35:39+00:00

system time current-timezone "UTC (UTC, +0000)"

REMOTE               REFID   ST    T         WHEN       POLL       REACH         DELAY
```

```
     OFFSET        JITTER
==============================================================================

*calo-timeserver     .GPS.    1          u           4   64           1         69.423
     2749736        0.000

* sys.peer and synced, o pps.peer, # selected, + candidate,

- outlyer, . excess, x falseticker, space reject
```

If the NTP server is invalid, it will not be displayed in the table. Also, when an NTP server is queried, if a response is not received before the timeout, the NTP server is not displayed in the table.

System Time APIs and Commands

APIs	Commands
• /api/operations/system/set-manual-time • /api/config/system/time/ntp/preferred_server • /api/config/system/time/ntp/backup_server • /api/config/system/time/timezone • /api/operational/system/time?deep	• system time • show system time • system set-manual-time

Enable or Disable Access to NFVIS Portal

The Cisco Enterprise NFVIS portal access is enabled by default. You can disable the access if required.

To disable the portal access:

```
configure terminal
system portal access disabled
commit
```

Note You can enable the portal access using the **enabled** keyword with the **system portal access** configuration.

Verifying the Portal Access

Use the **show system portal status** command to verify the portal access status as shown below:

```
nfvis# show system portal status
system portal status "access disabled"
```

Portal Access APIs and Commands

Portal Access APIs	Portal Access Commands
• /api/config/system/portal • /api/operational/system/portal/status	• system portal access • show system portal status

Configuring System Logs

NFVIS generates log files for troubleshooting issues. The configuration log and the operational log are the two main system log files. The configuration log has information related to configurations and actions performed on the system such as creation of networks. The operational log has information related to system operation such as statistics collection and monitoring.

Log entries can be one of the following types:

Log Level	Purpose
DEBUG	Information, typically of interest only when diagnosing problems.
INFO	Confirmation that things are working as expected.
WARNING	An indication that something unexpected happened, or indicative of some problem in the near future (for example, ‘disk space low’). The software application is still working as expected.
ERROR	Due to a serious problem, the software application is not able to perform some function.
CRITICAL	A serious error, indicating that the program itself may not be able to continue running.

By default, the configuration log has a log-level of INFO. All logs of type INFO, WARNING, ERROR and CRITICAL are logged.

By default, the operational log has a log-level of WARNING. All logs of type WARNING, ERROR and CRITICAL are logged.

The log-level for these log files can be changed using the **system set-log** command:

```
system set-log level error logtype configuration
```

The change to the log level is not persistant across a reboot. After a reboot, the default log levels are used.

The current log files are kept in the */var/log* directory in the system:

- show log - To display the list of available log files
- show log {filename} - To display the contents of a specific log file

Log Rotation

There is a size limit for the log files, under */var/log/* directory. When the log files reach the size limit, the location of logs is rotated to another place. The space limit for the total size of all rotated log files is 2 GB. The older log files are dropped automatically on reaching the space limit. You can also execute a command to trigger the log rotation procedure. The log files are monitored periordically and if a log file gets too big, it is rotated to another place.

There is a size limit for the log files stored in the */var/log* directory. The size of the log files is monitored periodically every fifteen minutes and if a log file gets too big, it is rotated to the */data/intdatastore/logs* directory. The space limit for the total size of all the rotated log files is 2 GB. The older log files are dropped automatically on reaching the space limit. You can also execute the **logrotate** command to trigger the log rotation procedure.

```
nfvis# logrotate
```

Verifying the System Log Configuration

To verify the system log configuration, use the **show system logging-level** command as shown below:

```
nfvis# show system logging-level
system logging-level configuration error
system logging-level operational warning
```

System Log APIs and Commands

System Log APIs	System Log Commands
• /api/operations/system/set-log • /api/operational/system/logging-level	• system set-log logtype [all/configuration/operational] level [critical/debug/error/info/warning] • show system logging-level

Network File System Support

Network File System (NFS) is an application where you can view, store, and update the files on a remote device. NFS allows you to mount all or a part of a file system on a server. NFS uses Remote Procedure Calls (RPC) to route requests between the users and servers.

Mount and Unmount NFS

The following example shows how to mount NFS:

```
configure terminal
system storage nfs_storage
nfs
100
10.29.173.131
/export/vm/amol
commit
```

To unmount NFS use the **no system storage nfs_storage** command.

Image Registration on NFS

Images in tar.gz, ISO and qcow2 formats, remote images and images on mounted NFS can be registered on NFS.

To register tar.gz images on NFS:

```
configure terminal
vm_lifecycle images image myas10 src file:///data/mount/nfs_storage/repository/asav961.tar.gz
properties property placement value nfs_storage
commit
```

Similar configuration can be used for the various images formats.

To unregister an image from NFS use **no vm_lifecycle images** command.

Deploy VM on NFS

To deploy a VM on NFS, under deployment vm group, use the **placement type zone_host host nfs_storage** command.

Secure Boot of host



Note This feature is available only for NFVIS 3.9.1 release fresh install and supported only on ENCS 5400. Upgrade BIOS to version 2.6 for this feature.

The secure boot feature prevents malicious software applications and unauthorized operating systems from loading into the system during the system start up process. If secure boot feature is enabled, only the authorized software applications boots up from the device. Each device has keys that allow software with the correct signature to boot up on the device.

This feature ensures that the software applications that boot up on the device are certified by Cisco. The NFVIS 3.9.1 image is signed with Cisco key. If secure boot is enabled the signature is verified during the device boot up. If the verification fails, the image does not boot up.

Secure boot is disabled by default and to enable it you must change firmware configurations from CIMC. Secure boot needs to boot from a separate UEFI partition.

To enable secure boot:

1. Access CIMC and use **show bios detail** command to view the BIOS version.

```
ENCS# scope bios
ENCS/bios # show detail
BIOS:
    BIOS Version: " ENCS54_2.6 (Build Date: 07/12/2018)"
    Boot Order: EFI
    FW Update/Recovery Status: Done, OK
    Active BIOS on next reboot: main
    UEFI Secure Boot: disabled
ENCS/bios #
```

2. Enable secure boot of host.

```
ENCS/bios # set secure-boot enable
Setting Value : enable
Commit Pending.
ENCS/bios *# commit
ENCS/bios # show detail
BIOS:
    BIOS Version: "ENCS54_2.6 (Build Date: 07/12/2018)"
    Boot Order: EFI
    FW Update/Recovery Status: None, OK
    Active BIOS on next reboot: main
    UEFI Secure Boot: enabled
ENCS/bios #
```

Legacy boot, UEFI boot and UEFI secure boot are the three boot modes. Secure boot can only be used on a disk that has UEFI partition.

You can configure boot order from CIMC command or portal or from BIOS setup menu. You can only configure legacy boot order with CIMC . By default, **BootOrderRules** are set to **Strict**, so the boot order follows the CIMC configuration. Since CIMC cannot be used to configure UEFI boot order, to enable secure boot change the **BootOrderRules** setting to **Loose**.

If **BootOrderRules** is set to **Loose**, the boot order follows the BIOS setup menu. When an operating system is installed in secure boot mode, the new UEFI boot option for the OS automatically appears at the top of the BIOS menu boot order list, to boot the installed operating system.

To set **BootOrderRules** to **Loose**:

```
ENCS/bios # scope advanced
ENCS/bios/advanced # set BootOrderRules Loose
ENCS/bios/advanced *# commit
Changes to BIOS set-up parameters will require a reboot.
Do you want to reboot the system?[y|N]y
```

CIMC Access Control

NFVIS administrators have authoritative control control over ENCS 5400 devices. This includes capability to change the IP address used to reach the CIMC and modifying the CIMC and BIOS passwords.

CIMC Access using NFVIS



Note CIMC access using NFVIS is supported only on ENCS 5400.

When CIMC access is enabled on NFVIS, ISRv can gain access to the host CIMC and internal switch management console. You must have authorization from Cisco Interactive Debug (CID) to access both consoles.

To access CIMC using NFVIS WAN or management interface IP address, use the **system settings cimc-access enable** command. Once you configure CIMC access on NFVIS, the stand alone CIMC access using CIMC IP address is disabled and you will be able to access CIMC using NFVIS management interface IP address. The configurations remain on the device even after the device reboot.

When the CIMC access is configured, it enables a few ports to access services like SSH, SNMP, HTTP and HTTPs into the CIMC.

The following port numbers are being used for forwarding services to CIMC:

- 20226 for SNMP
- 20227 for SSH
- 20228 for HTTP
- 20229 for HTTPS

If you are unable to access CIMC using NFVIS, check the show log nfvis_config.log file.

Use **system settings cimc-access disable** to disable this feature.

BIOS-CIMC Update

For releases 3.8.1 and later, if the BIOS or CIMC versions on Cisco ENCS 5400 routers are lower than the image version in the NFVIS ISO or upgrade package, the BIOS and CIMC versions on the routers are automatically upgraded to the version of the bundled image during NFVIS upgrade or installation. The CPU microcode is also upgraded as part of this upgrade or installation. Note that the upgrade process takes longer than in previous releases and the process cannot be stopped midway.

For Cisco ENCS 5100 routers, BIOS is automatically upgraded to the new version, but the server needs to be rebooted manually for the upgrade to show.

NFVIS Password Recovery

1. Load the NFVIS ISO image, using the CIMC KVM console.
2. Select Troubleshooting from the Boot Selection menu.
3. Select Rescue a NFVIS Password.
4. Select Continue.
5. Press Return to get a shell.
6. Run the **chroot /mnt/sysimage** command.
7. Run the **./nfvis_password_reset** command to reset the password to admin.
8. Confirm the change in password and enter Exit twice.

   Disconnect the NFVIS ISO image in the CIMC KVM console and reboot NFVIS.
9. Login to NFVIS with the default credentials admin/Admin123#.

   After login to NFVIS, enter a new password at prompt.
10. Connect to NFVIS with the new password.



Note You can update and recover NFVIS 3.8.1 and older passwords using NFVIS 3.9.1.

DPDK Support on NFVIS

The Data Plane Development Kit (DPDK) is a set of data plane libraries and network interface controller drivers for fast packet processing.

DPDK Support for NFVIS 3.10.x

DPDK support is enabled only on ENCS 5400 from NFVIS 3.10.1 release. To enable DPDK use **system settings dpdk enable** command. Once DPDK is enabled it cannot be disabled. You can use **factory-default-reset all-except-images-connectivity** to disable DPDK.

To enable DPDK support:

```
configure terminal
system settings dpdk enable
commit
```

DPDK can be enabled if:

- No VMs are deployed.
- There are no other bridges created other than the default bridge which is wan-br, wan2-br or lan-br.
- The default bridges are not modified.

DPDK mode is enabled on a bridge, if the bridge is created as part of a network or bridge api without any NIC ports. NIC ports can also be added later to the bridge, if no VMs are deployed on the network associated to the bridge. If a NIC port is added to the bridge, the bridge will switch to non-dpdk mode. Once a bridge enters non-dpdk mode, it will not switch back to DPDK mode again. NFVIS supports DPDK for the interface with virtio driver only.

Note NFVIS 3.10.x release does not support **tcpdump packetcapture** command on DPDK enabled bridge.

If DPDK is enabled, all VMs deployed will have DPDK and HugePage support. The default hugepage size is 2MB. After DPDK is enabled the system reserves 512 hugepages for Openvswitch operations. Hugepages for VM are allocated dynamically. If the system is not able to allocate HugePages for a newly deployed VM, the VM will boot up in error state. Memory Fragmentation is the main reason why HugePage allocation fails. In this case a reboot can help solve the issue.

Note DPDK support is only enabled on the bridges without NIC ports.

For a system without Hyper-threading one additional core is reserved by the system and for a system with Hyper-threading two additional logical cores are reserved by the system.

NFVIS does not support changing Hyper-thread option such as disabling after DPDK is enabled with Hyper-thread. The system can be unstable if you change Hyper-thread setting after DPDK is enabled.

DPDK VM Migration for NFVIS 3.11.x

Enhancements to DPDK Support in NFVIS 3.12.x

Service Bridge DPDK Support for New Platforms

Starting from NFVIS 3.12.1 DPDK support is added for service bridges on UCS-C M5, UCS-E M3 and CSP 5000 series. Only the service bridge is supported and the default wan bridge and lan bridge do not support DPDK feature. DPDK is supported on service bridge with or without any pnic. The configuration for service bridge without any pnic is the same as ENCS platform on NFVIS 3.11.1 release, all service bridges will become dpdk-capable after DPDK is enabled. NFVIS system does not require reboot after DPDK is enabled.

DPDK and SR-IOV cannot co-exist on the pnic used for DPDK-capable bridge. The configuration for service bridge with pnic is as followed:

1. (Optional) If pnic ethX-Y is still attached to lan-br, remove it from lan-br.

```
configure terminal
bridges bridge lan-br
no port ethX-Y
commit
```

2. Remove SR-IOV networks associated with ethX-Y (usually ethX-Y-SRIOV-*).

```
configure terminal
no networks network ethX-Y-SRIOV-1
no networks network ethX-Y-SRIOV-2
...
commit
```

3. Disable ethX-Y SRIOV configuration.

```
configure terminal
pnic ethX-Y
no sriov
nnumvf 0
commit
```

4. Create service bridge with port ethX-Y.

```
configure terminal
bridges bridge service-br port ethX-Y
networks network service-net bridge service-br
commit
```

5. Enable system-wide DPDK setting.

```
configure terminal
system settings dpdk enable
commit
```

Note

- You can also enable DPDK before creating service bridge with port ethX-Y.
- Except ENCS 5400 platform, all default bridges (lan-br, wan-br, wan-br2) does not enter DPDK-mode after DPDK is enabled system-wide.
- If you do not remove SR-IOV configuration from pnic ethX-Y, and ethX-Y is a port of service bridge, this service bridge cannot enter DPDK-mode after DPDK is enabled system-wide. If you have to use pnic ethX-Y in DPDK bridge, after removing SR-IOV configuration from ethX-Y, delete old bridge and create new bridge and attach ethX-Y to the new bridge.

Migration for VMs with PNIC associated to DPDK bridge

If DPDK is enabled, all VMs deployed will have DPDK and HugePage support. The default hugepage size is 2MB. After DPDK is enabled, the system will reserve additional 512~2048 contiguous hugepages for DPDK operations. Though hugepages for VM are allocated dynamically, it will cause memory fragmentation and the main reason DPDK enabling process will fail. In this case, a system reboot might help to solve the issue.

If system can find enough hugepages for DPDK process, not only service bridges will enter DPDK-mode, all VMs' vnic attached to those bridges will be migrated from virtio to dpdk type (dpdkvhostuserclient):

```
nfvis# support ovs vsctl show | begin service-br
    Bridge service-br
        Port "vnic2"
            Interface "vnic2"
                type: dpdkvhostuserclient
                options: {vhost-server-path="/run/vhostfd/vnic2"}
        Port service-br
            Interface service-br
                type: internal

nfvis# support virsh dumpxml ROUTER | begin interface | until serial
    <interface type='network'>
      <mac address='52:54:00:ae:14:57'/>
      <source network='wan-net'/>
      <target dev='vnic1'/>
      <model type='virtio'/>
      <address type='pci' domain='0x0000' bus='0x00' slot='0x04' function='0x0'/>
    </interface>
    <interface type='vhostuser'>
      <mac address='52:54:00:3d:ee:1a'/>
      <source type='unix' path='/run/vhostfd/vnic2' mode='server'/>
      <target dev='vnic2'/>
      <model type='virtio'/>
      <address type='pci' domain='0x0000' bus='0x00' slot='0x05' function='0x0'/>
    </interface>
```

DPDK and Factory Reset

The three commonly-used factory reset configurations disable DPDK.

Factory-default-reset configuration	DPDK enabled	Networks kept	Deployment deleted	Registered image deleted
all	No	No	Yes	Yes

Factory-default-reset configuration	DPDK enabled	Networks kept	Deployment deleted	Registered image deleted
all-except-images	No	No	Yes	No
all-except-images -connectivity	No	Yes	Yes	No

Frequently Seen Failures

- Error when DPDK is enabled.
  - When enabling DPDK, lack of system memory results in an error. Ensure that you have 4 GB contiguous HugePages to enable DPDK. You can also reboot the system.
  - DPDK requires one physical core from each socket. A multi-sockets system (like UCS-C or CSP) can encounter deployments that occupy all cores on one socket and DPDK can fail to enable. Power-off some VM (resource-locking VMs likes ISRv/ASAv/vWAAS) and enable VM after DPDK is enabled successfully.
- If the system is not able to allocate HugePages for a newly deployed VM, the VM will boot up in error state.

Import and Export NFVIS VM

Starting from NFVIS 3.10.1 release, you can backup or export (vmExportAction) and restore or import (vmImportAction) VMs. To backup or restore the whole NFVIS system, refer Backup and Restore NFVIS and VM Configurations.

VM Export and Import Limitations

- The imported VM cannot change datastore.
- The original registered image must exist.
- The OVS network name must be identical to the one used by original deployment.
- VM export is dependant on the amount of free space available in the deployed datastore, regardless of the free space available in the destination datastore. For example, when the VM is deployed in the intdatastore (default), you should ensure that the available free space is at least twice that of the deployed VM.

To export a VM ensure that:

- Backup file must be saved to NFVIS datastore or USB.
- Provide a backup name for NFVIS to append .vmbkp extension to the backup name.

You can only create and save a VM backup to datastores. The backup file has .vmbkp extension. To verify the backup:

```
nfvis# show system file-list disk local | display xpath | include backup
```

```
/system/file-list/disk/local[si-no='84']/name tiny_backup.vmbkp
nfvis# show system file-list disk local 84
SI NO  NAME              PATH               SIZE  TYPE               DATE MODIFIED
-------------------------------------------------------------------------------------------
84     tiny_backup.vmbkp  /mnt/extdatastore1  17M   VM Backup Package  2019-01-31 19:31:32
```

To import a VM ensure that:

- The Backup file is placed under NFVIS datastores or USB.
- The registered image used by the original deployed VM is in the same datastore, with same properties.
- The exported VM does not exist on the system.
- OVS network used by the original deployment should exist.
- Restored VM is created with the same datastore with same deployment properties.
- The full path name to backup file is used (for example, /mnt/extdatastore1/backup.vmbkp, not extdatastore1:backup)

```
nfvis# vmImportAction importPath /mnt/extdatastore1/tiny_backup.vmbkp
System message at 2019-01-31 19:53:32...
Commit performed by admin via ssh using maapi.
```

Starting from NFVIS 4.1 release, an optional unique MAC UID support is added to VM import.

```
vmImportAction importPath <vm backup file with location> uniqueMacUid
```

Specifying the uniqueMACUid flag ensures that the imported VM is not deployed with the same UID and interface MAC addresses.

The following examples show export failures:

- Original deployment is not deleted

```
nfvis# vmImportAction importPath /mnt/extdatastore1/tiny_backup.vmbkp
Error: Exception from action callback: Deployment Configuration :
'SystemAdminTenantIdtiny' already exists , can not be imported/restored due to conflict!
```

- 2. OVS network used by original deployment is deleted.

```
nfvis# vmImportAction importPath /mnt/extdatastore1/tiny_backup.vmbkp
Error: Exception from action callback: Restoration Request rejected, see logs for root
 cause
```

Feature Comparison Table for Import and Export

VM backup using vmExportAction:

Features	NFVIS 4.1.1 and Earlier Releases	NFVIS 4.2.1 Release	NFVIS 4.4.1 Release
Default file location for backup vmExportAction vmName sample exportName vmbackup exportPath <datastore>:	/data/intdatastore/vmbackup.vmbkp /mnt/extdatastore1/vmbackup.vmbkp /mnt/extdatastore2/vmbackup.vmbkp	Same	Same
Default file location for backup on USB vmExportAction vmName sample exportName backup02 exportPath usb:usb1	/mnt-usb/usb1/vmbackup.bkup	Same	Same
Check disk space before backup	Not supported	Supported	Supported
VM backup format	Full backup	Diff disk backup	Diff disk backup
Backup image and flavor	Not supported	Not supported	Supported
VM live export snapshot	Not supported	Not supported	Supported
VM Export with Selective Disk	Not supported	Not supported	Supported

VM restore using vmImportAction:

Features	NFVIS 4.1.1 and Earlier Releases	NFVIS 4.2.1 Release	NFVIS 4.4.1 Release
Default file location for backup vmImportAction importPath <datastore>:vmbackup.vmbkp	/data/intdatastore/vmbackup.vmbkp /mnt/extdatastore1/vmbackup.vmbkp /mnt/extdatastore2/vmbackup.vmbkp	Same	Same
Default file location for restore on USB vmImportAction importPath usb:usb1/vmbackup.vmbkp	/mnt-usb/usb1/vmbackup.vmbkp	Same	Same
Check disk space before backup	Not supported	Supported	Supported
Restore backing images and flavors	Not supported	Not supported	Supported

Features	NFVIS 4.1.1 and Earlier Releases	NFVIS 4.2.1 Release	NFVIS 4.4.1 Release
VM Export with Selective Disk	Not supported	Not supported	Supported

Backup and Restore NFVIS and VM Configurations

Table 2: Feature History

Feature Name	Release Information	Description
Enhancements to backup and restore of configurations	NFVIS 4.2.1	New commands are introduced to view the overall status of backup and restore process. Enhancements to backup file location and factory default options are introduced. Information on how to troubleshoot failure to restore NFVIS configurations is added.

Starting from NFVIS 3.10.1 release, you can backup and restore NFVIS configurations and VMs. You can also restore a backup from one NFVIS device to another if they are running on the same version of NFVIS and have the same platform.

Note

- To backup and restore a single VM, use vmExportAction (for VM backup) and vmImportAction (for VM restore) APIs.
- Perform the following hostaction backup that avoids loss of VMs during hostaction restore due to insufficient disk space:
  1. Stop the functioning of the VMs that are associated with Cisco NFVIS.
  2. Perform individual image backups of the VMs using the **vmExportAction** command.
  3. Once the backup is successful, delete the VMs and the images from Cisco NFVIS.
  4. When you delete the VMs and the images, perform a host level backup with configurations-only option using the command **hostaction backup configuration-only file-path** *extdatastore2:sample-dir/sample*.
  5. Copy the backup files to a file server.
  6. Perform a factory reset using the **factory-default-reset** command.
  7. Paste the backup copied to a file server and restore the host level backup file using the **hostaction restore** *file-path extdatastore2:sample-dir/sample.bkup* command.
  8. When the restore fails due to disk storage issues, restore the configurations-only backup. When the restore is successful, restore the VMs and their images using the **vmImportAction** *importPath /mnt/extdatastore1/tiny_backup.vmbkp* command.

Restrictions for Backup and Restore on NFVIS

- The backup includes all deployed VMs and the registered images except uploaded files.
- VM restore using hostaction restore and vmImportAction requires original registered image to be on the system, on the same datastore. Missing registered image or image registered in a different datastore results in VM restore failure.

  For NFVIS 4.2 release, onlyVM restore using hostaction restore does not require original registered images on the system.
- For NFVIS 4.1 and earlier releases, NFVIS VM backup does not support differential disk backup and every backup is a full VM backup.
- For NFVIS 4.1 and earlier releases, in case of multiple deployments based on a single registered image, every VM backup includes the registered image disk.
- The time taken to backup a VM depends on the option you choose:
  - *configuration-only* - within 1 min.
  - *configuration-and-vms* - depends on the number of VM deployments on your system, system disk write speed, and compress the VM disks into one bundle.
- You can either backup all the VMs or none.

- The final backup is a compressed file which requires temporary disk space to create the VM backup file. If the system has only one datastore, the maximum deployment backups in a single file is around one-third to half of the datastore disk space. If the deployments occupies more disk space, use *vmExportAction* to backup an individual VM instead of relying on host backup for all VM deployments.
- NFVIS only supports backup or restore on the same release. For example, backup created in NFVIS 4.1.1 cannot be used to restore on NFVIS 4.2.1.
- Starting from NFVIS 4.5 release, secure-overlay configuration with EAP authentication is supported. However, it will be discarded if restored on a different box or on the same box after fresh-install because of encrypted password.
- Starting from NFVIS 4.5 release, single ip configuration is supported. However, it will be discarded if restored on a different box because the single ip bootstrap is tailored towards a particular box.

Feature Comparison Table for Backup and Restore

Backup using hostaction backup:

Feature	NFVIS 4.1.1 and Earlier Releases	NFVIS 4.2.1 Release
Default file location for backup	/data/intdatastore/uploads/backup.bkup /mnt/extdatastore1/uploads/backup.bkup /mnt/extdatastore2/uploads/backup.bkup	/data/intdatastore/backup.bkup /mnt/extdatastore1/backup.bkup /mnt/extdatastore2/backup.bkup
VM backup format	Full backup	Diff disk backup
Registered Image and Flavors	No	Yes
Status monitoring	No	Yes
Check disk space before backup	No	Yes

Restore using hostaction restore:

Feature	NFVIS 4.1.1 and Earlier Releases	NFVIS 4.2.1 Release
Default file location for backup	/data/intdatastore/uploads/backup.bkup /mnt/extdatastore1/uploads/backup.bkup /mnt/extdatastore2/uploads/backup.bkup	/data/intdatastore/backup.bkup /mnt/extdatastore1/backup.bkup /mnt/extdatastore2/backup.bkup
Restore images and flavors	No	Yes
Unique Mac Uid for VM	No for NFVIS 3.12.3 and earlier release	Yes
Status monitoring	No	Yes

Feature	NFVIS 4.1.1 and Earlier Releases	NFVIS 4.2.1 Release
SNMP v3 user/passphrase restore (with uniqMacUid)	v3 user/passphrase restore	If system engine ID is the same as backup, restore all v3 users. If system engine ID is different from backup, ignore v3 users restoration.
SNMP engine ID restore on different system	No	Engine ID changed to same as backup bundle

VM backup using vmExportAction:

Feature	NFVIS 4.1.1 and Earlier Releases	NFVIS 4.2.1 Release
VM backup format	Full backup	Diff disk backup

Backup and Restore

To backup and save NFVIS and all VM configurations use **configuration-only** option. To backup and save VM disks, NFVIS and VM configurations use **configuration-and-vms** option.

You can only create a backup and save into datastore, or mounted USB storage device. Without specifying, the backup file will have *.bkup* extension.

	Backup configuration-only	Backup configuration-and-vms
Save system configurations	Yes	Yes
Save system upgrade configurations	Yes	Yes
Save system upgrade file	No	No
Save images and flavors configurations	Yes	Yes
Save image disks	No	Yes
Save deployments configurations	Yes	Yes
Save deployments disks	No	Yes

The following examples shows the backup options:

```
nfvis# hostaction backup configuration-and-vms file-path intdatastore:sample
```

```
nfvis# hostaction backup configuration-only file-path extdatastore2:sample-dir/sample
```

The following example shows the backup stored on a USB:

```
nfvis# hostaction backup configuration-only file-path usb:usb1/sample
```

Use the **hostaction backup force-stop** command to stop the running backup.

Starting from NFVIS 4.2 release, use the **show hostaction backup status** command to view the status of the overall backup process and each components like system, image and flavors, vm and so on. The following is an example of the show command output after the backup process is complete:

```
nfvis# show hostaction backup status
hostaction backup status 2020-07-16T07:02:44-00:00
destination intdatastore:backup_20200704.bkup
status      BACKUP-SUCCESS
size        "2798.0 MB"
components FIREWALL
  status      BACKUP-SUCCESS
  last update 2020-07-16T07:07:38-00:00
  size        "20.49 MB"
  details     ""
components Linux
  status      BACKUP-SUCCESS
  last update 2020-07-16T07:07:36-00:00
  size        "0.01 MB"
  details     ""
components NFS
  status      BACKUP-SUCCESS
  last update 2020-07-16T07:06:44-00:00
  size        "0.01 MB"
  details     ""
components NFVIS
  status      BACKUP-SUCCESS
  last update 2020-07-16T07:02:48-00:00
  size        "0.72 MB"
  details     ""
components ROUTER
  status      BACKUP-SUCCESS
  last update 2020-07-16T07:07:35-00:00
  size        "579.89 MB"
  details     ""
components VM_Images_Flavors
  status      BACKUP-SUCCESS
  last update 2020-07-16T07:06:45-00:00
  size        "2197.73 MB"
  details     ""
nfvis#
```

To restore a previous backup on an existing NFVIS setup or on a new NFVIS setup use except-connectivity option which preserves connectivity of the NFVIS and restores everything else from backup.

The restore is based on the system condition created during backup.

	Restore configuration-only	Restore configuration-and-vms
Restore system configurations	Yes	Yes

	Restore configuration-only	Restore configuration-and-vms
Restore upgrade configurations	yes, requires same upgrade files in system if the host backup was taken has such upgrade files. No, if host where backup was taken did not have any upgrade files registered. Restoree will fail.	Yes, requires same upgrade files in system if the host backup was taken has such upgrade files. No, if host where backup was taken did not have any upgrade files registered. Restore will fail.
Restore registered images and flavors	Yes, if images sources are still available (URL link is still valid, or uploaded files are still in the same locations). No, if images sources are not available (URL link is invalid, upload files are deleted or moved to new location). The restore process will fail.	Yes, restore from backup file.
Restore deployments	No	Yes, restore from backup file.

Note This means if there are upgrade files registered in the NFVIS. The backup create on this host will contain those information. If using this backup on new host or same host after factory-default-reset, the restore will fail.

	dpdk-disabled while backup	dpdk-enable while backup
dpdk-disabled while restore	Yes (system is dpdk-disabled)	Yes (system will beconverted to dpdk enabled, and VM vnic will be convereted inf needed)
dpdk-enabled while restore	No support	Yes (system is dpdk-enabled)

Note In hostaction restore process, the full file name (with *.bkup* extension) is required in the CLI.

```
nfvis# hostaction restore file-path intdatastore:sample.bkup
```

The following example shows how to restore a backup on a different NFVIS device:

```
nfvis# hostaction restore except-connectivity file-path extdatastore2:sample-dir/sample.bkup
```

Starting from NFVIS 4.2 release, use the **show hostaction restore-status** command to view the status of the overall restore process and each components like system, image and flavors, vm and so on. The following is an example of the show command output after the restore process is complete:

```
nfvis# show hostaction restore-status
hostaction restore-status 2020-07-16T07:18:54-00:00
source intdatastore:backup_20200704.bkup
status RESTORE-SUCCESS
components FIREWALL.vmbkp
  status      RESTORE-SUCCESS
  last update 2020-07-16T07:26:34-00:00
  details     ""
components Linux.vmbkp
  status      RESTORE-SUCCESS
  last update 2020-07-16T07:26:03-00:00
  details     ""
components NFS.vmbkp
  status      RESTORE-SUCCESS
  last update 2020-07-16T07:25:36-00:00
  details     ""
components NFVIS
  status      RESTORE-SUCCESS
  last update 2020-07-16T07:22:03-00:00
  details     ""
components ROUTER.vmbkp
  status      RESTORE-SUCCESS
  last update 2020-07-16T07:26:55-00:00
  details     ""
components VM_Images_Flavors
  status      RESTORE-SUCCESS
  last update 2020-07-16T07:26:01-00:00
  details     ""
components intdatastore:backup_20200704.bkup
  status      VERIFICATION-SUCCESS
  last update 2020-07-16T07:18:54-00:00
  details     ""
nfvis#
```

Starting from NFVIS 4.2 release, you can backup registered images and flavors into backup package and restore these images and flavors into the system. The new system does not require a pre-registered image before system restore. If the system has existing images, flavors or deployments, the system restore erases them all and restores from its own backup. VM backup is now faster and uses less disk space compared to NFVIS 4.1 release, but it also takes up additional process, time or disk space to backup registered images and flavors.

Backup, Restore, and Factory-Default-Reset

To perform **hostaction backup** -> **factory-default-reset** -> **hostaction restore** on the same box without any external storage (like USB or NFS mount), check the following issues:

	NFVIS 4.1.x and earlier releases	NFVIS 4.2.x and later releases
Backup file location	• The system backup bundle is saved under */datastore/uploads/* by default. • Factory-default-reset cleans up all files under */datastore/uploads/*, but leave files under */datastore/* intact. • **hostaction restore** requires backup bundle saved under */datastore/uploads/*. The restore process will not start if the backup bundle is saved in another location (bundle saved on USB or NFS should be copied to *datastore/uploads/ folder*).	• The system backup bundle is saved under */datastore/* by default.
System requirements if system backup bundle contains VM backups	• VM restoration requires the original image or template registered in NFVIS. • Factory-default-reset all clean ups all registered images and uploaded files. You need to configure minimum setup, like host connection and upload registered images to the same datastore.	• The backup package includes original registered images.

	NFVIS 4.1.x and earlier releases	NFVIS 4.2.x and later releases

Prevent backup bundle from deleting with factory-default-reset	• Save the backup bundle in remote locations. Then restore the connectivity and upload the backup bundle after reset.
	• Save backup bundle in local */datastore/* and not in */datastore/uploads/* or copy backup bundle from */datastore/uploads/* to */datastore/*:
	`# Backup & Restore on the same NFVIS box without NFS & USB`
	`# [[ BACKUP ]]`
	`# before executing factory-default-reset`
	`nfvis# nfvis# hostaction backup configuration-only file-path extdatastore1:configBackup-01.bkup`
	`nfvis# system file-copy source /mnt/extdatastore1/uploads/configBackup-01.bkup destination /mnt/extdatastore2/`
	`# after factory-default-reset all-except-images or all-except-images-connnectivity,`
	`# file /mnt/extdatastore1/uploads/configBackup-01.bkup will be deleted`
	`# but /mnt/extdatastore2/configBackup-01.bkup won't.`
	`# [[RESTORE]]`
	`# after NFVIS rebooted and login to console, copy file to uploads/ directory`
	`nfvis# system file-copy source /mnt/extdatastore2/configBackup-01.bkup destination /mnt/extdatastore2/uploads/`
	`nfvis# hostaction restore file-path extdatastore2:configBackup-01.bkup`
		• Save backup bundle in local */intdatastore/* and not in */intdatastore/uploads/* or copy backup bundle from */datastore/uploads/* to */datastore/*

You can copy backup file to intdatastore/ if there is sufficient storage space. If the backup is larger than free disk space in intdatastore/, you can copy to a remote server like scp or NFVIS web portal.

The following table lists the data erased and retained upon using NFVIS factory default reset options:

	Factory-default-reset all	Factory-default-reset all-except-images	Factory-default-reset all-except-images-connectivity
files under intdatastore	Retain	Retain	Retain
files under intdatastore/uploads/	Delete	Delete	Delete

	Factory-default-reset all	Factory-default-reset all-except-images	Factory-default-reset all-except-images-connectivity
files under extdatastore${1,2}	Delete	Retain	Retain
files under extdatastore${1,2}/uploads/	Delete	Delete	Delete
files under USB	Retain	Retain	Retain
files under NFS mounted datastore	Retain	Retain	Retain
Deployments	Delete	Delete	Delete
Registered Images and Flavors	Delete	Retain	Retain

Failure to Restore

NFVIS configurations fails to restore if:

- There is no sufficient disk space. Restore requires temporary disk space to save un-compressed files. You can move, copy or upload the backup file to a larger datastore and run system restore.

```
nfvis# show hostaction restore-status
hostaction restore-status 2020-07-16T21:29:08-00:00
source intdatastore:encs07-configVms-dpdk-2020-07101600.bkup
status RESTORE-ERROR
components intdatastore:encs07-configVms-dpdk-2020-07101600.bkup
  status      VERIFICATION-ERROR
  last update 2020-07-16T21:49:18-00:00
  details     "Backup package could not be inflated. No space left on device"
nfvis#
```

- The application communication fails. You can see this error after the first restore attempt has failed, and when you try to restore for the second time. You can reboot NFVIS before you attempt restore again.

```
nfvis# hostaction restore file-path extdatastore2:backup_20200704.bkup
Error: application communication failure
```

Dynamic SR-IOV

Dynamic Single-root input/output virtualization (SR-IOV) allows you to enable or disable SR-IOV on a Physical Network Interface Controller (PNIC). To disable SR-IOV on a PNIC, set the SR-IOV value to 0. To enable SR-IOV on a PNIC, set the SR-IOV value between 1 and the maximum number of virtual functions (maxvfs) supported on that PNIC. You can also create and delete SR-IOV networks based on the number of virtual functions (numvfs) set on that PNIC while enabling SR-IOV. The existing fresh installation behavior has not changed. Each PNIC has a number of VFs and SR-IOV networks created by default. You can use CLI, API, or the GUI to enable and disable SR-IOV on a PNIC and to create and delete SR-IOV networks.

Note The number of SR-IOV networks, numvfs or inusevfs, created per PNIC on fresh installation of NFVIS depends on the link speed of that particular pnic.

Restrictions and Limitations

- The supported platforms are CSP-2100, CSP-5228, CSP-5436, CSP-5444 (Beta), Cisco Catalyst 8200 UCPE, UCSC-C220-M5X, and UCS-E-M3.

  Dynamic SR-IOV is not supported on ENCS 5000 series.

- Dynamic SR-IOV is not supported on certain PNICs:
  - PNIC with driver i40e

Note PNIC with driver i40e is supported on default SR-IOV.

  - PNIC that does not support SR-IOV

- NFVIS release 3.12.1 supports Virtual Ethernet Bridge (VEB) in switch mode only.
- Resizing the number of virtual functions is not supported. SR-IOV should be disabled and then enabled with desired number of virtual functions.

Disable SR-IOV on a PNIC

To disable SR-IOV on a PNIC, ensure that:

- All SR-IOV networks on a PNIC are deleted.
- The PNIC is not attached to a bridge.

```
configure terminal
no pnic eth0-1 sriov
commit
```

Enable SR-IOV on a PNIC

To enable SR-IOV on a PNIC, ensure that:

- The PNIC supports SR-IOV.
- The numvfs field is populated with a value that is less than the maximum number of virtual functions (maxvfs) supported on that PNIC.
- The PNIC is not attached to a bridge.

```
configure terminal
pnic eth0-1 sriov numvfs 20
commit
```

To display the SR-IOV status of all PNICs, use the **show pnic sriov** command. To display the SR-IOV state of an individual PNIC use the **show pnic eth0-1 sriov** command.

Create SR-IOV Networks

To create SR-IOV networks, the PNIC must have SR-IOV enabled and configured with numvfs. The SR-IOV network name must have the following format: <pnic_name>-SRIOV-<num> where <pnic_name> is a valid PNIC name and <num> is a value that is greater than 0 and less than the number of VFs (numvfs).

To create an SR-IOV network in trunk mode:

```
configure terminal
networks network eth0-1-SRIOV-1 sriov true
commit
```

To create an SR-IOV network in access mode:

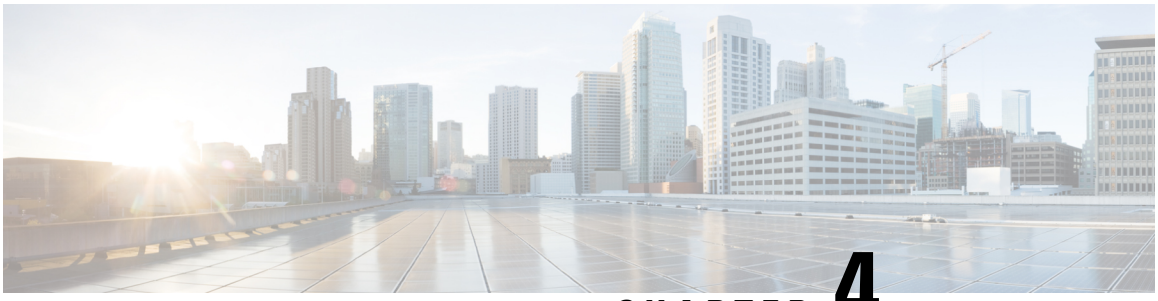
```
configure terminal
networks network eth0-1-SRIOV-1 sriov true trunk false vlan 30
commit
```

Delete SR-IOV Networks

To delete an SR-IOV network, ensure that no VMs are attached to the network.

```
configure terminal
no networks network eth0-1-SRIOV-1
commit
```

To verify the system networks, use the **show system networks** command.



CHAPTER 4

Cisco Network Plug-n-Play Support

Note Starting from 3.10.1 release, NFVIS is integrated with PnP 1.8.

The Cisco Network Plug and Play (Cisco Network PnP) solution provides a simple, secure, unified, and integrated offering for enterprise network customers to ease new branch or campus device rollouts, or for provisioning updates to an existing network. The solution provides a unified approach to provision enterprise networks comprising Cisco routers, switches, and wireless devices with a near zero touch deployment experience. This solution uses Cisco Application Policy Infrastructure Controller Enterprise Module (APIC-EM) to centrally manage remote device deployments.

Currently, you can use the Cisco Network Plug and Play client to:

- Auto discover the server
- Provide device information to the server
- Bulk provisioning of user credentials

Note NFVIS SD-Branch Zero Touch Provisioning (ZTP) PnP redirect fails if the WAN DHCP IP address is in the same subnet as management IP address 192.168.1.0/24. For a successful PnP redirect, you can change the DHCP pool to a different subnet or manually change the management IP address to a different subnet.

Bulk Provisioning of User Credentials

You can change the default user name and password of the devices using the Cisco Network PnP client. The Cisco Network PnP server sends the configuration file to Cisco Network PnP clients residing on multiple devices in the network, and the new configuration is automatically applied to all the devices.

Note For bulk provisioning of user credentials, ensure that you have the necessary configuration file uploaded to the Cisco APIC-EM. The following are the supported configuration formats:

Sample Format 1

```
<config xmlns="http://tail-f.com/ns/config/1.0">
  <rbac xmlns="http://www.cisco.com/nfv/rbac">
    <authentication>
      <users>
        <user>
          <name>admin</name>
          <password>Cisco123#</password>
          <role>administrators</role>
        </user>
        <user>
          <name>test1</name>
          <password>Test1239#</password>
          <role>administrators</role>
        </user>
        <user>
          <name>test2</name>
          <password>Test2985#</password>
          <role>operators</role>
        </user>
      </users>
    </authentication>
  </rbac>
</config>
```

Sample Format 2

If you use format 2, the system will internally convert this format into format 1.

```
<aaa xmlns="http://tail-f.com/ns/aaa/1.1">
    <authentication>
      <users>
        <user>
          <name>admin</name>
          <password>User123#</password>
         </user>
      </users>
    </authentication>
</aaa>
```


- PnP Discovery Methods, on page 60
- Configuring PnP Discovery Methods, on page 61
- PnP Action, on page 64

PnP Discovery Methods

When a device is powered on for the first time, the Cisco Network PnP agent discovery process, which is embedded in the device, starts in the absence of the startup configuration file, and discovers the IP address of the Cisco Network PnP server located in the Cisco APIC-EM. The Cisco Network PnP agent uses the following discovery methods:

- Static IP address—The IP address of the Cisco Network PnP server is specified using the **set pnp static ip-address** command.

- DHCP with option 43—The Cisco PnP agent automatically discovers the IP address of the Cisco Network PnP server specified in the DHCP option 43 string. For more details on how to configure DHCP for PnP server auto-discovery, see the Solution Guide for Cisco Network Plug and Play
- Domain Name System (DNS) lookup—If DHCP discovery fails to get the IP address of the PnP server, for example, because option 43 is not configured, the Cisco Plug and Play Agent falls back on a DNS lookup method. Based on the network domain name returned by the DHCP server, it constructs a fully qualified domain name (FQDN) for the PnP server, using the preset hostname "pnpserver". For more details on how to configure DNS for PnP server auto-discovery, see the Solution Guide for Cisco Network Plug and Play.

Note DNS FQDN Only lookup method is supported since 3.10.1 release.

- Cloud Redirection—This method uses the Cisco Cloud Device Redirect tool available in the Cisco Software Central.The Cisco Plug and Play Agent falls back on the Cloud Redirection method if DNS lookup is not successful.

Configuring PnP Discovery Methods

To enable static mode for PnP discovery using IPv4:

```
configure terminal
pnp automatic dhcp disable dhcp-ipv6 disable dns disable dns-ipv6 disable cco disable
cco-ipv6 disable
pnp static ip-address 192.0.2.8 port 80 transport http
commit
pnp action command restart
```

To enable static mode for PnP discovery using IPv6:

```
configure terminal
pnp automatic dhcp disable dhcp-ipv6 disable dns disable dns-ipv6 disable cco disable
cco-ipv6 disable
pnp static ipv6-address 0:0:0:0:0:ffff:c000:208 port 80 transport http
commit
pnp action command restart
```

Note Either IPv4 or IPv6 can be enabled at a time.

To enable static mode for PnP discovery using FQDN:

```
configure terminal
pnp static ip-address apic-em-fqdn.cisco.com port 80 transport http
commit
```

Note In FQDN support for PnP, domain names can be specified as an input. FQDN that is configured with IPv6 on a DNS server is not supported.

To enable automatic mode for PnP discovery using IPv4:

```
configure terminal
pnp automatic dhcp enable
pnp automatic dns enable
pnp automatic cco enable
pnp automatic timeout 100
commit
```

Note By default, the automatic discovery mode for DHCP, DNS, and CCO is enabled. You can enable or disable the options as required. For example, you can enable all options or keep one enabled, and the rest disabled.

To enable automatic mode for PnP discovery using IPv6:

```
configure terminal
pnp automatic dhcp-ipv6 enable
pnp automatic dns-ipv6 enable
pnp automatic cco-ipv6 enable
pnp automatic timeout 30
commit
```

Note You cannot disable both static and automatic PnP discovery modes at the same time. You must restart PnP action every time you make changes to the PnP discovery configuration. You can do this using the **pnp action command restart**.

Verifying the PnP Status

Use the **show pnp** command in privileged EXEC mode to verify the configuration of PnP discovery methods. The following sample output shows that the static discovery mode is enabled, and the automatic discovery mode is disabled.

```
nfvis# show pnp
pnp status response "PnP Agent is running\n"
pnp status ip-address 192.0.2.8
pnp status ipv6-address ""
pnp status port 80
pnp status transport http
pnp status cafile ""
pnp status created_by user
pnp status dhcp_opt43 0
pnp status dns_discovery 0
pnp status cco_discovery 0
pnp status dhcp-ipv6 0
pnp status dns-ipv6 0
pnp status cco-ipv6 0
pnp status timeout 100
```

```
nfvis#

FQDN
nfvis# show pnp
pnp status response "PnP Agent is running\nserver-connection\n    status: Success\n    time:
 06:23:11 Jun 17\ndevice-info\n    status: Success\n    time: 06:23:06 Jun 17\nbackoff\n
  status: Success\n    time: 06:23:11 Jun 17\ncertificate-install\n    status: Success\n
  time: 06:21:38 Jun 17\ncli-exec\n    status: Success\n    time: 06:22:50 Jun 17\ntopology\n
    status: Success\n    time: 06:23:00 Jun 17\n"
pnp status ip-address apic-em-fqdn.cisco.com
pnp status ipv6-address ""
pnp status port 443
pnp status transport https
pnp status cafile /etc/pnp/certs/trustpoint/pnplabel
pnp status created_by user
pnp status dhcp_opt43 0
pnp status dns_discovery 0
pnp status cco_discovery 0
pnp status dhcp-ipv6 0
pnp status dns-ipv6 0
pnp status cco-ipv6 0
pnp status timeout 0
nfvis#
```

The following sample output shows that the static discovery mode is disabled, and the automatic discovery mode is enabled for DHCP, DNS, and CCO:

```
DHCP:
nfvis# show pnp
pnp status response "PnP Agent is running\nserver-connection\n    status: Success\n    time:
 05:05:59 Jun 17\ninterface-info\n    status: Success\n    time: 05:05:56 Jun
17\ndevice-info\n    status: Success\n    time: 05:05:38 Jun 17\nbackoff\n    status:
Success\n    time: 05:05:59 Jun 17\ncapability\n    status: Success\n    time: 05:05:44 Jun
 17\ncertificate-install\n    status: Success\n    time: 05:01:19 Jun 17\ncli-exec\n
status: Success\n    time: 04:58:29 Jun 17\ntopology\n    status: Success\n    time: 05:05:49
 Jun 17\n"
pnp status ip-address 192.0.2.8
pnp status ipv6-address ""
pnp status port 443
pnp status transport https
pnp status cafile /etc/pnp/certs/trustpoint/pnplabel
pnp status created_by dhcp_discovery
pnp status dhcp_opt43 1
pnp status dns_discovery 1
pnp status cco_discovery 1
pnp status dhcp-ipv6 1
pnp status dns-ipv6 1
pnp status cco-ipv6 1
pnp status timeout 60

DNS:
nfvis# show pnp
pnp status response "PnP Agent is running\nserver-connection\n    status: Success\n    time:
 05:13:55 Jun 17\ndevice-info\n    status: Success\n    time: 05:13:49 Jun 17\nbackoff\n
  status: Success\n    time: 05:13:55 Jun 17\ncertificate-install\n    status: Success\n
  time: 05:12:26 Jun 17\ncli-exec\n    status: Success\n    time: 05:13:34 Jun 17\ntopology\n
    status: Success\n    time: 05:13:45 Jun 17\n"
pnp status ip-address pnpserver.apic-em-fqdn.cisco.com
pnp status ipv6-address ""
pnp status port 443
pnp status transport https
pnp status cafile /etc/pnp/certs/trustpoint/pnplabel
pnp status created_by dns_discovery
```

```
pnp status dhcp_opt43 1
pnp status dns_discovery 1
pnp status cco_discovery 1
pnp status dhcp-ipv6 1
pnp status dns-ipv6 1
pnp status cco-ipv6 1
pnp status timeout 60
```

CCO:

```
nfvis# show pnp
pnp status response "PnP Agent is running\nserver-connection\n    status: Success\n    time:
 05:24:25 Jun 17\ninterface-info\n    status: Success\n    time: 05:23:13 Jun
17\ndevice-info\n    status: Success\n    time: 05:23:01 Jun 17\nbackoff\n    status:
Success\n    time: 05:24:25 Jun 17\ncapability\n    status: Success\n    time: 05:23:06 Jun
 17\nredirection\n    status: Success\n    time: 05:09:43 Jun 17\ncli-exec\n    status:
Success\n    time: 05:09:53 Jun 17\ncertificate-install\n    status: Success\n    time:
05:18:43 Jun 17\ntopology\n    status: Success\n    time: 05:23:10 Jun 17\n"
pnp status ip-address 192.0.2.8
pnp status ipv6-address ""
pnp status port 443
pnp status transport https
pnp status cafile /etc/pnp/certs/trustpoint/pnplabel
pnp status created_by cco_discovery
pnp status dhcp_opt43 1
pnp status dns_discovery 1
pnp status cco_discovery 1
pnp status dhcp-ipv6 1
pnp status dns-ipv6 1
pnp status cco-ipv6 1
pnp status timeout 60
```

PnP Server APIs and Commands

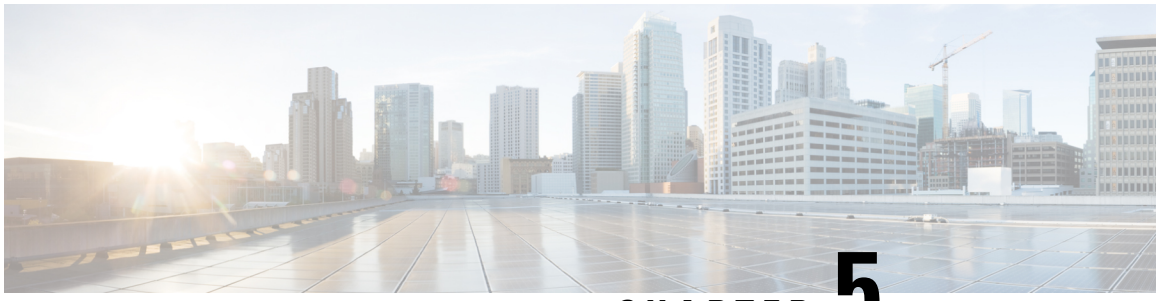
PnP Server APIs	PnP Server Commands
• /api/config/pnp • /api/config/pnp?deep	• pnp static ip-address • pnp automatic • show pnp

PnP Action

You can start, stop, and restart any PnP action using the PnP action command or API.

PnP Action API and Command

PnP Action API	PnP Action Command
• /api/operations/pnp/action	• pnp action command



CHAPTER 5

Overview of VM Life Cycle Management

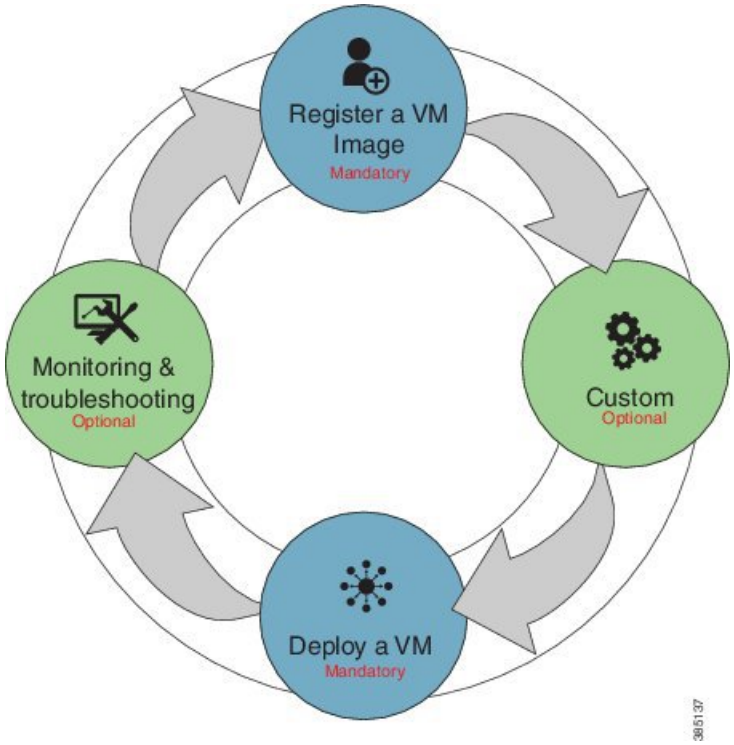
VM life cycle management refers to the entire process of registering, deploying, updating, monitoring VMs, and getting them service chained as per your requirements. You can perform these tasks and more using a set of REST APIs or NETCONF commands or the Cisco Enterprise NFVIS portal.

- Workflow of VM Life Cycle Management, on page 65
- Uploading VM Images to an NFVIS Server, on page 67
- VM Bootstrap Configuration Options with a VM Deployment, on page 68
- OpenStack Configuration Drive Support for Third Party VMs, on page 69
- Performing Resource Verification, on page 70
- Configuring Management IP Subnet, on page 71
- VM States, on page 71

Workflow of VM Life Cycle Management

The following diagram depicts the basic workflow of the VM life cycle management using REST APIs:

Figure 5: VM Life Cycle Management



1. **Register a VM Image—**To create a VM instance, you must register the source image in the NFVIS image repository. Registering the image is a one-time activity. The source image is expected to be hosted on a remote server. In the Registration API you specify the URL on the remote server, where the file is located. When invoked, device pulls the API info, completes the file transfer and registration. Once registered, you can create one or more VM instances. Remote servers supported for hosting the source images (.tar.gz files) are HTTP/HTTPS servers. Starting with NFVIS 4.12.1 release, remote servers can also be FTP servers or SCP servers.

2. **Customizing the VM—**After registering a VM image, you can optionally create a custom profile or flavor for the VM image if the profiles defined in the image file do not match your requirement. The flavor creation option lets you provide specific profiling details for a VM image, such as the virtual CPU on which the VM will run, and the amount of virtual memory the VM will consume.

   Depending on the topology requirement, you can create additional networks and bridges to attach the VM to during deployment.

3. **Deploy a VM—** A VM can be deployed using the deployment API. The deployment API allows you to provide values to the parameters that are passed to the system during deployment. Depending on the VM you are deploying, some parameters are mandatory and others optional.

4. **Manage and Monitor a VM—**You can monitor a VM using APIs and commands that enable you to get the VM status and debug logs. Using VM management APIs, you can start, stop, or reboot a VM, and view statistics for a VM such as CPU usage.

   A VM can also be managed by changing or updating its profile. You can change a VM's profile to one of the existing profiles in the image file; alternatively, you can create a new custom profile for the VM.

   The vNICs on a deployed VM can also be added or updated.

Note Before performing the VM life cycle management tasks, you will have to upload the VM images to the NFVIS server or http/s server.

For details on APIs, see the VM Lifecycle Management APIs chapter in the *API Reference for Cisco Enterprise Network Function Virtualization Infrastructure Software*.

Uploading VM Images to an NFVIS Server

You can upload VM images to an NFVIS server in the following ways. The files are copied to the default location (/data/intdatastore/uploads) on the host server.

- Copy the images from your local system to the NFVIS server—Use the **Image Upload** option from the Cisco Enterprise NFVIS portal.
- Copy the images using the USB drive—Ensure that you have plugged the USB drive that contains the required images into the server before mounting the USB drive.
- Copy using the **scp** command (scp username@external_server:/path/image.tar.gz intdatastore:image.tar.gz).

To copy an image using the USB device:

```
configure terminal
system usb-mount mount ACTIVE
system file-copy usb file name usb1/package/isrv-universalk9.16.03.01.tar.gz
commit
```

Note Use the **show system file-list disk usb** command in privileged EXEC mode to view a list of files available with the mounted USB drive. To save space, you can delete all unwanted text and TAR files from the default location using the **system file-delete** command in global configuration mode.

Verifying the Image Copied from the USB Drive

After copying the file from the USB drive to the host server, you can verify the file using the **show system file-list disk local** command:

```
nfvis# show system file-list disk local

SI NO     NAME     PATH                                                    SIZE  TYPE    DATE MODIFIED
-------------------------------------------------------------------------------------------------------------
1 lastlog-20170314.gz /data/intdatastore/logs/2017-03/14/10-00 337 Other 2017-03-14 21:55:42

2 escmanager-tagged-log.log-20170314.gz /data/intdatastore/logs/2017-03/14/10-00 167K Other
 2017-01-18 05:58:26
3 confd_audit.log-20170317.gz /data/intdatastore/logs/2017-03/17/09-30 4.6K Other 2017-03-17
 21:29:59
4 esc_postinit.log-20170317.gz /data/intdatastore/logs/2017-03/17/05-00 605K Other 2017-03-17
 16:40:19
5 error.log-20170317.gz /data/intdatastore/logs/2017-03/17/05-00 1.3K Other 2017-03-17
16:40:15
```

```
6 ovs-ctl.log-20170317.gz /data/intdatastore/logs/2017-03/17/12-00 20 Other 2017-03-16
00:00:01 4:01
!
!
!
62 ovs-ctl.log-20170323.gz /data/intdatastore/logs/2017-03/23/12-00 20 Other 2017-03-22
00:00:01
63 CentOS-7-x86_64-Everything-1511.ova /data/intdatastore/uploads 1.1G VM 2017-03-15 19:20:03
 Package
64 TinyLinux.tar.gz /data/intdatastore/uploads 17M VM 2017-03-15 18:25:00 Package
65 Cisco-KVM-vWAAS-1300-6.3.0-b98.tar.gz /data/intdatastore/uploads 979M VM 2017-03-15
19:19:11 Package
66 ubuntu_14.04.3-server-amd64-disk1.tar /data/intdatastore/uploads 527M VM 2017-03-15
19:20:17.gz Package
67 asav961.tar.gz /data/intdatastore/uploads 164M VM 2017-03-15 18:24:57 Package
68 isrv-universalk9.16.03.01.tar.gz /data/intdatastore/uploads 1.3G VM 2017-03-15 19:19:53
```

Related APIs and Commands

APIs	CLI Commands
• /api/operations/system/file-copy/usb/file • /api/config/system/usb-mount	• system file-copy usb file name • system usb-mount mount ACTIVE • system file-delete • show system file-list disk usb • show system file-list disk local

VM Bootstrap Configuration Options with a VM Deployment

You can include the bootstrap configuration (day zero configuration) of a VM in the VM deployment payload in the following ways:

- Bundle bootstrap configuration files into the VM package: In this method, the bootstrap configuration variables can be assigned tokens. Token names must be in bold text. For each variable with a token assigned, key-value pairs must be provided during deployment in the deployment payload.
- Bootstrap configuration as part of the deployment payload: The entire bootstrap configuration is copied to the payload without tokens.
- Bootstrap configuration file in the NFVIS server: In this method, the configuration file is copied or downloaded to the NFVIS server, and referenced from the deployment payload with the filename, which includes the full path.

For examples on how to use bootstrap configuration options in the deployment payload, see the API Reference for Cisco Enterprise Network Function Virtualization Infrastructure Software.

OpenStack Configuration Drive Support for Third Party VMs

To enable staging of bootstrap configuration files at the time of a third party VM deployment as per OpenStack standards, the following cloud init format is supported:

```
openstack/content
openstack/content/0000
openstack/content/0001
openstack/latest/meta_data.json
```

In the above sample, the "0000" and "0001" files are the actual bootstrap files from the deployment payload. A third party VM can use the init file to fetch its configuration files.

The following metadata file is used to provide the file path on the configuration drive and reference to the actual bootstrap configuration files.

```
{
    "files": [
        {
            "content_path": "/content/0000",
            "path": "/config/day-0.txt"
        },
        {
            "content_path": "/content/0001",
            "path": "/sample/path/iosxe_config.txt"
        }
    ]
}
```

With this implemenation, two copies of the same bootstrap configuration file will be present on the virtual CD-ROM package. The first version at the root (iosxe_config.txt) and the second inside the "openstack/content" folder.

The admin will also have to sepcify the bootstrap configuartion file in the image properties file before packaging the VM.

Example for the Bootstrap Configuartion File in the Image Properties File

```
--optimize=OPTIMIZE [REQUIRED] optimized VM: --optimize=true/false;
 --root_file_disk_bus=ROOT_FILE_DISK_BUS root disk file type:
 --root_file_disk_bus=virtio/ide; default is virtio
 --virtual_interface_model=VIRTUAL_INTERFACE_MODEL
 --virtual_interface_model=rtl8139; default is none
 --thick_disk_provisioning=THICK_DISK_PROVISIONING
 --thick_disk_provisioning=true; default is false
 --bootstrap_cloud_init_bus_type=BOOTSTRAP_CLOUD_INIT_BUS_TYPE
 --bootstrap_cloud_init_bus_type=virtio; default is ide
 --bootstrap_cloud_init_drive_type=BOOTSTRAP_CLOUD_INIT_DRIVE_TYPE
 --bootstrap_cloud_init_drive_type=disk; default is cdrom
 --bootstrap=BOOTSTRAP bootstrap file/s for VM (two parameters required in the format of
dst:src; dst filename including path has to match exactly to what the VM expects;
 upto 20 bootstrap files are accepted.)
 examples:
 --bootstrap ovf-env.xml:file1,ios-xe.txt:file2 for ISRv; both files get mounted at the
root level on the VM.
 --bootstrap day0-config:filename1 for ASAv
 --bootstrap
 /:bootstrap.xml,/license/lic.txt:license.txt
```

```
 bootstrap.xml get mounted as bootstrap.xml at root, and license.txt get mounted as
/license/lic.txt.
```

Note If any of the strings in the configuration file has wild characters, wrap the string with this #[[ ]]# so that the token/key replacement engine does not consider wild characters as key or token, and looks for key value pairs to replace during a VM deployment.

For details on the OpenStack standards, visit *http://docs.openstack.org*.

Performing Resource Verification

Given below are the APIs and commands to perform different types of resource verification:

Task	API	Command
To display CPU information for each CPU or the user specified CPU, and the VMs pinned to the CPU	• /api/operational/resources/cpu-info/cpus • /api/operational/resources/cpu-info/cpus/cpu • /api/operational/resources /cpu-info/cpus/cpu/<cpu-id>	show resources cpu-info cpus
To display information on the VMs running in all the physical CPUs or a specific physical CPU in the system	• /api/operational/resources/cpu-info/vnfs • /api/operational/resources/cpu-info/vnfs/vnf • /api/operational/resources/cpu-info/vnfs/vnf/ <deployment_name>.<vm_group_name>	show resources cpu-info vnfs
To get information on the number of CPUs allocated to VMs and the CPUs that are already used by the VMs	/api/operational/resources/cpu-info/allocation	show resources cpu-info allocation

Note To display information on all CPUs, VMs pinned to the CPUs, and VMs allocated to the CPUs, use the **show resources cpu-info** command.

CPU Over-Subscription

Cisco Enterprise NFVIS does not allow CPU over-subscription for low-latency network appliance VMs (for example, Cisco ISRv and Cisco ASAv). However, the CPU over-subscription is allowed for non low-latency VMs (for example, Linux Server VM and Windows Server VM).

Configuring Management IP Subnet

By default, all VMs with management interfaces will be provisioned with an IP address in the subnet of 10.20.0.1. To change the default subnet, the following commands needs to be executed in a sequence to first delete an exiting subnet and then add a new subnet in the network. For these commands to work, ensure there is no managed VNF's in the system before you change management network address.

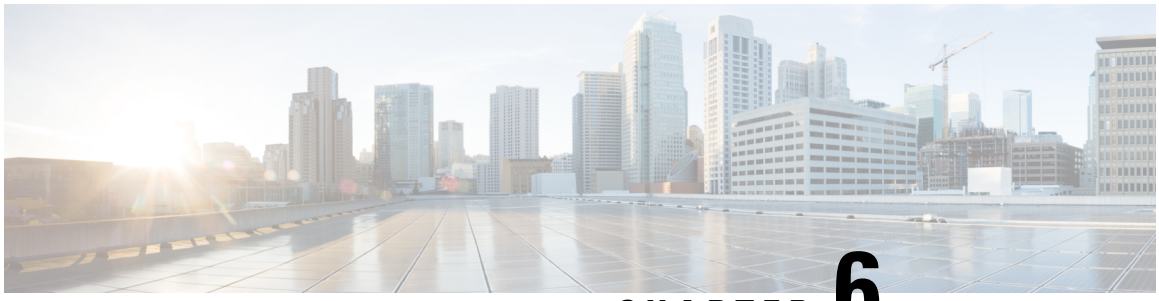
To delete an existing subnet use **no vm_lifecycle networks network int-mgmt-net subnet int-mgmt-net-subnet** command.

To create a new subnet:

```
configure terminal
vm_lifecycle networks network int-mgmt-net subnet int-mgmt-net-subnet address 105.20.0.0
gateway 105.20.0.1 netmask 255.255.255.0 dhcp false
commit
```

VM States

VM States	Description
VM_UNDEF_STATE	The initial state of VM or VNF before deployment of this VM.
VM_DEPLOYING_STATE	VM or VNF is being deployed on to the NFVIS.
VM_MONITOR_UNSET_STATE	VM or VNF is deployed in the NFVIS but the monitoring rules are not applied.
VM_MONITOR_DISABLED_STATE	Due to a VM action request or recovery workflow, the monitoring or KPI rules applied on the VM or VNFs were not enabled.
VM_STOPPING_STATE	VM or VNF is being stopped.
VM_SHUTOFF_STATE	VM or VNF is in stopped or shutoff state.
VM_STARTING_STATE	VM or VNF is being started.
VM_REBOOTING_STAT	VM or VNF is being rebooted.
VM_INERT_STATE	VM or VNF is deployed but not alive. The KPI monitor is applied and waiting for the VM to become alive.
VM_ALIVE_STATE	VM or VNF is deployed and successfully booted up or alive as per the monitor or kpi metric.
VM_UNDEPLOYING_STATE	VM or VNF is being undeployed or terminated.
VM_ERROR_STATE	VM or VNF will be in error state if deployment or any other operation is failed.



CHAPTER 6

VM Deployment Scenarios

This chapter provides details on the following deployment scenarios using REST APIs. As an example, the Cisco ENCS is used to illustrate these scenarios.

- Single VM deployment
- Service chaining with two VMs
- Service chaining of multiple VMs with Windows or Linux servers

  The following VM images are used to explain the deployment scenarios:

  - Cisco Integrated Services Router (ISRv) —isrv-03.16.02
  - Cisco Adaptive Security Virtual Appliance (ASAv)— asav951-201
  - Linux server—ubuntu-14.04.3-server-amd64-disk1

- Registering VM Images, on page 73
- Service Chaining of VMs, on page 77

Registering VM Images

You must register all VM images before deploying them.



Note Register all the VM images required for the VM deployment depending on the topology. A VM image registration is done only once per VM image. You can perform multiple VM deployments using the registered VM image.

To register a Cisco ISRv image:

1. Set up the http/https server to host the VM image, or upload the image to the NFVIS server using the local portal or the **scp** command.
2. Register the Cisco ISRv image using the following API method:

```
curl -k -v -u admin:admin -H Accept:application/vnd.yang.data+xml -H
Content-Type:application/vnd.yang.data+xml -X
```

```
POST https://<NFVIS_IP>/api/config/vm_lifecycle/images -d
'<image><name>isrv-k9.16.03.01</name><src>http://filename_with_full-path-of
the-file/isrv-universalk9.16.03.01.tar.gz</src></image>'
```

3. Verify the image status using the following API method:

```
curl -k -v -u admin:admin -H Accept:application/vnd.yang.data+xml -H
Content-Type:application/vnd.yang.data+xml -X
GET
https://<NFVIS_IP>/api/operational/vm_lifecycle/opdata/images/image/isrv-9.16.03.01?deep
```

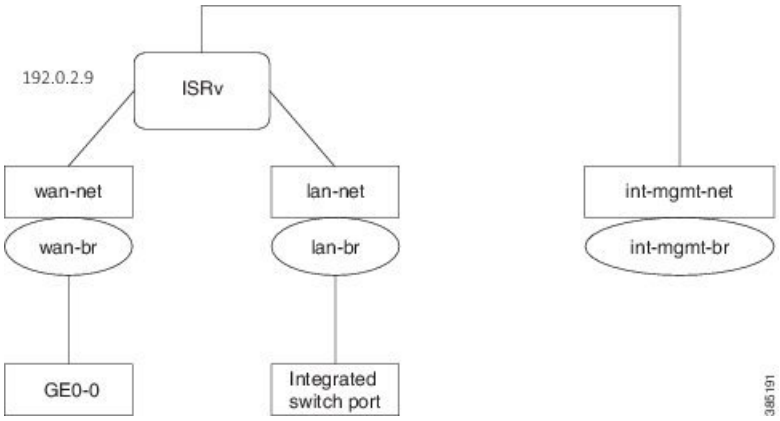
4. Now, repeat Steps 1 to 3 to register the Cisco ASAv and Linux server images. Ensure that you provide the exact image name and source file location when running the API commands.

Note You can run API commands from any console/server that can reach Cisco Enterprise NFVIS.

Single VM Deployment

In this example, a Cisco ISRv image with three network interfaces is deployed. The following diagram illustrates the deployment topology:

Figure 6: Single VM Deployment



Steps for Deploying a VM

To deploy a Cisco ISRv image:

1. Verify that all networks required for your deployment are configured.

```
curl -k -v -u admin:admin -H content-type:application/vnd.yang.data+xml -X
GET https://<NFVIS_IP>/api/config/networks?deep
```

2. Before deploying the VM, you can perform a resource check to ensure that you have sufficient resources for the deployment.

```
curl -k -v -u "admin:admin" -X GET
https://<NFVIS_IP>/api/operational/resources/precheck/vnf/newvnf,isrv-small,true
?deep
```

3. Deploy the Cisco ISRv VM.

```
curl -k -v -u admin:admin -H Accept:application/vnd.yang.data+xml -H
Content-Type:application/vnd.yang.data+xml -X
POST https://<NFVIS_IP>/api/config/vm_lifecycle/tenants/tenant/admin/deployments --data
<deployment>
    <name>ISR</name>
    <vm_group>
        <name>ISR</name>
        <image>isrv-universalk9.16.03.01/image>
        <bootup_time>600</bootup_time>
        <recovery_wait_time>0</recovery_wait_time>
        <recovery_policy>
            <action_on_recovery>REBOOT_ONLY</action_on_recovery>
        </recovery_policy>
        <flavor>isrv-small</flavor>
        <interfaces>
            <interface>
                <nicid>0</nicid>
                <network>int-mgmt-net</network>
                <port_forwarding>
                    <port>
                        <type>ssh</type>
                        <protocol>tcp</protocol>
                        <vnf_port>22</vnf_port>
                        <external_port_range>
                            <start>20022</start>
                            <end>20022</end>
                        </external_port_range>
                    </port>
                </port_forwarding>
            </interface>
            <interface>
                <nicid>1</nicid>
                <network>lan-net</network>
                <ip_address>209.165.201.0</ip_address>
            </interface>
            <interface>
                <nicid>2</nicid>
                <network>wan-net</network>
                <ip_address>209.165.201.1</ip_address>
            </interface>
        </interfaces>
        <scaling>
            <min_active>1</min_active>
            <max_active>1</max_active>
        </scaling>
        <kpi_data>
            <kpi>
                <event_name>VM_ALIVE</event_name>
                <metric_value>1</metric_value>
                <metric_cond>GT</metric_cond>
                <metric_type>UINT32</metric_type>
                <metric_collector>
                    <type>ICMPPing</type>
```

```
                    <nicid>0</nicid>
                    <poll_frequency>3</poll_frequency>
                    <polling_unit>seconds</polling_unit>
                    <continuous_alarm>false</continuous_alarm>
                </metric_collector>
            </kpi>
        </kpi_data>
        <rules>
            <admin_rules>
                <rule>
                    <event_name>VM_ALIVE</event_name>
                    <action>ALWAYS log</action>
                    <action>TRUE servicebooted.sh</action>
                    <action>FALSE recover autohealing</action>
                </rule>
            </admin_rules>
        </rules>
        <config_data>
            <configuration>
                <dst>bootstrap_config</dst>
                <variable>
                    <name>TECH_PACKAGE</name>
                    <val>security</val>
                </variable>
                <variable>
                    <name>ngio</name>
                    <val>enable</val>
                </variable>
            </configuration>
        </config_data>
    </vm_group>
</deployment>
```

Note If two VM's are connected to the same physical interface, one over SRIOV and another over virtio net from a bridge connected to the packet filter, you cannot ping between the two VM's between these interfaces. Use SRIOV or virtio net on both the VM's to connect to the packet filter over a bridge.

4. Verify the deployment status.

```
curl -k -v -u admin:admin -H Accept:application/vnd.yang.data+xml -H
Content-Type:application/vnd.yang.data+xml -X
GET
https://NFVIS_IP/api/operational/vm_lifecycle/opdata/tenants/tenant/admin/deployments/ISR,-,-?deep
```

Note To enable NIM support on a Cisco ISRv running on Cisco ENCS, you must use the following variable in the ISRv deployment payload.

```
<variable>
     <name>ngio</name>
     <val>enable</val>
</variable>
```

Service Chaining of VMs

Service chaining here refers to a set of network services in the form of VMs using an intermediate network. Cisco Enterprise NFVIS supports service chaining of two or more VMs eliminating the need of dedicated hardware devices for different types of network services.

To service chain traffic between two or more VMs, you will have to create the following:

- Bridge—For example, you can create a new bridge called sc-br.
- Network—For example, you can create a new network called sc-net.
- Launch VM1 and VM2 with an interface from each VM to the service chain network (sc-net).

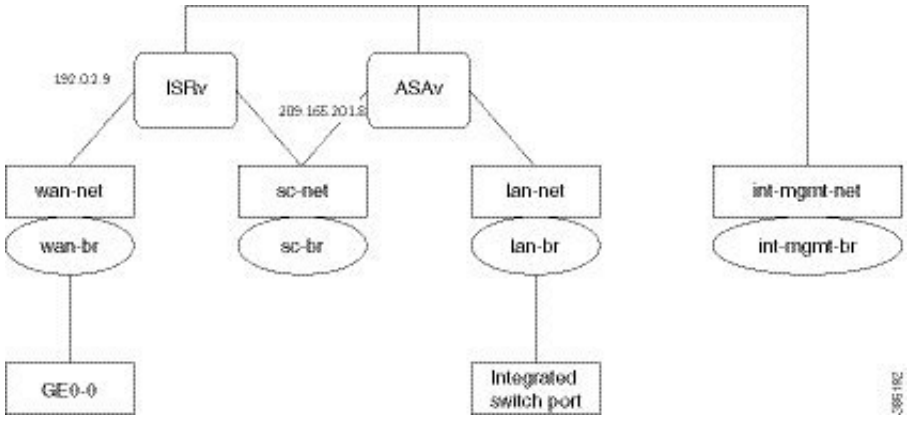
For more details on how to configure service chaining using APIs, see the following topics:

- Service Chaining with two VM Images, on page 77
- Service Chaining of Multiple VMs with Windows or Linux Servers, on page 78

Service Chaining with two VM Images

In this example, a Cisco ISRv VM and a Cisco ASAv VM are service chained. For that, you will have to deploy both VMs.

Figure 7: Service Chaining with two VM Images



Steps for Service Chaining with Two VM Images

1. Create a new bridge for service chaining.

```
curl -k -v -u admin:admin -H content-type:application/vnd.yang.data+xml -X POST
https://<NFVIS_IP>/api/config/bridges --data
'<bridge><name>sc-br</name></bridge>'
```

2. Create a new network for service chaining, and attach the bridge to the network.

```
curl -k -v -u admin:admin -H content-type:application/vnd.yang.data+xml -X POST
https://<NFVIS_IP>/api/config/networks --data
```

```
'<network><name>sc-net</name><bridge>sc-br</bridge> </network>'
```

3. Verify that all bridges and networks are configured.
4. Deploy the Cisco ISRv VM, and verify the deployment status.
5. Deploy the cisco ASAv VM, and verify the deployment status.

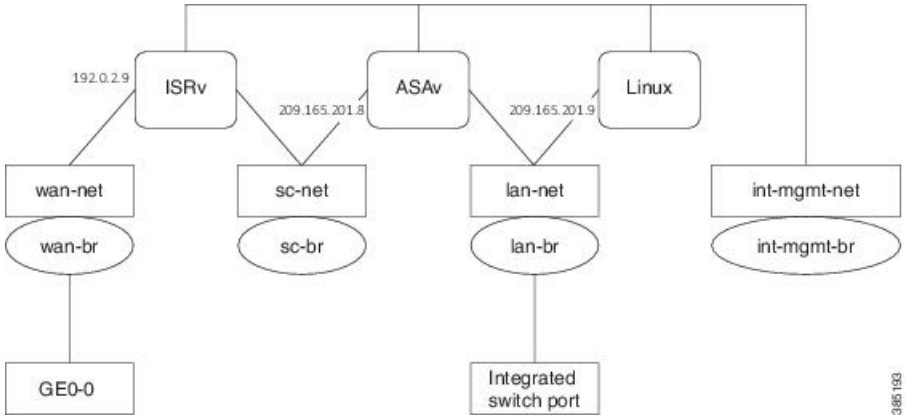
   See Steps for Deploying a VM , on page 74 for API command details for Steps 3 to 5.

Service Chaining of Multiple VMs with Windows or Linux Servers

In this example, multiple VMs will be service chained. Cisco ISRv and Cisco ASAv VMs can be deployed as explained in Service Chaining with two VM Images, on page 77.

This section covers Linux server deployment (Windows 2012 server can also be deployed using the same steps.)

Figure 8: Service Chaining of Multiple VMs with Windows or Linux Servers

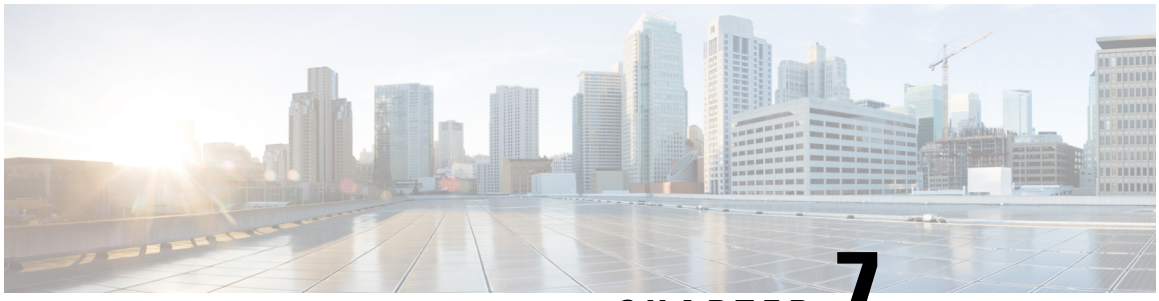


Steps for Service Chaining of Multiple VMs with Windows or Linux Servers

1. Create networks and bridges as required.

   See Steps 1 and 2 in Steps for Service Chaining with Two VM Images, on page 77 for details on creating networks and bridges.
2. Deploy Cisco ISRv and Cisco ASAv, and verify their deployment status.
3. Deploy the Linux server VM.
4. Verify the server deployment status.

   See the Steps for Deploying a VM , on page 74 for API command details for Steps 2 to 4.



CHAPTER 7

SPAN Session or Port Mirroring

- About SPAN Sessions, on page 79
- Configuring SPAN Sessions, on page 79
- Configuration Examples for SPAN Session Scenarios, on page 81

About SPAN Sessions

The Switched Port Analyzer (SPAN) or Port Mirroring feature helps you analyze network traffic passing through interfaces or VLANs by using SPAN sessions. The SPAN sessions send a copy (mirror) of the traffic to another interface or VLAN on the switch that has been connected to a network analyzer or monitoring device. SPAN does not affect the switching of network traffic on the source interfaces.

Note You must dedicate a destination port for SPAN use. Except for traffic that is required for the SPAN session, destination ports do not receive or forward traffic. When the SPAN is configured on the system, there might be some performance hit.

SPAN Session Interfaces

The interface can be:

- Physical interface
- LAN SRIOV
- VM's vNIC (virtio net)

In the case of virtio net or SRIOV VF, you have to specify the VM group name and NIC ID of the VM interface. If the VM vNIC is virtio net type, then the SPAN session is applied on the OVS bridge. If VM vNIC is SRIOV VF, then the mirror is applied to the hardware bridge. The interface name is specified for a physical interface, for example, GE0-0 or eth0.

Configuring SPAN Sessions

The SPAN session configuration has the following four parameters:

- Session number—Each SPAN session is identified with a unique number.
- Bridge name—The SPAN session is applied to a bridge. For VLAN mirroring, the bridge must be specified. The bridge name is optional if the source or destination interface is configured for the session.
- Source configuration—The source of the mirror traffic can be one of the following:
  - Packets entering (Rx), or exiting (Tx), or both. You can specify multiple interfaces of any type.
  - You can also specify all interfaces on the OVS bridge.
  - All packets entering a VLAN. You can also specify a list of VLANs.
- Destination configuration—The destination for the mirrored traffic can be one of the following:
  - You can send the mirrored traffic to interfaces of Cisco ENCS switch port when you use a physical interface.
  - The mirrored traffic can be sent to a speciifc VLAN. In this case, the original VLAN tag is stripped in the mirrored traffic in favor of the destination VLAN. This loss of original VLAN information might make the mirrored traffic hard to interpret.



Note Use VM Group Name for **source vm-vnic** and **destination vm-vnic** commands. Use the **show vm_lifecycle deployments all | include Group** command to get Group Name.

```
ENCS# show vm_lifecycle deployments all | include Group
VM Group Name : deployment-ASAv
VM Group Name : deployment-ISRv
VM Group Name : deployment-Linux
ENCS#
```

To configure a SPAN session:

```
configure terminal
monitor session  2
bridge wan-br
source interface GE0-0
destination vm-vnic Linux2 0
commit
```

Verifying the SPAN Session Configuration

Use the **show system monitor session** command to verify the SPAN session configuration.

```
nfvis# show system monitor session
system monitor session 2
 bridge                wan-br
 destination_vlan      ""
 destination_interface vnic0
 source_vlans          ""
 source_rx_interfaces  "GE0-0"
 source_tx_interfaces  "GE0-0"
 source_all            false
 statistics            "tx_bytes=142660, tx_packets=1380"
```

Use the **show running-config monitor session** command to verify the interface configuration for a SPAN session:

```
nfvis# show running-config monitor session
monitor session 2
 destination vm-vnic Linux2 0
 source vm-vnic Linux1 0 both
 source interface GE0-0 both
```

SPAN Session APIs and Commands

APIs	Commands
• /api/config/monitor • /api/operational/monitor\?deep • /api/config/monitor\?deep • /api/operational/system/monitor/session\?deep	• monitor session • bridge • source • destination • show system monitor session • show monitor session status • show running-config monitor session

Configuration Examples for SPAN Session Scenarios

Example: SPAN Session Traffic on a Physical Interface

The following example shows how to configure all traffic coming in or going out on GE0-0 (physical interface) and VM Linux1 (vnic0). And traffic is mirrored to the VM Linux2 (vnic1). With this configuraton, any traffic arriving on vnet1 will be dropped.

Note An existing SPAN session will be in FAIL state after the system reboot. In this case, you need to recreate (delete and create) the SPAN session after the system bootup.

VM deployment interfaces:

- SPAN source: GE0-0 (traffic in both directions)
- SPAN source: Linux1/vnic0, and wan-net (traffic in both directions)
- • SPAN destination: Linux2/vnic0, and wan-net

```
nfvis# show running-config monitor session
monitor session 20
 destination vm-vnic Linux2 0
```

```
 source vm-vnic Linux1 0 both
 source interface GE0-0 both
!
nfvis#

nfvis# show system monitor session
system monitor session 20
 bridge                wan-br
 destination_vlan      ""
 destination_interface vnic11
 source_vlans          ""
 source_rx_interfaces  "vnic10, GE0-0"
 source_tx_interfaces  "vnic10, GE0-0"
 source_all            false
 statistics            "tx_bytes=142660, tx_packets=1380"
nfvis#

nfvis# show monitor session status
NUMBER  STATUS
------------------------
20      CREATE_SUCCESS
```

Example: SPAN Session Traffic on a LAN SRIOV

The following example shows how to configure all traffic coming in or going out on an SRIOV interface (VF0). It is also mirrored to VF1.

Note This scenario is applicable only to the Cisco ENCS.

VM deployment for VF-VF scenario:

CentOS_SRIOV, C3, and C5 are CentOS VMs with SRIOV support.

- CentOS_SRIOV: vnic0: wan-net/vnic1: LAN-SRIOV-1 (192.168.1.36)
- C3: vnic0: LAN-SRIOV3 (192.168.1.3)
- C5: vnic0: LAN-SRIOV5 (192.168.1.5)

SPAN destination and source:

- SPAN destination: CentOS_SRIOV (vnic0: wan-net/vnic1: LAN-SRIOV-1)
- SPAN source: C3 (vnic0: LAN-SRIOV-3); traffic in both directions (rx, tx)
- Ping target: C5 (vnic0: LAN-SRIOV-5)

```
nfvis# show running-config monitor session
monitor session 6
 destination vm-vnic CentOS_SRIOV 1
 source vm-vnic C3 0
!
nfvis#

nfvis# show system monitor session
system monitor session 6
```

```
 bridge                ""
 destination_vlan      ""
 destination_interface LAN-SRIOV-1
 source_vlans          ""
 source_rx_interfaces  LAN-SRIOV-3
 source_tx_interfaces  LAN-SRIOV-3
 source_all            ""
 statistics            ""
nfvis#

nfvis# show monitor session status
NUMBER  STATUS
------------------------
6       CREATE_SUCCESS
```

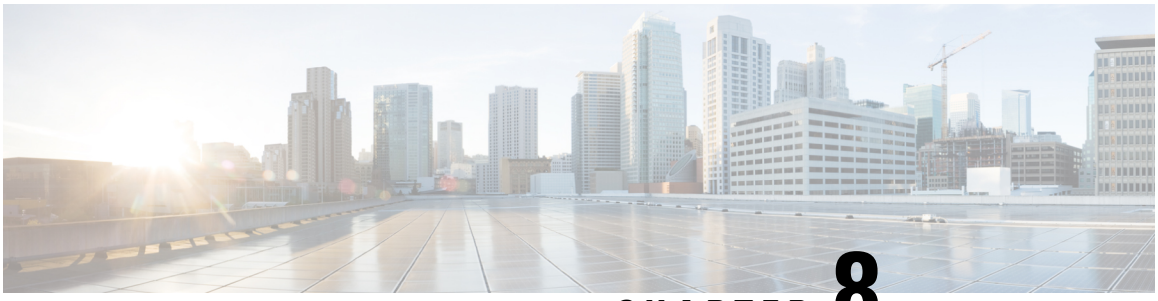
Example: SPAN Session Traffic on a VLAN

The following example shows how to configure the SPAN session for all traffic entering in VLAN 10 and 11. It is also mirrored to VLAN 20.

```
nfvis# show running-config monitor session
monitor session 11
 bridge lan-br
 destination vlan 20
 source vlan [ 10 11 ]
!

nfvis# show system monitor session
system monitor session 11
 bridge                lan-br
 destination_vlan      20
 destination_interface ""
 source_vlans          "10, 11"
 source_rx_interfaces  ""
 source_tx_interfaces  ""
 source_all            true
 statistics            "tx_bytes=0, tx_packets=0"

nfvis# show monitor session 11
NUMBER  STATUS
------------------------
11      CREATE_SUCCESS
```



CHAPTER 8

Configuring Packet Capture

The Packet Capture feature helps you capture all packets being transmitted and received over physical and virtual network interface controllers (physical port and vNIC) for analysis. These packets are inspected to diagnose and solve network problems. Packets are stored in the*/data/intdatastore/pktcaptures* folder on the host server.

Benefits

- You can customize the configuration to capture specific packets such as Internet Control Message Protocol (ICMP), TCP, UDP, and Address Resolution Protocol (ARP).
- You can specify a time period over which packets are captured. The default is 60 seconds.

To configure packet capture on a physial port:

```
configure terminal
tcpdump port eth0

Output: pcap-location /data/intdatastore/pktcaptures/tcpdump_eth0.pcap
```

To configure packet capture on a vNIC:

```
configure terminal
tcpdump vnic tenant-name admin deployment-name 1489084431 vm-name ROUTER vnic-id 0 time 30

Output: pcap-location /data/intdatastore/pktcaptures/1489084431_ROUTER_vnic0.pcap
```

Types of Errors

Error	Scenario
Port/vnic not found	When non-existing interface is given as input.
File/directory not created	When the system is running out of disk space.
The **tcpdump** command fails	When the system is running out of disk space.

These errors are logged in the *nfvis_config.log*. By default, warnings and errors are logged,

Example: Debug Built-in Switch Issues

To monitor traffic problems related to built-in switch on an internal interface:

The regular traffic flow between int-LAN and GE1/0 is:

GE0-0-- vnic1---(VM)--vnic2--intLAN--GE1/0

The NFVIS portal has the capability to capture packets. In the network diagram, right click on any vertical line and a window pops up where you can select the duration of the capture. The packet capture starts on the selected interface link. At the end of the capture, a file is downloaded to your local machine. SPAN sessions are supported on both NFVIS host and the built-in switch.

The following is an example of SPAN in built-in switch:

1. From NFVIS system shell-access, get the password which can be used later.

```
cd /opt/switch-confd/
python decrypt_switch.py

<it will print out a string, it will be the password you need to use later>
8H7)gR348V4Byq4mwjiNt
```

2. From Cisco IMC complete the challenge-response authentication:

```
#connect debug-shell
#sldp
login <hit return>
it will print out the challenge string
enter the  respond string
# switch-con ge
user-name:cisco
password: <enter the sring we get from nfvis system shell>

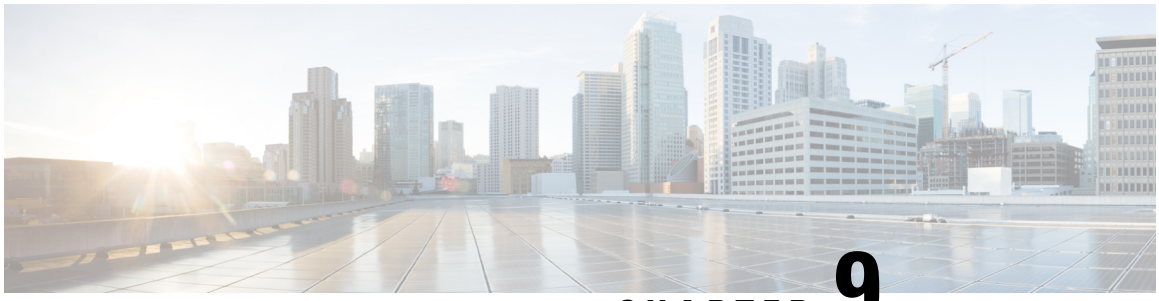
User Name:cisco
Password:**********.   <this is the password you get from step 1 above>
```

3. To configure SPAN specify the source and distribution interface and direction of the packet flow. For example, if you want to mirror XG2 output packet to Ge0, connect an external packet capture tool in GE1/0 and you will see all packets flow from internal XG2. In the following example, the traffic between int_LAN and GE1/0 go through internal interface XE2 and traffic for XE2 interface is monitored:

```
nfvis(config)#monitor session 1 source interface XG 2 out
nfvis(config)#monitor session 1 destination interface GigabitEthernet 0
remember to unconfig it once you finish debuging.
nfvis(config)#no monitor session 1 destination
nfvis(config)#no monitor session 1 source interface XG 2
```

Packet Capture APIs and Commands

APIs	Commands
• /api/operations/packet-capture/tcpdump	• tcpdump port • tcpdump vnic



CHAPTER 9

VM Image Packaging

VM Image Packaging is a tool for converting qcow2 and img images into a tar.gz format with additional properties and profiles. VM image packaging can be done in two ways:

- VM Image Packaging Utility: This is an enhanced packaging process that allows the VM owner to run the **nfvpt.py** utility as a command with a combination of parameters to package the VM.
- Standard Image Packaging: This is a manual process in which a raw disk image (qcow2, img) is packaged along with the image properties file and bootstrap files (if needed) into a TAR archive file.

- VM Image Packaging Utility, on page 87
- Standard VM Image Packaging, on page 96
- Appendix, on page 96

VM Image Packaging Utility

A VM image package is a TAR archive file with the root disk image and other descriptor files. This packaging method simplifies the process of a VM image registration and deployment. The attributes specified for the image enable resource requirement specification, creation of VM profiles, and a host of other properties for the VM.

The Cisco Enterprise NFVIS VM image packaging tool, nfvpt.py, helps VM owners package their VMs. The tool takes one or more qcow2 images (raw disk file) as the input file along with VM specific properties, bootstrap configuration files (if any), and generates a compressed TAR file.

Contents

The VM image packaging utility contains the following:

- nfvpt.py—It is a python based packaging tool that bundles the VM raw disk image/s along with VM specific properties.
- image_properties_template.xml—This is the template file for the VM image properties file, and has the parameters with default values. If the user provides new values to these parameters while creating the VM package, the default values get replaced with the user-defined values.
- tool-usage-examples.txt—This file contains examples on how to use the image packaging utility to package a VM image.

- scaffold_template.cfg— This is a json format configuration file that specifies the meta-data for a VM to be packaged.

Usage

VM Image Packaging Tool with a UI

Table 3: Feature History Table

Feature	Description
VM Image Packaging Tool with a UI	Starting from Cisco NFVIS Release 4.9, a new VM packaging tool is introduced. This VM image packaging tool helps in generating a full VM package. In addition, scaffolding and repacking can be performed using the UI elements in the tool. **Note** The VM image packaging tool is not a part of Cisco NFVIS. The tool needs to be downloaded and installed in your system.

Download VM Image Packaging Tool

Use the following steps to download and install the VM image packaging tool:

1. Navigate to the Software Download page of NFVIS Release 4.9.1.
2. Click the download icon next to **VM Packaging tool for NFVIS** file.
3. Unzip the **CISCO-NFVIS-vmpackagingtool-4.9.1.tar** file.
4. The unzipped folder contains **nfvpt_ui_README.txt** and **nfvpt_ui.tar** files.
5. Follow the instructions in the **nfvpt_ui_README.txt** file to install the tool and perform VM image packaging.

VM Packaging using the CLI

To get the list of parameters that can be included in the command, and to get an explanation of each of the parameters, run the **help** command for the tool.

Starting from NFVIS 4.6 and later releases, use the following **help** command:

```
nfvpt.py --help

Cisco NFVIS Packaging Tool

usage: nfvpt.py packaging-v2 [-h] [--config CONFIG] [-o PACKAGE_FILENAME]
                             [--log LOG] [-d PACKAGE_OUTPUT_DIR]
                             [--legacy_help LEGACY_HELP]

Version: 4.6+ Cisco NFVIS: VNF image packaging utility

optional arguments:
```

```
  -h, --help            show this help message and exit
  --config CONFIG       [REQUIRED] config file template for creation of the
                        package.
  -o PACKAGE_FILENAME, --package_filename PACKAGE_FILENAME
                        [REQUIRED] file name for the target VNF package name-
                        default is root disk image name with extension .tar.gz
  --log LOG             [REQUIRED] file name for logfiles
  -d PACKAGE_OUTPUT_DIR, --package_output_dir PACKAGE_OUTPUT_DIR
                        destination path where package is required
  --legacy_help LEGACY_HELP
                        show the help message for the options available in
                        legacy packaging version and exit
```



Note You can also use the **nfvpt.py --legacy_help** command to get information about the previous version of the tool command help section.

For NFVIS 4.5 and earlier releases, use the following **help** command:

```
nfvpt.py --help
optional arguments:
  -h, --help            show this help message and exit
  --json JSON           Provide JSON input for bootstrap variables; mutually
                        exclusive with custom and bootstrap configs
  --newjson NEWJSON     Provide JSON input for bootstrap variables; mutually
                        exclusive with custom and bootstrap configs
  --log_dir LOG_DIR     Log Directory to for logfiles
  --multi_use           Add options for use in multiple use-cases
  --console_type_serial {true,false}
                        Attach the console serial to the VM; default is false;
                        --console_type_serial=true/false;
  --root_file_disk_bus {virtio,ide}
                        root disk file type: --root_file_disk_bus=virtio/ide;
                        default is virtio
  --virtual_interface_model {rtl8139}
                        --virtual_interface_model=rtl8139; default is none
  --thick_disk_provisioning {true,false}
                        --thick_disk_provisioning=true; default is false
  --eager_zero {true,false}
                        --eager_zero=true; default is false
  --nocloud {true,false}
                        --nocloud=true/false; default is false
  --bootstrap_cloud_init_bus_type {ide,virtio}
                        --bootstrap_cloud_init_bus_type=virtio; default is ide
  --bootstrap_cloud_init_drive_type {cdrom,disk}
                        --bootstrap_cloud_init_drive_type=disk; default is
                        cdrom
  --bootstrap BOOTSTRAP
                        Every bootstrap file should be a different option Non
                        HA format: --bootstrap
                        <mountpoint>:<file1>,<mountpoint>:<file2>... See
                        usage.txt for more details HA format for SDWAN
                        NetworkHub: --bootstrap mount_point:<value>,file:<file
                        2mount>[,<attrib>:<value>] mount_point:<value> and
                        file:<file2mount> are mandatory followed by one or
                        more attributes in the format <attrib>:<value>
  --interface_hot_add {true,false}
                        VM supports interface add without power off. Default
                        is set to true; --interface_hot_add=true/false
  --interface_hot_delete {true,false}
                        VM supports interface delete without power off.
```

```
                        Default is set to false;
                        --interface_hot_delete=true/false
  -v, --verbose         verbose
  -q, --quiet           quiet
  --no_compress         creates tar file without compressing the input files
  --cleanup             deletes all the input and configuration files upon tar
                        file created
  --tablet {true,false}
                        : Add input device of type tablet --tablet=true/false;
  --ha_package          enable HA packaging
  --mgmt_vnic MGMT_VNIC
                        VM management interface identifier
  --pack_dir <DIR> PACK
                        package all files in directory

Required:
  -o PACKAGE_FILENAME, --package_filename PACKAGE_FILENAME
                        [REQUIRED] file name for the target VNF package name-
                        default is root disk image name with extension .tar.gz
  -i ROOT_DISK_IMAGE, --root_disk_image ROOT_DISK_IMAGE
                        [REQUIRED] List of root disk images to be bundled
                        example: --root_disk_image isrv.qcow2;
                        --root_disk_image isrv1.qcow2,isrv2.qcow2
  --prop_template PROP_TEMPLATE
                        image properties template file name including path
                        default path is the current dir of the tool and name
                        is image_properties_template.xml if the user doesn't
                        input this option example: --prop_template
                        /usr/bin/image_properties_template.xml
  -t VNF_TYPE, --vnf_type VNF_TYPE
                        [REQUIRED] VNF type, e.g. ROUTER, FIREWALL, vWAAS,
                        vWLC, and OTHER
  -n NAME, --vnf_name NAME
                        [REQUIRED] Name of the VNF image
  -r VNF_VERSION, --vnf_version VNF_VERSION
                        [REQUIRED] VNF version, e.g. --vnf_version 1.0 or
                        --vnf_version 0.9
  --app_vendor APP_VENDOR
                        Application Vendor e.g. Cisco, Juniper etc
  --monitored {true,false}
                        [REQUIRED] Monitored VNF: --monitored=true/false;
  --optimize {true,false}
                        [REQUIRED] optimized VM: --optimize=true/false;

HA options:
  --ha_capable
  --ha_vnic HA_VNIC     VM HA vnic
  --ha_vnic_count HA_VNIC_COUNT
                        Number of ha_vnics

Resources:
  Resources: min and max - vCPU, memory and disk

  --min_vcpu VCPU_MIN   min #vCPU : min number of vCPU supported by VM
                        example:--min_vcpu 2
  --max_vcpu VCPU_MAX   max #vCPU : max number if vCPU required for VM
                        example:--max_vcpu 4
  --min_mem MEMORY_MB_MIN
                        min mem : min mem in MB required for VM
                        example:--min_mem 1024
  --max_mem MEMORY_MB_MAX
                        max mem : max mem in MB required for VM
                        example:--max_mem 4196
  --min_disk ROOT_DISK_GB_MIN
```

```
                        min disk : min disk in GB required for VM
                        example:--min_disk 8
  --max_disk ROOT_DISK_GB_MAX
                        max disk : max disk in GB required for VM
                        example:--max_disk 8
  --vnic_max VNIC_MAX   max number of Vnics allowed for VM example:--vnic_max
                        8
  --vnic_names VNIC_NAMES
                        list of vnic number to name mapping in format
                        number:name example --vnic_names
                        1:GigabitEthernet2,2:GigabitEthernet4

Profile Options:
  --profile PROFILE     enter the profile name, profile description, no of
                        vCPU required, min memory required in MB, min disk
                        space required in MB, example: --profile
                        profile1,"This is profile 1",2,2048,4096 --profile
                        profile2,"This is profile 2",4,4096,4096
  --default_profile DEFAULT_PROFILE
                        default profile

Driver Support Options:
  --sriov {true,false}  Enable/Disable SRIOV support: --sriov=true/false;
                        default is false
  --sriov_list SRIOV_DRIVER_LIST
                        list of SRIOV drivers example: --sriov_list
                        igb,igbvf,i40evf
  --pcie {true,false}   Not supported
  --pcie_list PCIE_DRIVER_LIST
                        Not supported

Privilege/Priority Options:
  --privileged {true,false}
                        Not supported

Custom Properties:
  --custom CUSTOM       custom properties format: --custom ["propattr_"<attr>:
                        <value>],key:<value>,[keyattr_<attr>:<value>],type:<va
                        lue>,val<N>:<value>,[val<N>attr_<attr>:<value>] Allows
                        specification of custom properties: 0 or more
                        propattr_<attr>:<value> pairs - 'propattr' is a
                        keyword and used to specify property attributes
                        key:<value> pairs 0 or more keyattr_<attr>:value pairs
                        - 'keyattr' is a keyword and is used to specify key
                        attributes type:<value> pair - type of value
                        valN:<value> pair - val1:value,val2:value etc 0 or
                        more valNattr_<attr>:<value> pairs - 'val<N>attr' is
                        an attribute for val<N> See usage_examples.txt
```

The table lists the parameters that can be passed to the **nfvpt.py** command.

Parameter	Mandatory/Optional	Description
version	Not applicable	Show program's version number and exit.
help	Not applicable	Show this help message and exit.
package_file_name	Mandatory	File name for the target VNF package. The default is the root disk image name with extension *.tar.gz*.

Parameter	Mandatory/Optional	Description
disk_img_names	Mandatory	List of root disk images to be bundled. Only the qcow2 images are supported.
img_name	Mandatory	Name of the VNF image.
vnf_type	Mandatory	VNF type Supported types are: ROUTER, FIREWALL, vWAAS, vWLC, and OTHER.
vnf_version	Mandatory	VNF version
monitored	Mandatory	VM health monitoring for those VMs that can be bootstrapped Options are: true/false Monitoring timeout period for a monitored VM is 600 seconds by default
optimize	Mandatory	Optimized VM Options are: true/false
virtual_interface_model	Optional	Default is none.
thick_disk_provisioning	Optional	Default is false.
eager_zero	Optional	Default is false.
bootstrap_cloud_init_bus_type	Optional	Default is IDE.
bootstrap_cloud_init_drive_type	Optional	Mounts the day0 configuration file as disk Default is CD-ROM.
bootstrap	Optional	Bootstrap files for VNF. Two parameters are required in the format of dst:src; dst filename including path has to match exactly to what the VM expects; up to 20 bootstrap files are accepted. For example: --bootstrap ovf-env.xml for ISRv and --bootstrap day0-config for ASAv.
min_vcpu	Optional	Minimum number of vCPUs supported by the VM. The default is 1.
max_vcpu	Optional	Maximum number of vCPUs required for the VM. The default is 8.

Parameter	Mandatory/Optional	Description
min_mem	Optional	Minimum memory in MB required for the VM. The default is 4 GB.
max_mem	Optional	Maximum memory in MB required for the VM. Physical memory: 2 GB The default is 8 GB.
min_disk	Optional	Minimum disk in GB required for the VM. The default is 8 GB.
max_disk	Optional	Maximum disk in GB required for the VM. Available disks are SSD and HDD: 15 GB The default is 16 GB
vnic_max	Optional	Maximum number of VNICs allowed for the VM. The default is 8.
profile	Optional	The profile name, profile description, number of vCPUs required, minimum memory required in MB and minimum disk space required in MB.
default_profile	Optional	The default profile.
sriov	Optional	Enable or disable SRIOV support. The default is false.
sriov_list	Optional	List of SRIOV drivers.
pcie	Optional	Not supported.
pcie_list	Optional	Not supported.
privileged	Optional	Not supported.
custom	Optional	Custom properties to be supported and/or passed to the bootstrap configuration with tokenized variables. This is only used for the local portal to display options for the user to choose while deploying.
pack_dir	Optional	package all files in directory

NFVIS Specific Enhancements

Note Use pack_dir option if the *.tar.gz already exists and you want to modify the bootstrap configuration file or image_properties.xml manually.

The following parameters are added as part of the NFVIS specific enhancements:

```
 --pack_dir <DIR> PACK

                        package all files in directory

Resources:

  --vnic_names VNIC_NAMES

                        list of vnic number to name mapping in format

                        number:name example --vnic_names

                        1:GigabitEthernet2,2:GigabitEthernet4
```

Usage

Follow the steps to change a single line in day-0 configuration file or add a single option in image_properties.xml:

1. Get the working VM packaging image - isrv*.tar.gz.
2. Extract the contents - tar -xvf isrv*.tar.gz.
3. Modify the file contents as required.
4. nfvpt.py --pack_dir current-working-dir-with-files -i isrv.qcow2 -o isrv.tar.gz

VM Packaging Utility Usage Examples

Given below are the contents of the file *tool-usage-examples.txt*:

Example 1: Usage for TinyLinux

```
nfvpt.py -o TinyLinux -i TinyLinux.qcow2 -n TinyLinux -t linux -r 1.0 --monitored false
--min_vcpu 1 --max_vcpu 2 --min_mem 1024 --max_mem 1024 --min_disk 1 --max_disk 2
--vnic_max 1 --optimize false
```

Example 2: Usage for ASAv

Note The bootstrap filename has to be *day0-config*. This cannot be modified as ASAv looks for the exact filename.

```
nfvpt.py -o asav961-201 -i asav961-201.qcow2 -n ASAv -t firewall -r 961-201 --monitored
true --bootstrap day0-config:filename1
```

```
--min_vcpu 1 --max_vcpu 4 --min_mem 1024 --max_mem 8192 --min_disk 8 --max_disk 16 --vnic_max
 8 --optimize true
--profile ASAv5,"ASAv5 profile",1,1024,8192 --profile ASAv10,"ASAv10 profile",1,4096,8192
--profile ASAv30,"ASAv30 profile",4,8192,16384
--default_profile ASAv5
```

Example 3: Usage for ISRv



Note The bootstrap filename has to be *ovf-env.xml* . This cannot be modified as ISRv looks for the exact filename.

```
nfvpt.py -o isrv.16.03.01 -i isrv-universalk9.16.03.01.qcow2 -n ISRv.16.03.01 -t ROUTER -r
 16.03.01 --monitored true --privileged true
--bootstrap ovf-env.xml:file1,ios-xe.txt:file2 --min_vcpu 2 --max_vcpu 8 --min_mem 4096
--max_mem 8192 --min_disk 8 --max_disk 8
--vnic_max 8 --optimize true --profile ISRv-small,"ISRv small profile",2,4096,8192 --profile
 ISRv-medium,"ISRv medium profile",4,4096,8192
--default_profile ISRv-small --sriov_list igb,igbvf,i40evf --custom tech_package,ax
```

Example 4: Usage for a third party VM with config drive (ISO) mounted at specific path on the VM:

```
nfvpt.py -o test.1.0 -i test-1.0.qcow2 -n TEST -t OTHER -r 1.0 --monitored true --privileged
 true
--bootstrap /:bootstrap.xml,/license/lic.txt:license.txt --min_vcpu 2 --max_vcpu 8 --min_mem
 4096 --max_mem 8192
--min_disk 8 --max_disk 8 --vnic_max 8 --optimize true --profile small,"small
profile",2,4096,8192
--profile medium,"medium profile",4,4096,8192 --default_profile small
```

In this case, *test.1.0.pkg : bootstrap.xml* gets mounted as *bootstrap.xml* at the root, and the *license.txt* gets mounted as */license/lic.txt*.

Example 5: Usage for Palo Alto Firewall

```
nfvpt.py -o PA_L3_HA -i PA-VM-KVM-8.0.5.qcow2 --json d.json -t firewall -n "PA FIREWALL"
-r 8.0.5 --app_vendor PA --monitor true --ha_package
```

Example 6: Usage for Asav

```
nfvpt.py -i foo.qcow2 -o asav.tar.gz --json pa1.json --app_vendor cisco -t firewall -r 10
--optimize true -n asav --monitored true --ha_package —ha_capable
```

Example 7: Usage for csr

```
nfvpt.py --ha_package --pack_dir /data/intdatastore -i csr1000v-universalk9.16.09.01.qcow2
  -o csr1000v-universalk9.16.09.01-ha.tar.gz
```

Usage examples for NFVIS 4.6 and later releases:

Example 1: Usage for configuration template based packaging

```
nfvpt.py packaging-v2 --config scaffold_template.cfg --log logger -o example.tar.gz -d dest
```

Example 2: Usage for adding the certifiate file to the Day 0 file (only supported for C8000v and vEdge)

```
nfvpt.py --modify_package repackage --file_path /usr/data/example1.tar.gz --cert
BranchCert.pem --package_output example2.tar.gz --mod_version X.X --root_image
/Users/data/example.qcow2
```

Example 3: Usage for adding and replacing the qcow2 file and modifying the version

```
nfvpt.py --modify_package repackage --file_path /usr/data/example1.tar.gz --package_output
 example2.tar.gz --mod_version X.X --root_image /Users/data/example.qcow2
```

Standard VM Image Packaging

The standard VM packaging is based on the Open Virtualization Format (OVF) packaging standard, in which a single file is distributed in open virtualization appliance (OVA) format. The VM image is shared using a TAR archive file with the root disk image and descriptor files.

Note Cisco Enterprise NFVIS supports VM packaging in *.tar.gz* (compressed form of OVA) format. Ensure that all supported third party VM images are available in the supported format.

Generating a VM Package

Package files are provided for Cisco ISRv, Cisco ASAv, and tiny Linux and Windows server 2000. Vendors are responsible for packaging all third party VMs in the supported format.

1. Create a VM qcow2 image.
2. Create an *image_properties.xml* file with the VM properties. Ensure that you add all mandatory fields. Include the profiles supported for the VM in this file, and select one default profile. If you do not want to monitor the VM bootup, make the bootup time as -1.
3. Create *bootstrap-config* or *day0-config*, if any bootstrap configuration is required for the VM. If the bootstrap configuration requires inputs from the user, use the tokens in the xml or text file. These tokens are populated during the VM deployment with the provided data.

Note A VM deployment may fail, if there are tokens in the configuration, and the user does not provide the token values in the deployment payload.

4. Create a *package.mf* file, which lists all the files to be bundled into the *.tar.gz* file along with checksums.
5. Generate the packaging file using "tar -cvzf ova_file_name list_of_files_to_be_bundled".

   For example, *tar -cvzf isrv.tar.gz isrv-universalk9.03.16.02.S.155-3.S1a-ext-serial.qcow2 image_properties.xml isr_ovf_env.xml package.mf*.

Appendix

VM Image Package Files

The table lists the contents of the VM package that are generated using the packaging tool:

Table 4: VM Image Package Files

File	Description	Mandatory/Optional
Package Manifest (package.mf)	Lists the files in the package and the expected checksum for the files.	Mandatory
VM image properties (vmname_properties.xml)	XML file with resources and features supported by the VM	Mandatory
VM image (vmname.qcow2)	Image file of the VM. Multiple images are supported. One root_disk image file is mandatory.	Mandatory
Bootstrap (bootstrap_file)	Optional	Bootstrap files for VNF. Two parameters are required in the format of dst:src; dst filename including path has to match exactly to what the VM expects; up to 20 bootstrap files are accepted. For example: --bootstrap ovf-env.xml for ISRv and --bootstrap day0-config for ASAv.

Package Manifest File

The package manifest XML file provides a list of the files in the package with their names and their expected checksum. SHA1 algorithm (sha1sum) is used to calculate the checksum. This is a mandatory file to be bundled in the VM package. The manifest file must be named as *package.mf*.

Table 5: Package Manifest File Details

Property Name	Description	Property Tag	Mandatory/Optional
File information	XML tree with details of file name, file type, and expected checksum. The root_image and image_properties files are required.	<file_info>	Mandatory
File name	Name of the file	<name>	Mandatory
File type	Describes the file type. Supported types: • root_image • image_properties • bootstrap_config_file • ephemeral_disk1_image • ephemeral_disk2_image	<type>	Mandatory

Expected checksum	The calculated SHA1 checksum to be validated.	<sha1_checksum>	Mandatory

Bootstrap Configuration File

The bootstrap configuration file is an XML or a text file, and contains properties specific to a VM and the environment. Properties can have tokens, which can be populated during deployment time from the deployment payload.

VM Image Properties File

This XML file provides information about the resources supported or required for the VM operation. All mandatory parameters have to be defined. It also supports custom attributes. This is a mandatory file to be bundled in the VM package. The VM package supports up to 10 disks to be bundled into the package.

Table 6: VM Image Properties File Details

Property Name	Description	Property Tag	Possible Values	Mandatory/Optional
VNF Type	VM functionality provided. Router and firewall are predefined types.	<vnf_type>	Router, firewall, Windows, Linux, and custom_type	Mandatory
Name	Name associated with the VM packaging. This name is referenced for VM deployment.	<name>	Any	Mandatory
Version	Version of the package	<version>	Any	Mandatory
Boot-up time	Boot-up time (in seconds) of the VNF before it can be reachable via ping.	<bootup_time>	Any in seconds, (-1) to not monitor boot-up	Mandatory
Root Disk Image Bus	Root image disk bus	<root_file_disk_bus>	virtio, scsi, and ide	Mandatory
Disk-1 bus type	Additional disk1 image disk bus	<disk_1_file_disk_bus>	virtio, scsi, and ide	Optional
Disk-2 bus type	Disk2 image disk bus	<disk_2_file_disk_bus>	virtio, scsi, and ide	Optional
Disk-10 bus type	Disk10 image disk bus	<disk_10_file_disk_bus>	virtio, scsi, and ide	Optional

Root Disk Image format	Root image disk format	<root_image_disk_format>	qcow2 and raw	Mandatory
Disk-1 Image format	Additional disk 1 image format	<disk_1_image_format>	qcow2 and raw	Optional
Disk-2 Image format	Disk 2 image format	<disk_2_image_format>	qcow2 and raw	Optional
Disk-10 Image format	Disk 10 image format	<disk_10_image_format>	qcow2 and raw	Optional
Serial Console	Serial console supported	<console_type_serial>	true, false	Optional
Minimum vCPU	Minimum vCPUs required for a VM operation	<vcpu_min>		Mandatory
Maximum vCPU	Maximum vCPUs supported by a VM	<vcpu_max>		Mandatory
Minimum memory	Minimum memory in MB required for VM operation	<memory_mb_min>		Mandatory
Maximum memory	Maximum memory in MB supported by a VM	<memory_mb_max>		Mandatory
Minimum root disk size	Minimum disk size in GB required for VM operation	<root_disk_gb_min>		Optional
Maximum root disk size	Maximum disk size in GB supported by a VM	<root_disk_gb_max>		Optional
Maximum vNICs	Maximum number of vNICs supported by a VM	<vnic_max>		Mandatory
SRIOV support	SRIOV supported by VM interfaces. This should have a list of supported NIC device drivers.	<sriov_supported>	true, false	Optional

SRIOV driver list	List of drivers to enable SRIOV support	< sriov_driver_list>		Optional
PCI passthru support	PCI passthru support by VM interfaces	<pcie_supported>	true, false	Optional
PCIE driver list	List of VNICS to enable PCI passthru support	< pcie _driver_list>		Optional
bootstrap_cloud_init_drive_type	Mounts day0 config file as disk (default is CD-ROM)	<bootstrap_cloud_init_drive_type>	disk, cdrom	Optional
bootstrap_cloud_init_bus_type	Default is IDE	<bootstrap_cloud_init_bus_type>	virtio, ide	Optional
BOOTSTRAP	Bootstrap files for the VNF. Two parameters are required in the format of dst:src; dst filename including path has to match exactly to what the VM expects; up to 20 bootstrap files are accepted. For example: --bootstrap ovf-env.xml for ISRv and --bootstrap day0-config for ASAv	< bootstrap_file>	File name of the bootstrap file	Optional

Custom properties	List of properties can be defined within the custom_property tree. (Example: For ISRv, the technology packages are listed in this block.) If the Cisco Enterprise NFV portal is used to deploy the VM, the portal prompts you for inputs for custom properties fields, and can pass the values to the bootstrap configuration.	<custom_property>		Optional
Profiles for VM deployment	List of VM deployment profiles. Minimum one profile is required	<profiles>		Optional
Default profile	The default profile is used when no profile is specified during deployment.	<default_profile>		Optional
Monitoring Support	A VM supports monitoring to detect failures.	<monitoring_supported>	true, false	Mandatory
Monitoring Method	A method to monitor a VM. Currently, only ICMP ping is supported.	<monitoring_methods>	ICMPPing	Mandatory if monitoring is true
Low latency	If a VM's low latency (for example, router and firewall) gets dedicated resource (CPU) allocation. Otherwise, shared resources are used.	<low_latency>	true, false	Mandatory

Privileged-VM	Allows special features like promiscuous mode and snooping . By default, it is false.	<privileged_vm>	true, false	Optional
Virtual interface model		<virtual_interface_model>		Optional
Thick disk provisioning	By default, it is false.	<thick_disk_provisioning>	true, false	Optional
Profile for VM deployment	A profile defines the resources required for VM deployment. This profile is referenced during VM deployment.	<profile>		Optional
Name	Profile name	<name>	Any	Mandatory
Description	Description of the profile	<description>	Any	Mandatory
vCPU	vCPU number in a profile	<vcpus>		Mandatory
Memory	Memory - MB in profile	<memory_mb>		Mandatory
Root Disk Size	Disk size - MB in profile .	<root_disk_mb>		Mandatory
VNIC Offload	List of properties that can be set for vnic offload	<vnic_offload>		Optional
Generic Segmentation Offload	Turn generic segmentation offload on or off	<generic_segmentation_offload> (parent: <vnic_offload>)	on, off	Optional
Generic Receive Offload	Turn generic receive offload on or off	<generic_receive_offload> (parent: <vnic_offload>)	on, off	Optional
RX Checksumming	Turn RX checksumming on or off	<rx_checksumming> (parent: <vnic_offload>)	on, off	Optional
TX Checksumming	Turn TX checksumming on or off	<tx_checksumming> (parent: <vnic_offload>)	on, off	Optional

TCP Segmentation Offload	Turn TCP segmentation offload on or off	<tcp_segmentation_offload> (parent: <vnic_offload>)	on, off	Optional



Note A virtual console is supported by default. Specify the root disk size as zero for multiple disks (for example, vWaas deployment) as the system does not support populating multiple disk sizes. Actual disk sizes are calculated from the root_disk files.

Example: Package.mf

```
** sha1sum - for calculating checksum
<PackageContents>
  <File_Info>
    <name>ISRv_serial_3.16.02.qcow2</name>
    <type>root_image</type>
    <sha1_checksum>93de73ee3531f74fddf99377972357a8a0eac7b</sha1_checksum>
  </File_Info>
  <File_Info>
    <name>image_properties.xml</name>
    <type>image_properties</type>
    <sha1_checksum>c5bb6a9c5e8455b8698f49a489af3082c1d9e0a9</sha1_checksum>
</File_Info>
  <File_Info>
    <name>ISRv_ovf_env.xml</name>
    <type> bootstrap_file_1</type>
    <sha1_checksum>c5bb6a9c5e8455b8698f49a489af3082c1d9e0a9</sha1_checksum>
   </File_Info>
   <File_Info>
     <name>ISRv_disk1_image.qcow2</name>
     <type>ephemeral_disk1_image</type>
     <sha1_checksum>aac24513098ec6c2f0be5d595cd585f6a3bd9868</sha1_checksum>
   </File_Info>
</PackageContents>
```

Example: Image Properties

```
<?xml version="1.0" encoding="UTF-8"?>
<image_properties>
    <vnf_type>ROUTER</vnf_type>
    <name>isrv-universalk9</name>
    <version>03.16.02</version>
    <bootup_time>600</ bootup_time >
    <root_file_disk_bus>virtio</root_file_disk_bus>
    <root_image_disk_format>qcow2</root_image_disk_format>
    <vcpu_min>1</vcpu_min>
    <vcpu_max>8</vcpu_max>
    <memory_mb_min>4096</memory_mb_min>
    <memory_mb_max>8192</memory_mb_max>
    <vnic_max>8</vnic_max>
    <root_disk_gb_min>8</root_disk_gb_min>
    <root_disk_gb_max>8</root_disk_gb_max>
    <console_type_serial>true</console_type_serial>
    <sriov_supported>true</sriov_supported>
    <sriov_driver_list>igb</sriov_driver_list>
```

```
    <sriov_driver_list>igbvf</sriov_driver_list>
    <sriov_driver_list>i40evf</sriov_driver_list>
    <pcie_supported>true</pcie_supported>
    <pcie _driver_list> igb </pcie_driver_list>
    <pcie _driver_list> igbvf</pcie_driver_list>
    <pcie _driver_list> i40evf</pcie_driver_list>
    <bootstrap_file_1> ovf-env.xml </bootstrap_file_1>
    <monitoring_supported>true</monitoring_supported>
    <monitoring_methods>ICMPPing</monitoring_methods>
    <low_latency>true</low_latency>
    <privileged_vm>true</privileged_vm>
    <cdrom>true</cdrom>
    <custom_property>
        <tech_package>ax</tech_package>
        <tech_package>sec</tech_package>
        <tech_package>ipbase</tech_package>
        <tech_package>appx</tech_package>
    </custom_property>
    <profiles>
         <profile>
            <name>ISRv1kv-small</name>
            <description>ISRv upto 50MBPS performance</description>
            <vcpus>1</vcpus>
            <memory_mb>4096</memory_mb>
            <root_disk_mb>8</root_disk_mb>
         </profile>
         <profile>
            <name>ISRv1kv-medium</name>
            <description>ISRv upto 250MBPS performance</description>
            <vcpus>2</vcpus>
            <memory_mb>4096</memory_mb>
            <root_disk_mb>8</root_disk_mb>
         </profile>
    </profiles>
    <default_profile>small</default_profile>
 </image_properties>
```

Example: Bootstrap Configuration File

```
<?xml version="1.0" encoding="UTF-8"?>
<Environment
xmlns:oe="http://schemas.dmtf.org/ovf/environment/1">
  <PropertySection>
    <Property oe:key="com.cisco.ISRv.config-version.1" oe:value="1.0"/>
    <Property oe:key="com.cisco.isrv.enable-ssh-server.1" oe:value="True"/>
    <Property oe:key="com.cisco.isrv.login-password.1" oe:value="admin"/>
    <Property oe:key="com.cisco.isrv.login-username.1" oe:value="lab"/>
    <Property oe:key="com.cisco.isrv.mgmt-interface.1" oe:value="GigabitEthernet1"/>
   <Property oe:key="com.cisco.isrv.mgmt-ipv4-addr.1" oe:value="${NICID_0_IP_ADDRESS}/24"/>

    <Property oe:key="com.cisco.isrv.mgmt-ipv4-network.1" oe:value=""/>
    <Property oe:key="com.cisco.isrv.license.1" oe:value="${TECH_PACKAGE}"/>
   <Property oe:key="com.cisco.isrv.ios-config-0001" oe:value="vrf definition Mgmt-intf"/>

    <Property oe:key="com.cisco.isrv.ios-config-0002" oe:value="address-family ipv4"/>
    <Property oe:key="com.cisco.isrv.ios-config-0003" oe:value="exit-address-family"/>
    <Property oe:key="com.cisco.isrv.ios-config-0004" oe:value="address-family ipv6"/>
    <Property oe:key="com.cisco.isrv.ios-config-0005" oe:value="exit-address-family"/>
    <Property oe:key="com.cisco.isrv.ios-config-0006" oe:value="exit"/>
   <Property oe:key="com.cisco.isrv.ios-config-0007" oe:value="interface GigabitEthernet1"/>

    <Property oe:key="com.cisco.isrv.ios-config-0008" oe:value="vrf forwarding Mgmt-intf"/>
```

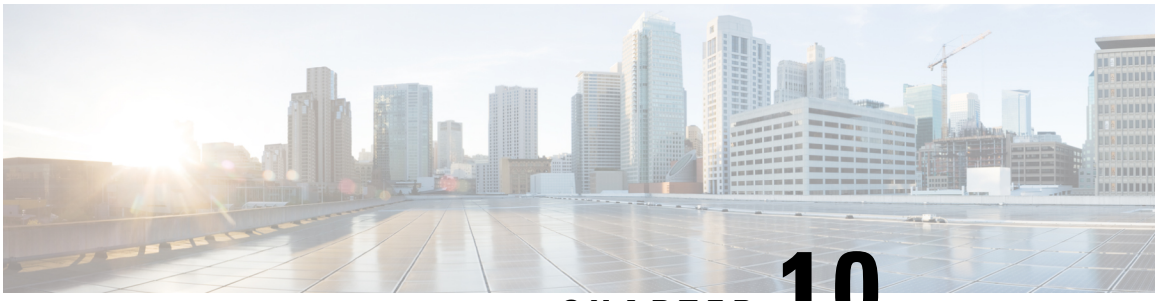
```
    <Property oe:key="com.cisco.isrv.ios-config-0009" oe:value="ip address
${NICID_0_IP_ADDRESS} ${NICID_0_NETMASK}"/>
    <Property oe:key="com.cisco.isrv.ios-config-0010" oe:value="no shut"/>
    <Property oe:key="com.cisco.isrv.ios-config-0011" oe:value="exit"/>
    <Property oe:key="com.cisco.isrv.ios-config-0012" oe:value="ip route vrf Mgmt-intf
0.0.0.0 0.0.0.0 ${NICID_0_GATEWAY}"/>
  </PropertySection>
</Environment>
```

Image Properties Template File

The parameters that go into the image properties file are listed in the code extract below.

```
<?xml version="1.0" encoding="UTF-8"?>
<image_properties>
    <vnf_type>ROUTER</vnf_type>
    <name>TEMPLATE</name>
    <version>1.0</version>
    <bootup_time>600</bootup_time>
    <root_file_disk_bus>virtio</root_file_disk_bus>
    <root_image_disk_format>qcow2</root_image_disk_format>
    <vcpu_min>1</vcpu_min>
    <vcpu_max>8</vcpu_max>
    <memory_mb_min>4096</memory_mb_min>
    <memory_mb_max>8192</memory_mb_max>
    <vnic_max>8</vnic_max>
    <root_disk_gb_min>8</root_disk_gb_min>
    <root_disk_gb_max>16</root_disk_gb_max>
    <console_type_serial>false</console_type_serial>
    <sriov_supported>true</sriov_supported>
    <sriov_driver_list>s1</sriov_driver_list>
    <sriov_driver_list>s2</sriov_driver_list>
    <sriov_driver_list>s3</sriov_driver_list>
    <pcie_supported>false</pcie_supported>
    <monitoring_supported>true</monitoring_supported>
    <monitoring_methods>ICMPPing</monitoring_methods>
    <low_latency>true</low_latency>
    <privileged_vm>false</privileged_vm>
    <cdrom>true</cdrom>
    <bootstrap_file_1>b1.xml</bootstrap_file_1>
    <bootstrap_file_2>b2.txt</bootstrap_file_2>
    <custom_property>
        <key>val</key>
    </custom_property>
    <profiles>
        <profile>
            <name>small</name>
            <description>small</description>
            <vcpus>1</vcpus>
            <memory_mb>1024</memory_mb>
            <root_disk_mb>4096</root_disk_mb>
        </profile>
        <profile>
            <name>medium</name>
            <description>medium</description>
            <vcpus>2</vcpus>
            <memory_mb>4096</memory_mb>
```

```
            <root_disk_mb>8192</root_disk_mb>
        </profile>
    </profiles>
    <default_profile>small</default_profile>
</image_properties>
```



CHAPTER 10

Upgrading Cisco Enterprise NFVIS

The Cisco Enterprise NFVIS upgrade image is available as a `.nfvispkg` file. Currently, downgrade is not supported. All RPM packages in the Cisco Enterprise NFVIS upgrade image are signed to ensure cryptographic integrity and authenticity. In addition, all RPM packages are verified during Cisco Enterprise NFVIS upgrade. For more information about the Image Signing and Verification feature, see Image Signing and Verification, on page 6.

Ensure that you copy the image to the Cisco Enterprise NFVIS server before starting the upgrade process. Always specify the exact path of the image when registering the image. Use the **scp** command to copy the upgrade image from a remote server to your Cisco Enterprise NFVIS server. When using the **scp** command, you must copy the image to the "`/data/intdatastore/uploads`" folder on the Cisco Enterprise NFVIS server. The following is an example on how to use the **scp** command to copy the upgrade image:

```
scp -P 22222 nfvis-351.nfvispkg admin@192.0.2.9:
/data/intdatastore/uploads/nfvis-351.nfvispkg
```

Alternatively, you can upload the image to the Cisco Enterprise NFVIS server using the **System Upgrade** option from the Cisco Enterprise NFVIS portal.

Note When the NFVIS upgrade is in progress, ensure that the system is not powered off. If the system is powered off during the NFVIS upgrade process, the system can become inoperable and you need to reinstall the system.

The upgrade process comprises two tasks:

- Registering the image using the **system upgrade image-name** command.
- Upgrading the image using the **system upgrade apply-image** command.

Registering an Image

To register an image:

```
configure terminal
system upgrade image-name nfvis-351.nfvispkg location /data/intdatastore/uploads
commit
```

Note You must verify the image registration status before upgrading the image using the **system upgrade apply-image** command. The package status must be valid for the registered image.

Verifying the Image Registration

Use the **show system upgrade reg-info** command in the privileged EXEC mode to verify the image registration.

```
nfvis# show system upgrade reg-info
PACKAGE
NAME              LOCATION                                    VERSION         STATUS UPLOAD DATE
-------------------------------------------------------------------------------------------------------------------
nfvis-351.nfvispkg /data/upgrade/register/nfvis-351.nfvispkg 3.6.1-722 Valid
2017-04-25T10:29:58.052347-00:00
```

Upgrading the Registered Image

To upgrade the registered image:

```
configure terminal
system upgrade apply-image nfvis-351.nfvispkg scheduled-time 5
commit
```

Verifying the Upgrade Status

Use the **show system upgrade apply-image** command in the privileged EXEC mode

```
nfvis# show system upgrade apply-image
UPGRADE
NAME   STATUS     FROM     UPGRADE TO
-------------------------------------------------------------------------------------------------------
nfvis-351.nfvispkg SUCCESS 3.5.0 3.5.1
```

The only upgrade supported when BIOS secured boot (UEFI mode) is enabled on ENCS 5400 platform is:

NFVIS 3.8.1 + BIOS 2.5(legacy) --> NFVIS 3.9.1 + BIOS 2.6(legacy)

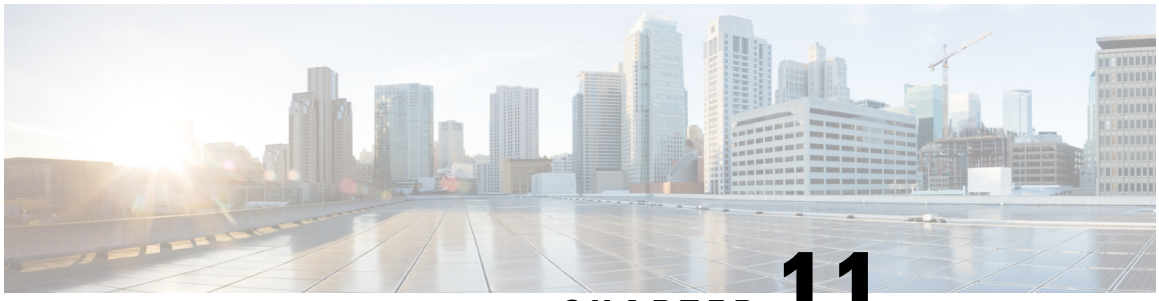
The following upgrade requires re-installation:

NFVIS 3.8.1 + BIOS 2.5(legacy) --> NFVIS 3.9.1 + BIOS 2.6(UEFI)

NFVIS 3.9.1 + BIOS 2.6(legacy) --> NFVIS 3.9.1 + BIOS 2.6(UEFI)

Upgrade APIs and Commands

Upgrade APIs	Upgrade Commands
• /api/config/system/upgrade • /api/config/system/upgrade/image-name • /api/config/system/upgrade/reg-info • /api/config/system/upgrade/apply-image	• system upgrade image-name • system upgrade apply-image • show system upgrade reg-info • show system upgrade apply-image



CHAPTER 11

Configuring vBranch High Availability

The vbranch high availability (HA) solution is a box-to-box HA. It is similar to the traditional branch, which uses physical boxes for routing and other services. This solution uses the Hot Standby Router Protocol (HSRP), a default gateway redundancy (or a first hop redundancy), which allows the network to recover from the failure of the device acting as the default gateway for the LAN side end points (devices). The routing protocols are configured to converge the traffic on the WAN side, when there are failures. So, this solution uses HSRP to provide redundancy for the branch connectivity on the LAN side. The Enhanced Interior Gateway Routing Protocol (EIGRP), Open Shortest Path First (OSPF) routing protocols, and Embedded Event manager (EEM) scripts are configured to converge on the WAN side. The following section explains the redundancy solutions for the branch, with each ENCS having separate active WAN link.

Note You can use this recommended HA design as is, or modify as per the field requirement.

- Prerequisites for vBranch HA, on page 111
- vBranch HA Design and Topology, on page 112
- Enable Virtual NIC Failure Detection with Track Feature , on page 112
- Isolating LAN and Transit Link Traffic for vBranch HA, on page 114
- Packet Flow for vBranch HA, on page 116
- Configuration Examples for vBranch HA, on page 117
- Cisco ENCS Failure Points, on page 120

Prerequisites for vBranch HA

- Cisco ISRv must run HSRP on the LAN facing interface.
- The WAN links are active on both Cisco ENCS1 and Cisco ENCS2. Each of the ENCS WAN link is connected to the WAN network (most cases with two SPs), with two ENCSs in an active-active mode.
- The LAN facing links of both Cisco ENCS devices are connected to an external switch (as an uplink), and all the devices on the LAN segment are also connected to the external switch. There should be no LAN device connecting directly to the Cisco ENCS internal switch.
- A transit link, which is L3 routed, is configured between the Cisco ENCS devices. Since the LAN HSRP makes only one device active, the transit link is used to forward traffic This link is used to forward traffic from the standby ENCS WAN to LAN or LAN to WAN. This link can be back-to-back connected on the ENCS internal switch ports.

- VMs and VNFs on both ENCS devices must be configured identical.

vBranch HA Design and Topology

Physical Devices Connection

Each Cisco ENCS has a WAN traffic connected to the Gigabit Ethernet interface, GE0-0, in this dual-WAN topology.

There are two Cisco ENCS devices namely ENCS1 and ENCS2. There is an external switch connecting one of the LAN ports from each Cisco ENCS. There is a back-to-back connection between ENCS1 and ENCS2 connecting one of the LAN ports from each Cisco ENCS. The WAN port from each Cisco ENCS is connected to the service providers network.

ISRv1 on ENCS1 and ISRv2 on ENCS2 are responsible for handling packets from LAN to WAN and WAN to LAN. If the WAN connection goes down or if the ISRv1 becomes unavailable, fast converging routing protocols, such as EIGRP and OSPF, can respond within seconds so that ISRv2 is prepared to transfer packets.

VM and Service Chain Network Connection

The Cisco ISRv should be created with an additional vNIC mapped to the transit link between two Cisco ENCS devices, apart from the regular WAN and LAN or service net links. The Cisco ISRv on both ENCS should have identical resource configurations (vNICs, vCPU, memory, etc.) and feature configurations.

Each Cisco ENCS is running an instance of service VNFs (for example, Cisco ASAv and Cisco vWAAS), and should have the identical service chain VNFs configured on both Cisco ENCS devices. Service VNFs should also have same features configured on both Cisco ENCS devices. The traffic goes through the service VNFs on the active Cisco ENCS only, even though both Cisco ENCS devices are actively forwarding on the WAN link. On a failover, the traffic will go over the service VNFs on the newly active ENCS (ENCS2).

This HA solution requires a transit link configured between two Cisco ENCS devices. One of the LAN ports from each of the Cisco ENCS can be connected back to back. This transit link port should be extended to the Cisco ISRv.

Enable Virtual NIC Failure Detection with Track Feature

You can enable the Track feature to detect virtual NIC failure in the following two scenarios:

- When the underlying physical link fails, the HSRP or routing protocols cannot detect the failure—This is because the line protocol does not go down when the underlying physical link fails if the Cisco ISRv is using a virtual NIC.
- With EEM scripts unconfigured, when the underlying physical link fails, the virtual NIC line protocol does not go down. In this case the routing protocol does not withdraw the routes.

Configuration Example for the Track Feature with Scenario 1 (HSRP)

In the virtual environment, you can enable higher protocols like HSRP to take action when the link failure happens. One way to achieve this is by configuring the Track feature on some object (ICMP ping) in Cisco IOS XE.

In Cisco ISRv, if the LAN interface where HSRP is running is a virtual NIC, then you can configure the track object to ping some device on the LAN segment, and monitor the connection failures. So, when the track object is down due to some connection failure, you can configure an action as to shut down the HSRP group, so that the peer will take over the active role making the default Gatway IP active. Without this track object, both Cisco ENCS devices will become active getting into a split-brain scenario.

The following example shows how to configure the track object on the active ISRv1, and monitor the connection failures by pinging the device IP in the network.



Note The Cisco ISRv should have AX license to configure the IP SLA.

```
track 1 ip sla 1 reachability
ip sla 1
 icmp-echo 192.0.2.1 source-ip 198.51.100.1
 frequency 5
ip sla schedule 1 life forever start-time now
!
track 5 ip sla 5 reachability
ip sla 5
 icmp-echo 192.0.2.2 source-ip 198.51.100.2
 frequency 5
ip sla schedule 5 life forever start-time now
!
```

The following output shows that the Track 1 reachability is failed, and Track 5 is up.

```
device1# show track
Track 1
  IP SLA 1 reachability
  Reachability is Down
    11 changes, last change 00:01:22
  Latest operation return code: Timeout
  Tracked by:
    HSRP GigabitEthernet3 25
Track 5
  IP SLA 5 reachability
  Reachability is Up
    4 changes, last change 00:02:32
  Latest operation return code: OK
  Latest RTT (millisecs) 1
  Tracked by:
    HSRP GigabitEthernet3 25
ISRv1#
```

The following example shows how to configure the Track object to monitor the line protocol state of the interface:

```
track 2 interface GigabitEthernet2 line-protocol
```

The following output shows that the line protocol state is down:

```
device# show track
Track 2
  Interface GigabitEthernet2 line-protocol
  Line protocol is Down ((hw down))
    8 changes, last change 00:01:25
```

```
  Tracked by:
    HSRP GigabitEthernet3 25
```

Configuration Example for the Track Feature with Scenario 2 (EEM Scripts)

With EEM scripts unconfigured, when an underlying link fails, the virtual NIC line protocol does not go down. This causes the problem as the routing protocol will not withdraw the routes. You can configure a Track object (can use the same object defined for HSRP above) to detect the failure. When the failure happens, the active Cisco ISRv has to withdraw the routes or network, so that the WAN link does not receive any traffic. One way to withdraw the routes is configure the EEM script, and delete the network from EIGRP.

The following example shows how to configure the EEM scripts, and remove the network from EIGRP:

```
track 5 ip sla 5 reachability
!
ip sla 5
 icmp-echo 192.0.2.1 source-ip 192.0.2.18
 frequency 5
ip sla schedule 5 life forever start-time now
!
event manager applet noshut_int
 event track 5 state up
 action 1.1 cli command "enable"
 action 1.2 cli command "config t"
 action 1.3 cli command "router eigrp 10"
 action 1.4 cli command "network 192.0.2.1 0.0.0.255"
 action 1.5 cli command "end"
event manager applet shut_int
 event track 5 state down
 action 1.1 cli command "enable"
 action 1.2 cli command "config t"
 action 1.3 cli command "router eigrp 10"
 action 1.4 cli command "no network 192.0.2.1 0.0.0.255"
 action 1.5 cli command "end"
```

In the virtual environment HSRP, make sure to use standby use-bia.

The following configuration example shows how to use the Track object (Track 5) to shut down HSRP group in ISRv1, when reachability is down for Track 5:

```
interface GigabitEthernet4
 description Service-NET-Virtio
 ip address 192.0.2.1 255.255.255.0
 standby use-bia
 standby 25 ip 192.0.2.22
 standby 25 timers 1 5
 standby 25 priority 105
 standby 25 preempt
 standby 25 track 5 shutdown
```

Isolating LAN and Transit Link Traffic for vBranch HA

LAN traffic and transit link traffic shall be isolated by configuring different VLANs for each traffic since both links are connected to the same ENCS internal switch. If you do not isolate these traffic, both LAN traffic and transit link will flow through the same internal switch on the Cisco ENCS.

The following Cisco ENCS switch configuration example shows how to isolate traffic. In this example, the Cisco ISRv is configured to send HSRP traffic as an untag and transit traffic in VLAN 46. So, to isolate HSRP traffic and transit traffic on the internal switch, the Gigabit Ethernet interface 1/0 is connected to a LAN network and Gigabit Ethernet interface 1/1 is configured as the transit link. The Gigabit Ethernet interface 1/1 allows the VLAN 46 to pass the transit traffic. It should also have non-default (other than 1) native VLAN (for example, VLAN 2), because the Cisco ENCS internal switch uplink (internal) has the native VLAN 1 configured.

Enable MSTP on all switches before isolating traffic.

```
switch
 interface gigabitEthernet1/0
  negotiation auto
  no shutdown
  switchport access vlan 1
  switchport mode access
  switchport trunk native vlan 1
  switchport trunk allowed vlan 1-2349,2450-4093
!
!
!
switch
 interface gigabitEthernet1/1
  negotiation auto
  no shutdown
spanning-tree mst 1 cost 200000000
 spanning-tree mst 2 cost 200000000
  switchport access vlan 46
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 1-2349,2450-4093
!
!
!
 spanning-tree enable
 spanning-tree mode mst
 spanning-tree mst configuration
 name     region1
  revision 1
  instance 1 vlan 1
  instance 2 vlan 46
 !
```

Use the **show switch vlan detailed** command to verify the configuration as shown below:

```
device# show switch vlan detailed

platform-detail hardware_info Manufacturer "Cisco Systems, Inc."
platform-detail hardware_info PID ENCS5412/K9
platform-detail hardware_info SN FGL212681GK
platform-detail hardware_info hardware-version M3
platform-detail hardware_info UUID 7BBEBDE0-CE3C-42E5-B564-CFEE8F18AE97
platform-detail hardware_info Version 3.8.1-FC3
platform-detail hardware_info Compile_Time "Sunday, April 15, 2018 [20:38:10 PDT]"
platform-detail hardware_info CPU_Information "Intel(R) Xeon(R) CPU D-1557 @ 1.50GHz 12
cores"
platform-detail hardware_info Memory_Information "16227148 kB"
platform-detail hardware_info Disk_Size "64.0 GB"
platform-detail hardware_info CIMC_IP NA
platform-detail hardware_info Entity-Name ENCS
platform-detail hardware_info Entity-Desc "Enterprise Network Compute System"
```

```
platform-detail software_packages Kernel_Version 3.10.0-514.21.1.1.el7.x86_64
platform-detail software_packages QEMU_Version 1.5.3
platform-detail software_packages LibVirt_Version 3.2.0
platform-detail software_packages OVS_Version 2.5.2
platform-detail switch_detail UUID NA
platform-detail switch_detail Type NA
platform-detail switch_detail Name NA
platform-detail switch_detail Ports 8
                                                                                PCI
NAME   TYPE      MEDIA         LINK  SPEED  MTU   MAC                 DETAIL
-------------------------------------------------------------------------------
GE0-0  physical  Twisted Pair  up    1000   9216  70:db:98:c3:f3:64   02:00.0
GE0-1  physical  Twisted Pair  up    1000   9216  70:db:98:c3:f3:65   02:00.1
MGMT   physical  Twisted Pair  up    1000   1500  70:db:98:c3:f3:d8   0e:00.0
```

Packet Flow for vBranch HA

This section explains high-level packet flow in failure and non-failure cases.

Non-Failure Case

In the non-failure case, both active and standby Cisco ENCS devices are up and running.

- LAN to WAN through the standby ENCS1 WAN link
  - The device in the LAN segment is configured with the default gateway as the HSRP virtual IP address, and since ENCS1 is an active HSRP, LAN traffic first comes to the active ENCS1.
  - LAN traffic goes through the service chain VM (Cisco ASAv), and then hits the Cisco ISRv. In this case, the destination IP is routable through the ENCS1 WAN interface. The Cisco ISRv sends traffic over the WAN link.
- LAN to WAN through the standby ENCS2 WAN link—In this case, the LAN to WAN traffic uses the transit link between the active and standby devices.
  - Devices in the LAN segment are configured with the default gateway as the HSRP virtual IP address, and since ENCS1 is an active HSRP, the LAN traffic first comes to the active ENCS1.
  - The LAN traffic goes through the service chain VMs (Cisco ASAv), and then hits the active Cisco ISRv. In this case, the destination IP is routable through the ENCS2 WAN interface. The traffic is sent to the Cisco ISRv on ENCS2 over the transit link, and then sent out over the WAN link to the destination.
- WAN to LAN through the active ENCS1
  - The WAN traffic hits the Cisco ISRv on ENCS1, then it goes through the service chain VMs, and sent to the LAN device.
- WAN to LAN through the standby ENCS2 WAN link—In this case, the WAN to LAN traffic uses the transit link between the active and standby devices.
  - The WAN traffic comes to the Cisco ISRv on ENCS2. The PBR/PFR configuration forces the traffic to use the transit link instead of the directly connected LAN port. So, the traffic is sent to the Cisco ISRv on ENCS1 over the transit link.

- Then, the traffic on ENCS1 goes through the service chain VMs, and sent to the LAN device.

Failure Case

In the failure case, the active device goes down, and the standby device becomes active.

The virtual IP (default gateway) address becomes active on ENCS2. The transit link will not be used. The traffic now goes through the service chain VMs on ENCS2, and gets forwarded directly between WAN and LAN interfaces. The PBR/PFR configuration should monitor the HSRP state, and use the LAN port instead of the transit link to forward LAN traffic.

Configuration Examples for vBranch HA

This sample configuration is for Cisco ENCS HA with a dual-WAN scenario. The Cisco ISRv is configured with vNICs connected to the wan-net, service-net, and transit link. HSRP is configured on the service-net interface. Each Cisco ENCS is provisioned with the Cisco ASAv (service-net) and Cisco vWAAS (service-net).



Note You can use this design as is, or modify as per the field requirement.

Example: Active Cisco ENCS Configuration with ISRv1

```
interface GigabitEthernet1
 vrf forwarding Mgmt-intf
 ip address 192.0.2.1 255.255.255.0
 negotiation auto
!
interface GigabitEthernet2
 description WAN-GE0-0-SRIOV-1
 ip address 192.0.2.2 255.255.255.0
 negotiation auto
!
interface GigabitEthernet3
 description LAN-NET
 no ip address
 shutdown
 negotiation auto
!
interface GigabitEthernet4
 description Service-NET-Virtio
 ip address 192.0.2.3 255.255.255.0
 standby use-bia
 standby 25 ip 192.0.2.20
 standby 25 timers 1 5
 standby 25 priority 105
 standby 25 preempt
 standby 25 track 1 decrement 10
 standby 25 track 2 decrement 10
 standby 25 track 3 decrement 10
 standby 25 track 5 shutdown
 standby 25 track 6 shutdown
 standby 25 track 7 shutdown
 negotiation auto
```

```
 bfd interval 9000 min_rx 9000 multiplier 3
!
interface GigabitEthernet5
 ip address 192.0.2.4 255.255.255.0
!
!
router eigrp stub 10
 network 25.25.25.0 0.0.0.255
 network 38.38.38.0 0.0.0.255
 network 46.46.46.0 0.0.0.255
  !
    !
track 1 ip sla 1 reachability
!
track 2 interface GigabitEthernet2 line-protocol
!
track 3 interface GigabitEthernet4 line-protocol
!
track 5 ip sla 5 reachability
!
track 6 ip sla 6 reachability
!
track 7 ip sla 7 reachability
!
ip sla 1
 icmp-echo 9.9.9.29 source-ip 192.0.2.2
 frequency 5
ip sla schedule 1 life forever start-time now
!
ip sla 5
 icmp-echo 25.25.25.11 source-ip 192.0.2.3
 frequency 5
ip sla schedule 5 life forever start-time now
!
ip sla 6
 icmp-echo 25.25.25.51 source-ip 192.0.2.3
 frequency 5
ip sla schedule 6 life forever start-time now
!
ip sla 7
 icmp-echo 25.25.25.75 source-ip 192.0.2.3
 frequency 5
ip sla schedule 7 life forever start-time now
!
event manager applet noshut_int
 event track 5 state up
 action 1.1 cli command "enable"
 action 1.2 cli command "config t"
 action 1.3 cli command "router eigrp 10"
 action 1.4 cli command "network 25.25.25.0 0.0.0.255"
 action 1.5 cli command "end"
event manager applet shut_int
 event track 5 state down
 action 1.1 cli command "enable"
 action 1.2 cli command "config t"
 action 1.3 cli command "router eigrp 10"
 action 1.4 cli command "no network 25.25.25.0 0.0.0.255"
 action 1.5 cli command "end"
event manager applet ASAv_noshut_int
 event track 6 state up
 action 1.1 cli command "enable"
 action 1.2 cli command "config t"
 action 1.3 cli command "router eigrp 10"
 action 1.4 cli command "network 25.25.25.0 0.0.0.255"
```

```
 action 1.5 cli command "end"
event manager applet ASAv_shut_int
event track 6 state down
 action 1.1 cli command "enable"
 action 1.2 cli command "config t"
 action 1.3 cli command "router eigrp 10"
 action 1.4 cli command "no network 25.25.25.0 0.0.0.255"
 action 1.5 cli command "end"
event manager applet vWAAS_noshut_int
 event track 7 state up
 action 1.1 cli command "enable"
 action 1.2 cli command "config t"
 action 1.3 cli command "router eigrp 10"
 action 1.4 cli command "network 25.25.25.0 0.0.0.255"
 action 1.5 cli command "end"
event manager applet vWAAS_shut_int
 event track 7 state down
 action 1.1 cli command "enable"
 action 1.2 cli command "config t"
 action 1.3 cli command "router eigrp 10"
 action 1.4 cli command "no network 25.25.25.0 0.0.0.255"
 action 1.5 cli command "end"
!
end
```

Example: Standby Cisco ENCS Configuration with ISRv2

```
interface GigabitEthernet1
 vrf forwarding Mgmt-intf
 ip address 192.0.2.1 255.255.255.0
 negotiation auto
!
interface GigabitEthernet2
 description WAN-GE0-0-SRIOV-1
 ip address 192.0.2.21 255.255.255.0
 negotiation auto
!
interface GigabitEthernet3
 no ip address
 shutdown
 negotiation auto
!
interface GigabitEthernet4
 description Service-NET-virtio
 ip address 192.0.2.22 255.255.255.0
 standby use-bia
 standby 25 ip 192.0.2.20
 standby 25 timers 1 5
 standby 25 preempt
 negotiation auto
 bfd interval 9000 min_rx 9000 multiplier 3
!
interface GigabitEthernet5
 ip address 192.0.2.23 255.255.255.0
!
!
router eigrp 10
network 8.8.8.0 0.0.0.255
 network 25.25.25.0 0.0.0.255
 network 46.46.46.0 0.0.0.255
```

!

Cisco ENCS Failure Points

Failure Points	Sequence of Events
ENCS chassis hardware failure: • Power down • Power cycle • Reboot Cisco Enterprise NFVIS software failure • Crash Cisco ISRv software failure • Stop (shutdown) • Reboot • Crash • Error	1. HSRP on ENCS2 detects the reachability failure to ENCS1, and triggers the failover. LAN virtual IP becomes active on ENCS2. 2. WAN-IP1 on ENCS1 becomes unreachable, and all the routes converge towards WAN-IP2 on ENCS2. WAN-IP2 is the only IP for branch connectivity. 3. All the WAN to LAN, and LAN to WAN traffic will now flow through ENCS2. 4. The PBR/PFR configuration will now select the LAN port as the preferred path instead of the transit link for the traffic heading to LAN.

ISRv1 (Active) Before the Failure

```
ISRv1# show platform software vnic-if interface-mapping
------------------------------------------------------------
 Interface Name        Driver Name         Mac Addr
------------------------------------------------------------
 GigabitEthernet5       i40evf             5254.003a.1020 (LAN-SRIOV-2)
 GigabitEthernet4       virtio             5254.0053.e392 (service-net)
 GigabitEthernet3       i40evf             5254.00c4.b925 (LAN-SRIOV-1)
 GigabitEthernet2       igbvf              5254.00d2.cc9a (GE0-0-SRIOV-1)
 GigabitEthernet1       virtio             5254.00d2.1b1c (int-mgmt-net)
------------------------------------------------------------

ISRv1# show standby brief
                     P indicates configured to preempt.
                     |
Interface   Grp  Pri P State   Active          Standby         Virtual IP
Gi4         25   105 P Active  local           192.0.2.1     192.0.2.222
#
```

ISRv2 (Standby) Before the Failure

```
ISRv2#show platform software vnic-if interface-mapping
------------------------------------------------------------
 Interface Name        Driver Name         Mac Addr
------------------------------------------------------------
 GigabitEthernet5       i40evf             5254.00cc.ce9f (LAN-SRIOV-2)
 GigabitEthernet4       virtio             5254.00e7.523f (Service-net)
```

```
 GigabitEthernet3       i40evf             5254.0055.ee45 (LAN-SRIOV-1)
 GigabitEthernet2       igbvf              5254.00a3.d443 (GE0-0-SRIOV-1)
 GigabitEthernet1       virtio             5254.0048.e84c (int-mgmt-net)
------------------------------------------------------------
```

```
ISRv2#show standby brief
                     P indicates configured to preempt.
                     |
Interface   Grp  Pri P State   Active          Standby         Virtual IP
Gi4         25   100 P Standby 192.0.2.20     local           192.0.2.222
```

ISRv2 After the Failure

ISRv1 becomes unreachable. ISRv2: The HSRP failover occurs, and the state changes from Standby to Active. The virtual IP (LAN side default gateway) becomes active on ENCS2 ISRv2.

```
ISRv2# show standby brief
                     P indicates configured to preempt.
                     |
Interface   Grp  Pri P State   Active          Standby         Virtual IP
Gi4         25   100 P Active  local           unknown         192.0.2.222
```

```
ISRv2# show logging
*Dec 13 21:22:17.138: %HSRP-5-STATECHANGE: GigabitEthernet4 Grp 25 state Speak -> Standby
*Dec 13 21:22:32.385: %HSRP-5-STATECHANGE: GigabitEthernet4 Grp 25 state Standby -> Active
```

Failure Points	Sequence of Events
WAN Net1 failure (WAN SRIOV VF connected to ISRv1) • Link down WAN Phy link failure • Switch failure • End-to-end connectivity failure	**1.** ISRv1 HSRP on ENCS1 detects the WAN connection failure. It reduces the LAN-HSRP priority. This failure is detected when the interface goes down due to VF going down or track object going down. **2.** WAN-IP1 becomes unreachable, and all the routes converge towards WAN-IP2 on ENCS2. WAN-IP2 is the only IP for branch connectivity. **3.** HSRP on ENCS2 becomes higher priority in the group, and takes over the active role. LAN-virtual IP becomes active on ENCS2. **4.** The PBR/PFR configuration will now select the LAN port as the preferred path instead of the transit link for the traffic destined to LAN. **5.** All the WAN to LAN, and LAN to WAN traffic will now flow through ENCS2.

ISRv1 After the Failure

ISRv1 becomes standby.

```
ISRv1# show ip interface brief
Interface              IP-Address      OK? Method Status                Protocol
GigabitEthernet1       192.0.2.1       YES NVRAM  up                    up
GigabitEthernet2       192.0.2.2     YES NVRAM  down                  down
```

```
GigabitEthernet3       unassigned      YES NVRAM  administratively down down
GigabitEthernet4       192.0.2.3     YES NVRAM  up                    up
GigabitEthernet5       unassigned      YES NVRAM  up                    up

ISRv1# show standby brief
                     P indicates configured to preempt.
                     |
Interface   Grp  Pri P State   Active          Standby         Virtual IP
Gi4         25   85  P Standby 192.0.2.22     local           192.0.2.222
ISRv1#
ISRv1#show logg
*Dec 14 03:41:52.307: %TRACK-6-STATE: 2 interface Gi2 line-protocol Up -> Down
*Dec 14 03:42:37.744: %HSRP-5-STATECHANGE: GigabitEthernet4 Grp 25 state Active -> Speak
*Dec 14 03:42:43.663: %HSRP-5-STATECHANGE: GigabitEthernet4 Grp 25 state Speak -> Standby
ISRv1#
ISRv1#show track
Track 1
  IP SLA 1 reachability
  Reachability is Down
    1405 changes, last change 00:03:08
  Latest operation return code: Timeout
  Tracked by:
    HSRP GigabitEthernet4 25
```

ISRv2 After the Failure

```
ISRv2# show standby brief
                     P indicates configured to preempt.
                     |
Interface   Grp  Pri P State   Active          Standby         Virtual IP
Gi4         25   100 P Active  local           192.0.2.3     192.0.2.222
```

Failure Points	Sequence of Events
LAN Phy link failure • Switch failure • End-to-end connectivity failure LAN connectivity failure • Switch failure • End-to-end connectivity failure SC Net failure (ISRv service-net down) • Link down VNFs (Cisco ASAv, Cisco vWAAS, and Windows/Linux) • Power down • Power cycle • Crash/reboot	**1.** ISRv1 HSRP on ENCS1 detects the LAN connection failure, and shut down the HSRP group. This failure is detected when the interface goes down due to the track object going down. **2.** EEM script on ISRv1 withdraws the routes (for example, delete EIGRP networks). All the branch traffic routes will now converge towards WAN-IP2 on ENCS2. WAN-IP2 is the only IP for branch connectivity. **3.** HSRP on ENCS-2 becomes active in the group. LAN virtual IP becomes active on ENCS2. **4.** On ISRv2, the PBR/PFR configuration will now select the LAN port as the preferred path, instead of the transit link for the traffic destined to LAN. **5.** All the WAN to LAN and LAN to WAN traffic will now flow through ENCS2.

ISRv1 After the Failure

```
ISRv1# show track
Track 7
  IP SLA 7 reachability
  Reachability is Down
    7 changes, last change 00:01:40
  Latest operation return code: Timeout
  Tracked by:
    HSRP GigabitEthernet3 25
    EEM 2450904616
    EEM 2450905656
ISRv1#

ISRv1# show ip eigrp topo
EIGRP-IPv4 Topology Table for AS(10)/ID(53.53.53.51)
Codes: P - Passive, A - Active, U - Update, Q - Query, R - Reply,
       r - reply Status, s - sia Status

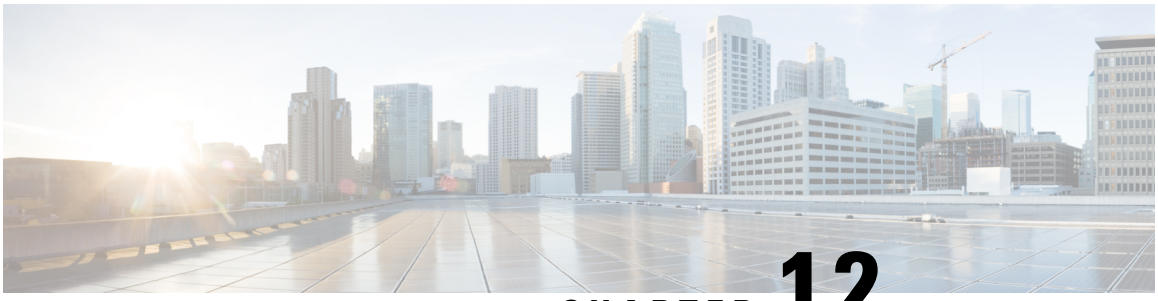
P 19.19.19.0/24, 1 successors, FD is 3328
        via 38.38.38.38 (3328/3072), GigabitEthernet2
P 9.9.9.0/24, 1 successors, FD is 3328
        via 38.38.38.38 (3328/3072), GigabitEthernet2
P 25.25.25.0/24, 0 successors, FD is Infinity
        via 38.38.38.38 (3840/3584), GigabitEthernet2
P 27.27.27.0/24, 1 successors, FD is 3328
        via 38.38.38.38 (3328/3072), GigabitEthernet2
P 38.38.38.0/24, 1 successors, FD is 2816
        via Connected, GigabitEthernet2
P 29.29.29.0/24, 1 successors, FD is 3072
        via 38.38.38.38 (3072/2816), GigabitEthernet2
P 33.33.33.0/24, 1 successors, FD is 3840
        via 38.38.38.38 (3840/3584), GigabitEthernet2
P 8.8.8.0/24, 1 successors, FD is 3584
        via 38.38.38.38 (3584/3328), GigabitEthernet2
P 53.53.53.0/24, 1 successors, FD is 2816
        via Connected, GigabitEthernet4
```

ISRv2 After the Failure

```
ISRv2# show standby brief
                     P indicates configured to preempt.
                     |
Interface   Grp  Pri P State   Active          Standby         Virtual IP
Gi3         25   100 P Active  local           unknown         192.0.2.222
```

Failure Points	Sequence of Events
WAN Net2 failure (WAN SRIOV VF connected to ISRv2 is down) • Link down	1. ISRv2 on ENCS2 detects the WAN connection failure. This failure is detected when the interface goes down due to VF going down or the track object going down. 2. WAN-IP2 becomes unreachable, and all the routes converge towards WAN-IP1 on ENCS1. WAN-IP1 is the only IP for branch connectivity. 3. All the WAN to LAN and LAN to WAN traffic will now flow through ENCS1.

Failure Points	Sequence of Events
Transit link between ENCS1 and ENCS2 fails • Link down	**1.** ISRv2 on ENCS2 detects the link going down due to VF going down or connection failure. The connection failure is detected by the track object. Then, ENCS2 WAN-IP2 link with EEM script is shut down. **2.** WAN-IP2 becomes unreachable, and all the routes converge towards WAN-IP1 on ENCS1. WAN-IP1 is the only IP for branch connectivity. **3.** All the WAN to LAN and LAN to WAN traffic will now flow through ENCS1.



CHAPTER 12

Cisco ENCS Single WAN IP Deployment Scenarios

- Single WAN IP Deployment, on page 125
- Preconfiguring the Cisco ENCS for a Single WAN IP Deployment, on page 126
- Single WAN IP Deployment with Gigabit Ethernet Interface 0/0, on page 127
- Single WAN IP Deployment with the 4G Interface, on page 128

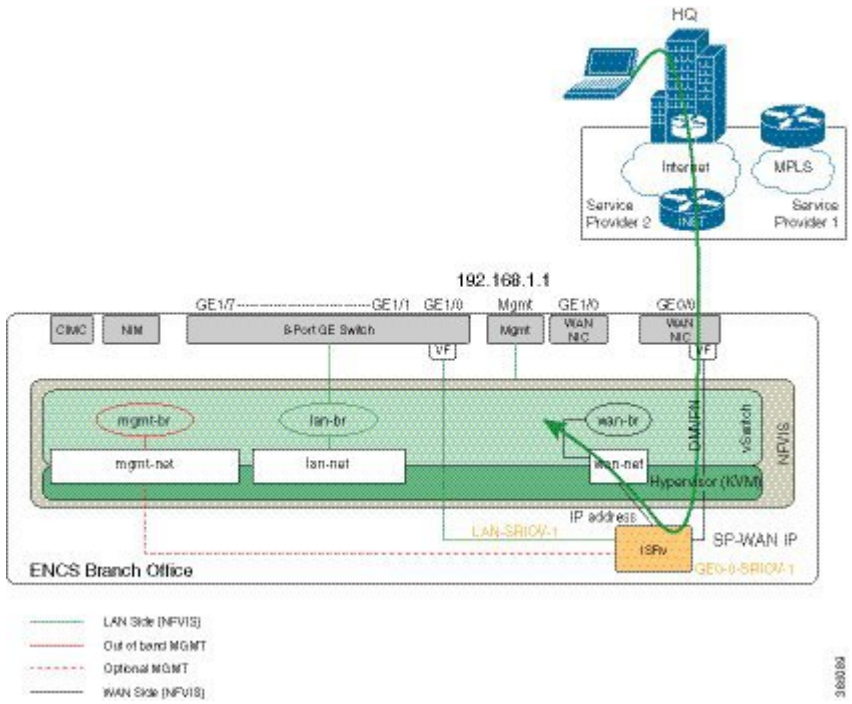
Single WAN IP Deployment

A single WAN IP deployment can be considered when the Cisco ENCS is preconfigured at the corporate main office with the service provider's WAN IP address, and shipped to the branch office for quick deployment. At the branch office, you do not have to perform any installation or configuration task. You just have to boot the system with the preconfigured setup. The single WAN IP deployment scenario could vary as per customer requirements. The following are two sample single WAN IP deployment scenarios with the Cisco ISRv:

Note Ensure that you preconfigure the Cisco ENCS at the main office before shippping the device to the branch office. You cannot connect to the remote branch office from your main office in a single WAN IP deployment scenario.

- Single WAN IP Deployment with Gigabit Ethernet Interface
- Single WAN IP Deployment with the 4G Interface

Figure 9: Single WAN IP Deployment Topology



Preconfiguring the Cisco ENCS for a Single WAN IP Deployment

To preconfigure the Cisco ENCS:

1. Install Cisco Enterprise NFVIS on the Cisco ENCS via CIMC. For details, see Installing Cisco Enterprise NFVIS on a Cisco ENCS 5100 and 5400, on page 12.
2. Connect your local system (laptop) to the local management interface of the host server.
3. Open the Cisco Enterprise NFVIS portal via https://192.168.1.1.
4. Upload the Cisco ISRv image using the portal, and register the VM.
5. From the portal, remove the default Gigabit Ethernet 0/0 or GE0-0 WAN interface.
6. Deploy Cisco ISRv with Gigabit Ethernet 2 for SRIOV-1 and Gigabit Ethernet 3 for the wan-net.
7. Open the Cisco ISRv VNC.
8. From the VNC console, configure ISRv Gigabit Ethernet 2 and Gigabit Ethernet 3 interfaces with appropriate IP addresses. Then, perform a "no shut" of the interfaces.
9. Set the WAN static IP address to be on the same subnet as ISRv Gigabit Ethernet 2 IP address, and use ISRv Gigabit Ethernet 2 interface IP address as the default gateway.
10. Ping with the Cisco ISRv IP address to ensure connectivity.
11. Configure Dynamic Multipoint VPN on the Cisco ISRv, and ensure the main server can access the portal.

For details, see the Dynamic Multipoint VPN Configuration Guide https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/sec_conn_dmvpn/configuration/xe-16/sec-conn-dmvpn-xe-16-book.html .

Single WAN IP Deployment with Gigabit Ethernet Interface 0/0

In this scenario, two Gigabit Ethernet interfaces are configured on the Cisco ISRv: Gigabit Ethernet2 as the outbound interface and Gigabit Ethernet3 as the internal interface. The outbound interface IP address is provided by the service provider. The internal interface is the WAN interface that serves as the default gateway for Cisco Enterprise NFVIS.

```
crypto isakmp policy 5
 authentication pre-share
 group 2
crypto isakmp key dmvpnkey address 0.0.0.0

crypto ipsec transform-set dmvpnset esp-3des esp-sha-hmac
 mode tunnel

crypto ipsec profile dmvpnprof
 set security-association lifetime seconds 1200
 set transform-set dmvpnset

! DMVPN tunnel configuration
interface Tunnel100
 ip address 192.0.2.3 255.255.255.0
 no ip redirects
 ip mtu 1440
 ip nhrp authentication dmvpnkey
 ip nhrp map 192.0.2.1 198.51.100.1
 ip nhrp network-id 90
 ip nhrp nhs 192.0.2.2
 tunnel source GigabitEthernet2
 tunnel mode gre multipoint
 tunnel key 100000
 tunnel protection ipsec profile dmvpnprof
!
interface GigabitEthernet2
 description this is the outbound interface
 ip address 198.51.100.2 255.255.0.0

interface GigabitEthernet3
 description this is the inside interface
 ip address 192.0.2.10 255.255.255.0
!

router eigrp 90
 network 10.4.76.0 0.0.0.255
 network 192.0.2.1
 eigrp stub connected
 no auto-summary
!
ip route 20.1.0.0 255.255.0.0 198.51.100.1
!
Smart license configuration

ip name-server 198.51.100.9
ip domain lookup
service internal
do test license smart dev-cert Enable
```

```
service call-home
call-home
 contact-email-addr callhome@cisco.com
 mail-server 192.0.2.8 priority 1
 alert-group-config snapshot
  add-command "show license tech su"
 profile "CiscoTAC-1"
  active
  no destination transport-method email
  destination transport-method http
  no destination address http https://example.com/
  destination address http http://10.22.183.117:8080/ddce/services/DDCEService
!
clock timezone PST -7
ntp server 192.0.2.9
do license smart register idtoken NDM1NjE1MDAtNDViZC00ZTQ5LTg4MGEtNmRj
Njg2Mjg5ZDVlLTE0OTg5NDk2%0ANjEzNzd8elk5SEtoL2pMTGtuNSs3Q3Jxd
GVoSUVpTmFnY2l0alVqR3B5MzFj%0AVWVrST0%3D%0A
```

Single WAN IP Deployment with the 4G Interface

In this scenario, a 4G interface (NIM card) is configured as the outbound interface and Gigabit Ethernet3 as the internal interface. The outbound interface IP address is provided by the service provider. The internal interface is the WAN interface that serves as the default gateway for Cisco Enterprise NFVIS.

```
License Level: ax
License Type: N/A(Smart License Enabled)
Next reload license Level: ax

service timestamps debug datetime msec
service timestamps log datetime msec
service internal
service call-home
no platform punt-keepalive disable-kernel-core
platform console virtual
platform hardware throughput level MB 1000
!
hostname ISRv
!
boot-start-marker
boot system bootflash:isrv-universalk9.16.03.02.SPA.bin
boot-end-marker

clock timezone PST -7 0
call-home
 contact-email-addr callhome@cisco.com
 mail-server 192.0.2.8 priority 1
 alert-group-config snapshot
  add-command "show license tech su"
 profile "CiscoTAC-1"
  active
  destination transport-method http
  no destination transport-method email
  destination address http
http://198.51.100.4/Transportgateway/services/DeviceRequestHandler
  no destination address http
https://example.com/
!
ip name-server 198.51.100.2
```

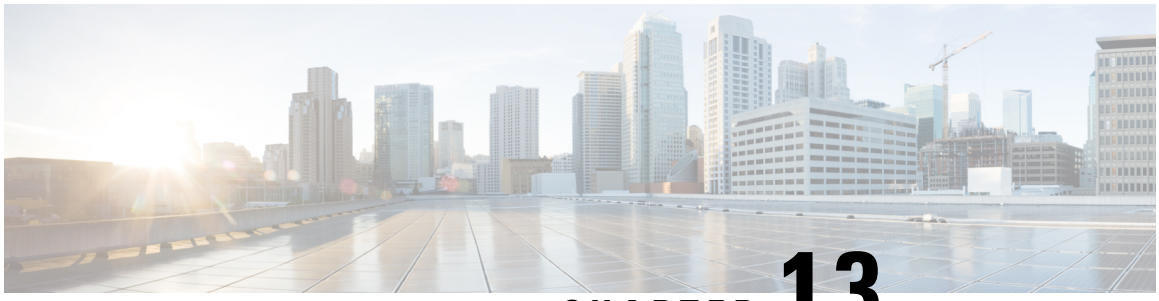
```
ip domain name cisco.com

! IPsec configuration

crypto isakmp policy 5
 authentication pre-share
 group 2
crypto isakmp key dmvpnkey address 0.0.0.0
!
!
crypto ipsec transform-set dmvpnset esp-3des esp-sha-hmac
 mode tunnel
!
!
crypto ipsec profile dmvpnprof
 set security-association lifetime seconds 1200
 set transform-set dmvpnset
!
!4G interface
controller Cellular 0/2/0
 lte modem link-recovery rssi onset-threshold -110
 lte modem link-recovery monitor-timer 20
 lte modem link-recovery wait-timer 10
 lte modem link-recovery debounce-count 6
!
!
!
no ip ftp passive
ip ftp username admin
ip ftp password admin
!DMVPN tunnel configuration

interface Tunnel100
 ip address 198.51.100.3 255.255.255.0
 no ip redirects
 ip mtu 1440
 ip nhrp authentication dmvpnkey
 ip nhrp map 198.51.100.5 192.0.2.7
 ip nhrp network-id 90
 ip nhrp nhs 198.51.100.5
 tunnel source Cellular0/2/0
 tunnel mode gre multipoint
 tunnel key 100000
 tunnel protection ipsec profile dmvpnprof
!
!
interface GigabitEthernet2
 ip address 198.51.100.6 255.255.255.0
 ip nat inside
 negotiation auto
!
interface GigabitEthernet3
 ip address 198.51.100.11 255.255.255.0
 negotiation auto
!
interface Cellular0/2/0
 ip address negotiated
 load-interval 30
 dialer in-band
 dialer idle-timeout 0
 dialer-group 1
 ipv6 address autoconfig
 pulse-time 1
!
```

```
interface Cellular0/2/1
 no ip address
!
!
router eigrp 90
 network 198.51.100.0 0.0.0.255
 network 198.52.100.0 0.0.0.255
 network 99.0.0.0
 eigrp stub connected
!
!
virtual-service csr_mgmt
 ip shared host-interface GigabitEthernet1
 activate
!
ip forward-protocol nd
ip http server
ip http authentication local
ip http secure-server
!
ip route 0.0.0.0 0.0.0.0 Cellular0/2/0
ip route 192.0.2.12 255.255.255.0 198.51.100.5
ip route 192.0.2.13 255.255.255.255 198.51.100.5
ip route 192.0.2.14 255.255.255.255 198.51.100.5
ip route 192.0.2.15 255.255.255.255 198.51.100.5
ip route vrf Mgmt-intf 0.0.0.0 0.0.0.0 198.51.100.20
ip ssh authentication-retries 5
ip ssh rsa keypair-name ssh-key
ip ssh version 2
ip scp server enable
!
dialer-list 1 protocol ip permit
!
!
line con 0
 stopbits 1
line vty 0 4
 password cisco123
 login local
 transport input telnet ssh
!
ntp server 198.51.100.17
```



CHAPTER 13

Resetting to Factory Default

Factory default reset is available on all NFVIS supported hardware platforms.

You can reset the host server to factory default with the following options:

- Reset all—Deletes VMs and volumes, files including logs, images, and certificates. Erases all configuration. Connectivity will be lost, and the admin password will be changed to factory default password.
- Reset all-except-images—Delete VMs and volumes, files including logs, user uploaded files and certificates. Erases all configuration except registered images. Connectivity will be lost, and the admin password will be changed to factory default password.
- Reset all-except-images-connectivity—Deletes VMs and volumes, files including logs and certificates. Erases all configuration except images, network, and connectivity.

Note Factory default reset must be used only for troubleshooting purpose. We recommend you contact Cisco Technical Support before performing factory default reset. This feature will reboot the system. Do not perform any operations until the system reboots successfully.

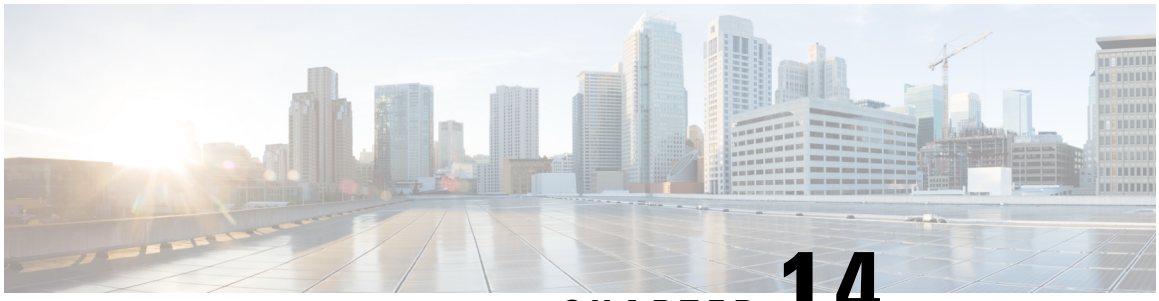
To reset to factory default:

```
nfvis#factory-default-resetall|all-except-images|all-except-images-connectivity
```

Note Enter **Yes** when you are prompted with the factory default warning message or **no** to cancel.

Factory Default APIs and Commands

Factory Default APIs	Factory Default Commands
• /api/operations/factory-default-reset/all • /api/operations/factory-default-reset/all-except-images • /api/operations/factory-default-reset /all-except-images-connectivity	• factory-default-reset



CHAPTER 14

NETCONF Event Notifications

Cisco Enterprise NFVIS generates event notifications for key events. A NETCONF client can subscribe to these notifications for monitoring the progress of configuration activation and the status change of the system and VMs.

There are two types of event notifications: nfvisEvent and vmlcEvent (VM life cycle event)

To receive event notifications automatically, you can run the NETCONF client, and subscribe to these notifications using the following NETCONF operations:

- --create-subscription=nfvisEvent
- --create-subscription=vmlcEvent

You can view NFVIS and VM life cycle event notifications using the **show notification stream nfvisEvent** and **show notification stream vmlcEvent** commands respectively. For more information see, Event Notifications.

- nfvisEvent, on page 134
- vmlcEvent, on page 150

nfvisEvent

Event Type	Notification Trigger	Notification Output Example
WAN_DHCP_RENEW	DHCP renew operation is performed.	<?xml version="1.0" encoding="UTF-8"?> <notification xmlns="urn:ietf:params:xml:ns:netconf:notification:1.0"> <eventTime>2017-04-26T18:06:46.142089+00:00</eventTime> <nfvisEvent xmlns="http://www.cisco.com/nfvis/notifier"> <user_id>NA</user_id> <config_change>false</config_change> <transaction_id>0</transaction_id> <status>SUCCESS</status> <status_code>0</status_code> <status_message>Wan DHCP IP address is being renewed</status_message> <details>NA</details> <event_type>**WAN_DHCP_RENEW**</event_type> </nfvisEvent> </notification>
BRIDGE_DHCP_RENEW	Bridge DHCP renew operation is performed.	<?xml version="1.0" encoding="UTF-8"?> <notification xmlns="urn:ietf:params:xml:ns:netconf:notification:1.0"> <eventTime>2018-04-26T09:47:06.066264+00:00</eventTime> <nfvisEvent xmlns="http://www.cisco.com/nfvis/notifier"> <user_id>NA</user_id> <config_change>false</config_change> <transaction_id>0</transaction_id> <status>SUCCESS</status> <status_code>0</status_code> <status_message>Bridge DHCP IP address is being renewed</status_message> <details>NA</details> <event_type>**BRIDGE_DHCP_RENEW**</event_type> </nfvisEvent> </notification>

Event Type	Notification Trigger	Notification Output Example
INTF_STATUS_CHANGE	Interface status is changed.	<?xml version="1.0" encoding="UTF-8"?> <notification xmlns="urn:ietf:params:xml:ns:netconf:notification:1.0"> <eventTime>2017-04-26T18:12:09.963556+00:00</eventTime> <nfvisEvent xmlns="http://www.cisco.com/nfvis/notifier"> <event_type>**INTF_STATUS_CHANGE**</event_type> <intf_name>eth7</intf_name> <intf_prv_op>up</intf_prv_op> <intf_op>down</intf_op> <intf_prv_link>down</intf_prv_link> <intf_link>down</intf_link> </nfvisEvent> </notification>
NETWORK_CREATE	A network is created.	<?xml version="1.0" encoding="UTF-8"?> <notification xmlns="urn:ietf:params:xml:ns:netconf:notification:1.0"> <eventTime>2016-09-22T12:41:04.564298+00:00</eventTime> <nfvisEvent xmlns="http://www.cisco.com/nfvis/notifier"> <user_id>admin</user_id> <config_change>true</config_change> <transaction_id>0</transaction_id> <status>SUCCESS</status> <status_message>Network created succesfully</status_message> <event_type>**NETWORK_CREATE**</event_type> <network_name>testn1</network_name> <network_bridge>test-net-br</network_bridge> <network_sriov>false</network_sriov> <network_vlan/> <network_trunk/> </nfvisEvent> </notification>

Event Type	Notification Trigger	Notification Output Example
NETWORK_UPDATE	A network is updated.	<?xml version="1.0" encoding="UTF-8"?> <notification xmlns="urn:ietf:params:xml:ns:netconf:notification:1.0"> <eventTime>2016-09-22T12:42:03.391986+00:00</eventTime> <nfvisEvent xmlns="http://www.cisco.com/nfvis/notifier"> <user_id>admin</user_id> <config_change>true</config_change> <transaction_id>0</transaction_id> <status>SUCCESS</status> <status_message>Network updated succesfully</status_message> <event_type>**NETWORK_UPDATE**</event_type> <network_name>testn1</network_name> <network_bridge/> <network_sriov/> <network_vlan/> <network_trunk/> </nfvisEvent> </notification>
NETWORK_DELETE	A network is deleted.	<?xml version="1.0" encoding="UTF-8"?> <notification xmlns="urn:ietf:params:xml:ns:netconf:notification:1.0"> <eventTime>2016-09-22T12:42:03.391986+00:00</eventTime> <nfvisEvent xmlns="http://www.cisco.com/nfvis/notifier"> <user_id>admin</user_id> <config_change>true</config_change> <transaction_id>0</transaction_id> <status>SUCCESS</status> <status_message>Network deleted succesfully</status_message> <event_type>**NETWORK_DELETE**</event_type> <network_name>testn1</network_name> <network_bridge/> <network_sriov/> <network_vlan/> <network_trunk/> </nfvisEvent> </notification>

Event Type	Notification Trigger	Notification Output Example
UPGRADE_REGISTER	System upgrade is registered.	<?xml version="1.0" encoding="UTF-8"?> <notification xmlns="urn:ietf:params:xml:ns:netconf:notification:1.0"> <eventTime>2017-04-26T15:57:50.434636+00:00</eventTime> <nfvisEvent xmlns="http://www.cisco.com/nfvis/notifier"> <user_id>NA</user_id> <config_change>true</config_change> <transaction_id>0</transaction_id> <status>SUCCESS</status> <status_code>0</status_code> <status_message>Upgrade package registration successful: Cisco_NFVIS_Upgrade-3.6.1-698-20170402_042811.nfvispkg</status_message> <event_type>**UPGRADE_REGISTER**</event_type> </nfvisEvent> </notification>
UPGRADE_APPLY	System upgrade is applied.	<?xml version="1.0" encoding="UTF-8"?> <notification xmlns="urn:ietf:params:xml:ns:netconf:notification:1.0"> <eventTime>2017-04-26T16:02:43.885516+00:00</eventTime> <nfvisEvent xmlns="http://www.cisco.com/nfvis/notifier"> <user_id>NA</user_id> <config_change>true</config_change> <transaction_id>0</transaction_id> <status>SUCCESS</status> <status_code>0</status_code> <status_message>Upgrade Process: In Progress</status_message> <event_type>**UPGRADE_APPLY**</event_type> </nfvisEvent> </notification>

Event Type	Notification Trigger	Notification Output Example
ROTATED_LOGS_DELETE	Rotated logs older than 30 days are deleted by the system.	<?xml version="1.0" encoding="UTF-8"?> <rpc-reply xmlns="urn:ietf:params:xml:ns:netconf:base:1.0" message-id="1"> <ok/> </rpc-reply> <?xml version="1.0" encoding="UTF-8"?> <notification xmlns="urn:ietf:params:xml:ns:netconf:notification:1.0"> <eventTime>2017-04-26T17:38:10.321152+00:00</eventTime> <nfvisEvent xmlns="http://www.cisco.com/nfvis/notifier"> <user_id>NA</user_id> <config_change>true</config_change> <transaction_id>0</transaction_id> <status>SUCCESS</status> <status_code>0</status_code> <status_message>Deleted rotated logs from archive older than 30 days</status_message> <details>NA</details> <event_type>**ROTATED_LOGS_DELETE**</event_type> </nfvisEvent> </notification>
ROTATED_LOGS_DELETE	Older logs deleted by the system when the total file size of rotated logs exceeds 2GB.	<?xml version="1.0" encoding="UTF-8"?> <rpc-reply xmlns="urn:ietf:params:xml:ns:netconf:base:1.0" message-id="1"> <ok/> </rpc-reply> <?xml version="1.0" encoding="UTF-8"?> <notification xmlns="urn:ietf:params:xml:ns:netconf:notification:1.0"> <eventTime>2017-04-26T17:42:10.321152+00:00</eventTime> <nfvisEvent xmlns="http://www.cisco.com/nfvis/notifier"> <user_id>NA</user_id> <config_change>true</config_change> <transaction_id>0</transaction_id> <status>SUCCESS</status> <status_code>0</status_code> <status_message>Rotated logs had exceeded 2G, older logs have been deleted to make space</status_message> <details>NA</details> <event_type>**ROTATED_LOGS_DELETE**</event_type> </nfvisEvent> </notification>

Event Type	Notification Trigger	Notification Output Example
SECURE_OVERLAY_CREATING	create secure overlay	<notification <eventTime> 2018-11-02T04:23:02.641317+00:00 <nfvisEvent <user_id>admin</user_id> <config_change>true</config_change> <transaction_id>0</transaction_id> <status>SUCCESS</status> <status_code>0</status_code> <status_message>Secure Overlay mgmthub initial creation. Active local bridge: wan-br</status_message> <details>NA</details> <event_type>SECURE_OVERLAY_CREATING<event_type> <severity> INFO</severity> <hostname>nfvis</hostname> </nfvisEvent> </notification>
SECURE_OVERLAY_UP	Secure Overlay is UP	<?xml version="1.0" encoding="UTF-8"?> <notification xmlns="urn:ietf:params:xml:ns:netconf:notification:1.0"> <eventTime>2018-04-26T09:47:06.066264+00:00</eventTime> <nfvisEvent xmlns="http://www.cisco.com/nfvis/notifier"> <user_id>admin</user_id> <config_change>true</config_change> <transaction_id>0</transaction_id> <status>SUCCESS</status> <status_code>0</status_code> <status_message>Secure Overlay mgmthub up. Active bridge: wan-br</status_message> <details>Secure overlay initial creation</details> <event_type>SECURE_OVERLAY_UP</event_type> <severity>INFO</severity> <hostname>NFVIS</hostname> </nfvisEvent> </notification>

Event Type	Notification Trigger	Notification Output Example
WAN_DHCP_SWITCHOVER	WAN bridge toggle	<?xml version="1.0" encoding="UTF-8"?> <notification xmlns="urn:ietf:params:xml:ns:netconf:notification:1.0"> <eventTime>2018-04-26T09:47:06.066264+00:00</eventTime> <nfvisEvent xmlns="http://www.cisco.com/nfvis/notifier"> <user_id>admin</user_id> <config_change>true</config_change> <transaction_id>0</transaction_id> <status>SUCCESS</status> <status_code>0</status_code> <status_message>Switch over to bridge wan2-br for auto DHCP enablement successful</status_message> <details>NA</details> <event_type>WAN_DHCP_SWITCHOVER</event_type> <severity>INFO</severity> <hostname>NFVIS</hostname> </nfvisEvent> </notification>
WAN_DHCP_TOGGLE_END	WAN bridge toggle	<?xml version="1.0" encoding="UTF-8"?> <notification xmlns="urn:ietf:params:xml:ns:netconf:notification:1.0"> <eventTime>2018-04-26T09:47:06.066264+00:00</eventTime> <nfvisEvent xmlns="http://www.cisco.com/nfvis/notifier"> <user_id>admin</user_id> <config_change>true</config_change> <transaction_id>0</transaction_id> <status>SUCCESS</status> <status_code>0</status_code> <status_message>Disabling bridge toggle for auto DHCP enablement.</status_message> <details>NA</details> <event_type>WAN_DHCP_TOGGLE_END</event_type> <severity>INFO</severity> <hostname>NFVIS</hostname> </nfvisEvent> </notification>

Event Type	Notification Trigger	Notification Output Example
ROUTE_DISTRIBUTION_START	To start route distribution	<?xml version="1.0" encoding="UTF-8"?> <notification xmlns="urn:ietf:params:xml:ns:netconf:notification:1.0"> <eventTime>2019-03-15T21:46:28.034133+00:00</eventTime> <nfvisEvent xmlns="http://www.cisco.com/nfvis/notifier"> <user_id>admin</user_id> <config_change>true</config_change> <transaction_id>0</transaction_id> <status>SUCCESS</status> <status_code>0</status_code> <status_message>Route Distribution initial creation. Neighbor Address: 172.25.221.106</status_message> <details>NA</details> <event_type>ROUTE_DISTRIBUTION_START</event_type> <severity>INFO</severity> <hostname>NFVIS</hostname> </nfvisEvent> </notification>
ROUTE_DISTRIBUTION_DOWN	Route distribution is down	<?xml version="1.0" encoding="UTF-8"?> <notification xmlns="urn:ietf:params:xml:ns:netconf:notification:1.0"> <eventTime>2019-03-15T21:46:28.034133+00:00</eventTime> <nfvisEvent xmlns="http://www.cisco.com/nfvis/notifier"> <user_id>admin</user_id> <config_change>true</config_change> <transaction_id>0</transaction_id> <status>SUCCESS</status> <status_code>0</status_code> <status_message>Neighbor Address: 172.25.221.106</status_message> <details>NA</details> <event_type>ROUTE_DISTRIBUTION_DOWN</event_type> <severity>INFO</severity> <hostname>NFVIS</hostname> </nfvisEvent> </notification>

Event Type	Notification Trigger	Notification Output Example
ROUTE_DISTRIBUTION_ERROR	Route distribution in error	<?xml version="1.0" encoding="UTF-8"?> <notification xmlns="urn:ietf:params:xml:ns:netconf:notification:1.0"> <eventTime>2019-03-15T21:46:28.034133+00:00</eventTime> <nfvisEvent xmlns="http://www.cisco.com/nfvis/notifier"> <user_id>admin</user_id> <config_change>true</config_change> <transaction_id>0</transaction_id> <status>SUCCESS</status> <status_code>0</status_code> <status_message>Neighbor Address: 172.25.221.106</status_message> <details>NA</details> <event_type>ROUTE_DISTRIBUTION_ERROR</event_type> <severity>INFO</severity> <hostname>NFVIS</hostname> </nfvisEvent> </notification>
ROUTE_DISTRIBUTION_DELETE	Route distribution deleted	<?xml version="1.0" encoding="UTF-8"?> <notification xmlns="urn:ietf:params:xml:ns:netconf:notification:1.0"> <eventTime>2019-03-15T21:46:28.034133+00:00</eventTime> <nfvisEvent xmlns="http://www.cisco.com/nfvis/notifier"> <user_id>admin</user_id> <config_change>true</config_change> <transaction_id>0</transaction_id> <status>SUCCESS</status> <status_code>0</status_code> <status_message>All Neighbor Addresses deleted</status_message> <details>NA</details> <event_type>ROUTE_DISTRIBUTION_DELETE</event_type> <severity>INFO</severity> <hostname>NFVIS</hostname> </nfvisEvent> </notification>

Event Type	Notification Trigger	Notification Output Example
ROUTE_DISTRIBUTION_UP	Route distribution up	<?xml version="1.0" encoding="UTF-8"?> <notification xmlns="urn:ietf:params:xml:ns:netconf:notification:1.0"> <eventTime>2019-03-15T21:46:28.034133+00:00</eventTime> <nfvisEvent xmlns="http://www.cisco.com/nfvis/notifier"> <user_id>admin</user_id> <config_change>true</config_change> <transaction_id>0</transaction_id> <status>SUCCESS</status> <status_code>0</status_code> <status_message>Neighbor Address: 172.25.221.106</status_message> <details>NA</details> <event_type>ROUTE_DISTRIBUTION_UP</event_type> <severity>INFO</severity> <hostname>NFVIS</hostname> </nfvisEvent> </notification>
OVS_DPDK_SUCCESS	Enable DPDK	<?xml version="1.0" encoding="UTF-8"?> <notification xmlns="urn:ietf:params:xml:ns:netconf:notification:1.0"> <eventTime>2019-03-15T21:46:28.034133+00:00</eventTime> <nfvisEvent xmlns="http://www.cisco.com/nfvis/notifier"> <user_id>admin</user_id> <config_change>false</config_change> <transaction_id>0</transaction_id> <status>SUCCESS</status> <status_code>0</status_code> <status_message>OVS-DPDK enabled</status_message> <details>NA</details> <event_type>OVS_DPDK_SUCCESS</event_type> <severity>INFO</severity> <hostname>NFVIS</hostname> </nfvisEvent> </notification>

Event Type	Notification Trigger	Notification Output Example
OVS_DPDK_FAILURE	DPDK failure	<?xml version="1.0" encoding="UTF-8"?> <notification xmlns="urn:ietf:params:xml:ns:netconf:notification:1.0"> <eventTime>2019-03-15T21:46:28.034133+00:00</eventTime> <nfvisEvent xmlns="http://www.cisco.com/nfvis/notifier"> <user_id>admin</user_id> <config_change>false</config_change> <transaction_id>0</transaction_id> <status>SUCCESS</status> <status_code>0</status_code> <status_message>Unable to allocate CPU</status_message> <details>NA</details> <event_type>OVS_DPDK_FAILURE</event_type> <severity>INFO</severity> <hostname>NFVIS</hostname> </nfvisEvent> </notification>
PNIC_SRIOV_ENABLE	Enable SR-IOV to 2 vfs	<?xml version="1.0" encoding="UTF-8"?> <notification xmlns="urn:ietf:params:xml:ns:netconf:notification:1.0"> <eventTime>2019-03-15T21:46:28.034133+00:00</eventTime> <nfvisEvent xmlns="http://www.cisco.com/nfvis/notifier"> <user_id>admin</user_id> <config_change>true</config_change> <transaction_id>0</transaction_id> <status>SUCCESS</status> <status_code>0</status_code> <status_message>Physical Interface: eth0-1 Number of VFs 2</status_message> <details>NA</details> <event_type>PNIC_SRIOV_ENABLE</event_type> <severity>INFO</severity> <hostname>NFVIS</hostname> </nfvisEvent> </notification>

Event Type	Notification Trigger	Notification Output Example
PNIC_SRIOV_DISABLE	Disable SR-IOV	<?xml version="1.0" encoding="UTF-8"?> <notification xmlns="urn:ietf:params:xml:ns:netconf:notification:1.0"> <eventTime>2019-03-15T21:46:28.034133+00:00</eventTime> <nfvisEvent xmlns="http://www.cisco.com/nfvis/notifier"> <user_id>admin</user_id> <config_change>true</config_change> <transaction_id>0</transaction_id> <status>SUCCESS</status> <status_code>0</status_code> <status_message>Physical Interface: eth0-1 Number of VFs 0</status_message> <details>NA</details> <event_type>PNIC_SRIOV_DISABLE</event_type> <severity>INFO</severity> <hostname>NFVIS</hostname> </nfvisEvent> </notification>
PNIC_SRIOV_ENABLE	Enable of disable SR-IOV in error	<?xml version="1.0" encoding="UTF-8"?> <notification xmlns="urn:ietf:params:xml:ns:netconf:notification:1.0"> <eventTime>2019-03-15T21:46:28.034133+00:00</eventTime> <nfvisEvent xmlns="http://www.cisco.com/nfvis/notifier"> <user_id>admin</user_id> <config_change>true</config_change> <transaction_id>0</transaction_id> <status>SUCCESS</status> <status_code>0</status_code> <status_message>Physical Interface: eth0-1 Number of VFs 2</status_message> <details>NA</details> <event_type>PNIC_SRIOV_ENABLE</event_type> <severity>INFO</severity> <hostname>NFVIS</hostname> </nfvisEvent> </notification>

Event Type	Notification Trigger	Notification Output Example
PNIC_SRIOV_UPDATE	Set switchmode to veb/vepa	<?xml version="1.0" encoding="UTF-8"?> <notification xmlns="urn:ietf:params:xml:ns:netconf:notification:1.0"> <eventTime>2019-03-15T21:46:28.034133+00:00</eventTime> <nfvisEvent xmlns="http://www.cisco.com/nfvis/notifier"> <user_id>admin</user_id> <config_change>true</config_change> <transaction_id>0</transaction_id> <status>SUCCESS</status> <status_code>0</status_code> <status_message>Physical Interface: eth0-1 Switchmode vepa</status_message> <details>NA</details> <event_type>PNIC_SRIOV_UPDATE</event_type> <severity>INFO</severity> <hostname>NFVIS</hostname> </nfvisEvent> </notification>
PNIC_SRIOV_UPDATE	Set switchmode to veb/vepa in error	<?xml version="1.0" encoding="UTF-8"?> <notification xmlns="urn:ietf:params:xml:ns:netconf:notification:1.0"> <eventTime>2019-03-15T21:46:28.034133+00:00</eventTime> <nfvisEvent xmlns="http://www.cisco.com/nfvis/notifier"> <user_id>admin</user_id> <config_change>true</config_change> <transaction_id>0</transaction_id> <status>SUCCESS</status> <status_code>0</status_code> <status_message>Physical Interface: eth0-1 Switchmode vepa</status_message> <details>NA</details> <event_type>PNIC_SRIOV_UPDATE</event_type> <severity>INFO</severity> <hostname>NFVIS</hostname> </nfvisEvent> </notification>

Event Type	Notification Trigger	Notification Output Example
PROMISC_ENABLED	Enable promiscuous mode	<?xml version="1.0" encoding="UTF-8"?> <notification xmlns="urn:ietf:params:xml:ns:netconf:notification:1.0"> <eventTime>2020-03-27T19:58:52.333682+00:00</eventTime> <nfvisEvent <user_id>admin</user_id> <config_change>false</config_change> <transaction_id>0</transaction_id> <status>SUCCESS</status> <status_code>0</status_code> <status_message>GE0-0: promiscuous mode enabled</status_message> <details>GE0-0: promiscuous mode enabled</details> <event_type>PROMISC_ENABLED</event_type> <severity>INFO</severity> <hostname>NFVIS</hostname> </nfvisEvent> </notification>
FIPS_CONFIGURED_REBOOT_TO_ENABLE	When FIPS mode is configured but a reboot is required to enable it.	<?xml version="1.0" encoding="UTF-8"?> <notification xmlns="urn:ietf:params:xml:ns:netconf:notification:1.0"> <eventTime>2020-06-07T21:26:06.231166+00:00</eventTime> <nfvisEvent <user_id>admin</user_id> <config_change>false</config_change> <transaction_id>0</transaction_id> <status>SUCCESS</status> <status_code>0</status_code> <status_message>Fips-mode configured and a manual reboot is required to enable it</status_message> <details>NA</details> <event_type>FIPS_CONFIGURED_REBOOT_TO_ENABLE</event_type> <severity>INFO</severity> <hostname>NFVIS</hostname> </nfvisEvent> </notification>

Event Type	Notification Trigger	Notification Output Example
FIPS_ENABLED	FIPS mode is enabled.	<?xml version="1.0" encoding="UTF-8"?> <notification xmlns="urn:ietf:params:xml:ns:netconf:notification:1.0"> <eventTime>2020-06-07T21:26:06.231166+00:00</eventTime> <nfvisEvent <user_id>admin</user_id> <config_change>false</config_change> <transaction_id>0</transaction_id> <status>SUCCESS</status> <status_code>0</status_code> <status_message>Fips-mode successfully enabled</status_message> <details>NA</details> <event_type>FIPS_ENABLED</event_type> <severity>INFO</severity> <hostname>NFVIS</hostname> </nfvisEvent> </notification>
FIPS_FAILED		<?xml version="1.0" encoding="UTF-8"?> <notification xmlns="urn:ietf:params:xml:ns:netconf:notification:1.0"> <eventTime>2020-06-07T21:26:06.231166+00:00</eventTime> <nfvisEvent <user_id>admin</user_id> <config_change>false</config_change> <transaction_id>0</transaction_id> <status>SUCCESS</status> <status_code>0</status_code> <status_message>Fips-mode configuration is unsuccessful</status_message> <details>NA</details> <event_type>FIPS_FAILED</event_type> <severity>INFO</severity> <hostname>NFVIS</hostname> </nfvisEvent> </notification>

Event Type	Notification Trigger	Notification Output Example
FIPS_DISABLED	FIPS mode is disabled.	<?xml version="1.0" encoding="UTF-8"?> <notification xmlns="urn:ietf:params:xml:ns:netconf:notification:1.0"> <eventTime>2020-06-07T21:26:06.231166+00:00</eventTime> <nfvisEvent <user_id>admin</user_id> <config_change>false</config_change> <transaction_id>0</transaction_id> <status>SUCCESS</status> <status_code>0</status_code> <status_message>Fips-mode successfully disabled</status_message> <details>NA</details> <event_type>FIPS_DISABLED</event_type> <severity>INFO</severity> <hostname>NFVIS</hostname> </nfvisEvent> </notification>
FIPS_UNCONFIGURED_REBOOT_TO_DISABLE		<?xml version="1.0" encoding="UTF-8"?> <notification xmlns="urn:ietf:params:xml:ns:netconf:notification:1.0"> <eventTime>2020-06-07T21:26:06.231166+00:00</eventTime> <nfvisEvent <user_id>admin</user_id> <config_change>false</config_change> <transaction_id>0</transaction_id> <status>SUCCESS</status> <status_code>0</status_code> <status_message>Fips-mode unconfigured and a manual reboot is required to disable it</status_message> <details>NA</details> <event_type>FIPS_UNCONFIGURED_REBOOT_TO_DISABLE</event_type> <severity>INFO</severity> <hostname>NFVIS</hostname> </nfvisEvent> </notification>

vmlcEvent

Event Type	Notification Trigger	Notification Output Example
CREATE_IMAGE	The VM image is registered.	<?xml version="1.0" encoding="UTF-8"?> <notification xmlns="urn:ietf:params:xml:ns:netconf:notification:1.0"> <eventTime>2016-11-17T17:12:30.76+00:00</eventTime> <vmlcEvent xmlns="http://www.cisco.com/nfvis/vm_lifecycle"> <status>SUCCESS</status> <status_code>200</status_code> <status_message>Image creation completed successfully.</status_message> <image>isrv-universalk9.16.03.01.tar.gz</image> <vm_source></vm_source> <vm_target></vm_target> <event> <type>**CREATE_IMAGE**</type> </event> </vmlcEvent> </notification>
DELETE_IMAGE	The VM image is unregistered.	<?xml version="1.0" encoding="UTF-8"?> <notification xmlns="urn:ietf:params:xml:ns:netconf:notification:1.0"> <eventTime>2016-11-17T17:14:51.169+00:00</eventTime> <vmlcEvent xmlns="http://www.cisco.com/nfvis/vm_lifecycle"> <status>SUCCESS</status> <status_code>200</status_code> <status_message>Image deletion completed successfully.</status_message> <image>isrv-universalk9.16.03.01.tar.gz</image> <vm_source></vm_source> <vm_target></vm_target> <event> <type>**DELETE_IMAGE**</type> </event> </vmlcEvent> </notification>

Event Type	Notification Trigger	Notification Output Example
CREATE_FLAVOR	A flavor is created.	<?xml version="1.0" encoding="UTF-8"?> <notification xmlns="urn:ietf:params:xml:ns:netconf:notification:1.0"> <eventTime>2016-11-17T17:12:29.685+00:00</eventTime> <vmlcEvent xmlns="http://www.cisco.com/nfvis/vm_lifecycle"> <status>SUCCESS</status> <status_code>200</status_code> <status_message>Flavor creation completed successfully.</status_message> <flavor>ISRv-small</flavor> <vm_source></vm_source> <vm_target></vm_target> <event> <type>**CREATE_FLAVOR**</type> </event> </vmlcEvent> </notification>
DELETE_FLAVOR	A flavor is deleted.	<?xml version="1.0" encoding="UTF-8"?> <notification xmlns="urn:ietf:params:xml:ns:netconf:notification:1.0"> <eventTime>2016-11-17T17:14:51.425+00:00</eventTime> <vmlcEvent xmlns="http://www.cisco.com/nfvis/vm_lifecycle"> <status>SUCCESS</status> <status_code>200</status_code> <status_message>Flavor deletion completed successfully.</status_message> <flavor>ISRv-small</flavor> <vm_source></vm_source> <vm_target></vm_target> <event> <type>**DELETE_FLAVOR**</type> </event> </vmlcEvent> </notification>

Event Type	Notification Trigger	Notification Output Example
VM_DEPLOYED	The VM is deployed.	<?xml version="1.0" encoding="UTF-8"?> <notification xmlns="urn:ietf:params:xml:ns:netconf:notification:1.0"> <eventTime>2016-11-17T17:19:16.927+00:00</eventTime> <vmlcEvent xmlns="http://www.cisco.com/nfvis/vm_lifecycle"> <status>SUCCESS</status> <status_code>200</status_code> <status_message>VIM Driver: VM successfully created, VM Name: [SystemAdminTena_ROUTER_0_df6733c1-0768-4ae6-8dce-b223ecdb036c]</status_message> <depname>1479341445</depname> <tenant>admin</tenant> <tenant_id>AdminTenantId</tenant_id> <depid>c64d79db-3a29-41a8-8114-c80d42731a5b</depid> <vm_group>ROUTER</vm_group> <vm_source> <vmid>d18dd252-80c8-44f2-ab66-d4481790bb79</vmid> <hostid>NFVIS</hostid> <hostname>NFVIS</hostname> <interfaces> <interface> <nicid>0</nicid> <port_id>vnet0</port_id> <network>int-mgmt-net</network> <subnet>N/A</subnet> <ip_address>10.20.0.2</ip_address> <mac_address>52:54:00:31:c5:7f</mac_address> <netmask>255.255.255.0</netmask> <gateway>10.20.0.1</gateway> </interface> <interface> <nicid>1</nicid> <port_id>vnet1</port_id> <network>wan-net</network> <subnet>N/A</subnet> <mac_address>52:54:00:59:52:41</mac_address> <netmask>255.255.255.0</netmask> <gateway>172.19.181.152</gateway> </interface> </interfaces> </vm_source> <vm_target></vm_target> <event> <type>**VM_DEPLOYED**</type> </event> </vmlcEvent> </notification>

Event Type	Notification Trigger	Notification Output Example
VM_ALIVE	The state of a monitored VM becomes ACTIVE.	<?xml version="1.0" encoding="UTF-8"?> <notification xmlns="urn:ietf:params:xml:ns:netconf:notification:1.0"> <eventTime>2016-11-17T17:22:47.306+00:00</eventTime> <vmlcEvent xmlns="http://www.cisco.com/nfvis/vm_lifecycle"> <status>SUCCESS</status> <status_code>200</status_code> <status_message>VM_Alive event received, VM ID: [SystemAdminTena_ROUTER_0_df6733c1-0768-4ae6-8dce-b223ecdb036c]</status_message> <depname>1479341445</depname> <tenant>admin</tenant> <tenant_id>AdminTenantId</tenant_id> <depid>c64d79db-3a29-41a8-8114-c80d42731a5b</depid> <vm_group>ROUTER</vm_group> <vm_source> <vmid>d18dd252-80c8-44f2-ab66-d4481790bb79</vmid> <hostid>NFVIS</hostid> <hostname>NFVIS</hostname> <interfaces> <interface> <nicid>0</nicid> <port_id>vnet0</port_id> <network>int-mgmt-net</network> <subnet>N/A</subnet> <ip_address>10.20.0.2</ip_address> <mac_address>52:54:00:31:c5:7f</mac_address> <netmask>255.255.255.0</netmask> <gateway>10.20.0.1</gateway> </interface> <interface> <nicid>1</nicid> <port_id>vnet1</port_id> <network>wan-net</network> <subnet>N/A</subnet> <mac_address>52:54:00:59:52:41</mac_address> <netmask>255.255.255.0</netmask> <gateway>172.19.181.152</gateway> </interface> </interfaces> </vm_source> <vm_target></vm_target> <event> <type>**VM_ALIVE**</type> </event> </vmlcEvent> </notification>

Event Type	Notification Trigger	Notification Output Example
VM_UNDEPLOYED	The VM is undeployed.	<?xml version="1.0" encoding="UTF-8"?> <notification xmlns="urn:ietf:params:xml:ns:netconf:notification:1.0"> <eventTime>2016-11-17T17:31:40.6+00:00</eventTime> <vmlcEvent xmlns="http://www.cisco.com/nfvis/vm_lifecycle"> <status>SUCCESS</status> <status_code>204</status_code> <status_message>VIM Driver: VM successfully deleted</status_message> <depname>1479341445</depname> <tenant>admin</tenant> <tenant_id>AdminTenantId</tenant_id> <depid>c64d79db-3a29-41a8-8114-c80d42731a5b</depid> <vm_group>ROUTER</vm_group> <vm_source> <vmid>d18dd252-80c8-44f2-ab66-d4481790bb79</vmid> <hostid>NFVIS</hostid> <hostname>NFVIS</hostname> <interfaces> <interface> <nicid>0</nicid> <port_id>vnet0</port_id> <network>int-mgmt-net</network> <subnet>N/A</subnet> <ip_address>10.20.0.2</ip_address> <mac_address>52:54:00:31:c5:7f</mac_address> <netmask>255.255.255.0</netmask> <gateway>10.20.0.1</gateway> </interface> <interface> <nicid>1</nicid> <port_id>vnet1</port_id> <network>wan-net</network> <subnet>N/A</subnet> <mac_address>52:54:00:59:52:41</mac_address> <netmask>255.255.255.0</netmask> <gateway>172.19.181.152</gateway> </interface> </interfaces> </vm_source> <vm_target></vm_target> <event> <type>**VM_UNDEPLOYED**</type> </event> </vmlcEvent> </notification>

Event Type	Notification Trigger	Notification Output Example
SERVICE_UPDATED	The VM is updated.	<?xml version="1.0" encoding="UTF-8"?> <notification xmlns="urn:ietf:params:xml:ns:netconf:notification:1.0"> <eventTime>2016-11-17T17:51:45.5+00:00</eventTime> <vmlcEvent xmlns="http://www.cisco.com/nfvis/vm_lifecycle"> <status>SUCCESS</status> <status_code>200</status_code> <status_message>Service group update completed successfully</status_message> <depname>1479342258</depname> <tenant>admin</tenant> <tenant_id>AdminTenantId</tenant_id> <depid>827e871a-30d5-4f5f-a05a-263b7ee3a734</depid> <vm_source></vm_source> <vm_target></vm_target> <event> <type>**SERVICE_UPDATED**</type> </event> </vmlcEvent> </notification>
VM_STOPPED	The VM is stopped per VM action requrest.	<?xml version="1.0" encoding="UTF-8"?> <notification xmlns="urn:ietf:params:xml:ns:netconf:notification:1.0"> <eventTime>2016-11-17T17:26:05.762+00:00</eventTime> <vmlcEvent xmlns="http://www.cisco.com/nfvis/vm_lifecycle"> <status>SUCCESS</status> <status_code>200</status_code> <status_message>Successfully stopped VM [SystemAdminTena_ROUTER_0_df6733c1-0768-4ae6-8dce-b223ecdb036c].</status_message> <depname>1479341445</depname> <tenant>admin</tenant> <tenant_id>AdminTenantId</tenant_id> <svcid>NULL</svcid> <depid>c64d79db-3a29-41a8-8114-c80d42731a5b</depid> <vm_group>ROUTER</vm_group> <vm_source> <vmid>d18dd252-80c8-44f2-ab66-d4481790bb79</vmid> <hostid>NFVIS</hostid> <hostname>NFVIS</hostname> </vm_source> <vm_target></vm_target> <event> <type>**VM_STOPPED**</type> </event> </vmlcEvent> </notification>

Event Type	Notification Trigger	Notification Output Example
VM_STARTED	The VM is started per VM action request.	<?xml version="1.0" encoding="UTF-8"?> <notification xmlns="urn:ietf:params:xml:ns:netconf:notification:1.0"> <eventTime>2016-11-17T17:26:40.398+00:00</eventTime> <vmlcEvent xmlns="http://www.cisco.com/nfvis/vm_lifecycle"> <status>SUCCESS</status> <status_code>200</status_code> <status_message>Started VM [SystemAdminTena_ROUTER_0_df6733c1-0768-4ae6-8dce-b223ecdb036c].</status_message> <depname>1479341445</depname> <tenant>admin</tenant> <tenant_id>AdminTenantId</tenant_id> <svcid>NULL</svcid> <depid>c64d79db-3a29-41a8-8114-c80d42731a5b</depid> <vm_group>ROUTER</vm_group> <vm_source> <vmid>d18dd252-80c8-44f2-ab66-d4481790bb79</vmid> <hostid>NFVIS</hostid> <hostname>NFVIS</hostname> </vm_source> <vm_target></vm_target> <event> <type>**VM_STARTED**</type> </event> </vmlcEvent> </notification>

Event Type	Notification Trigger	Notification Output Example
VM_REBOOTED	The VM is rebooted per VM action request.	<?xml version="1.0" encoding="UTF-8"?> <notification xmlns="urn:ietf:params:xml:ns:netconf:notification:1.0"> <eventTime>2016-11-17T17:36:56.5+00:00</eventTime> <vmlcEvent xmlns="http://www.cisco.com/nfvis/vm_lifecycle"> <status>SUCCESS</status> <status_code>200</status_code> <status_message>Rebooted VM [SystemAdminTena_ROUTER_0_f17fc494-8535-4b05-b88d-f0fd2effdc7d]</status_message> <depname>1479342258</depname> <tenant>admin</tenant> <tenant_id>AdminTenantId</tenant_id> <svcid>NULL</svcid> <depid>827e871a-30d5-4f5f-a05a-263b7ee3a734</depid> <vm_group>ROUTER</vm_group> <vm_source> <vmid>d918a3b1-f2a9-4065-9d8e-2135b0a37d87</vmid> <hostid>NFVIS</hostid> <hostname>NFVIS</hostname> </vm_source> <vm_target></vm_target> <event> <type>**VM_REBOOTED**</type> </event> </vmlcEvent> </notification>

Event Type	Notification Trigger	Notification Output Example
VM_RECOVERY_INIT	A monitored VM is not reachable.	<?xml version="1.0" encoding="UTF-8"?> <notification xmlns="urn:ietf:params:xml:ns:netconf:notification:1.0"> <eventTime>2016-11-17T16:27:51.627+00:00</eventTime> <vmlcEvent xmlns="http://www.cisco.com/nfvis/vm_lifecycle"> <status>SUCCESS</status> <status_code>200</status_code> <status_message>Recovery event for VM [SystemAdminTena_ROUTER_0_40ae18be-5930-4d94-95ff-dbb0b56ef12b] trigge red. Processing Auto healing. Proceeding with Recovery.</status_message> <depname>1479328919</depname> <tenant>admin</tenant> <tenant_id>AdminTenantId</tenant_id> <svcid>NULL</svcid> <depid>9e7fe4f8-a5f4-4a6d-aad7-121405be4ba4</depid> <vm_group>ROUTER</vm_group> <vm_source> <vmid>000883fc-77f3-4b9e-aaf6-0f31d88a8f67</vmid> <hostid>NFVIS</hostid> <hostname>NFVIS</hostname> </vm_source> <vm_target></vm_target> <event> <type>**VM_RECOVERY_INIT**</type> </event> </vmlcEvent> </notification>

Event Type	Notification Trigger	Notification Output Example
VM_RECOVERY_REBOOT	Recovery reboot starts for the monitored VM, which is not reachable.	<?xml version="1.0" encoding="UTF-8"?> <notification xmlns="urn:ietf:params:xml:ns:netconf:notification:1.0"> <eventTime>2016-11-17T16:27:53.979+00:00</eventTime> <vmlcEvent xmlns="http://www.cisco.com/nfvis/vm_lifecycle"> <status>SUCCESS</status> <status_code>200</status_code> <status_message>VM [SystemAdminTena_ROUTER_0_40ae18be-5930-4d94-95ff-dbb0b56ef12b] is being rebooted. </statu s_message> <depname>1479328919</depname> <tenant>admin</tenant> <tenant_id>AdminTenantId</tenant_id> <svcid>NULL</svcid> <depid>9e7fe4f8-a5f4-4a6d-aad7-121405be4ba4</depid> <vm_group>ROUTER</vm_group> <vm_source> <vmid>000883fc-77f3-4b9e-aaf6-0f31d88a8f67</vmid> <hostid>NFVIS</hostid> <hostname>NFVIS</hostname> </vm_source> <vm_target></vm_target> <event> <type>**VM_RECOVERY_REBOOT**</type> </event> </vmlcEvent> </notification>

Event Type	Notification Trigger	Notification Output Example
VM_RECOVERY_COMPLETE	Recovery reboot completes for the monitored VM, which is not reachable.	<?xml version="1.0" encoding="UTF-8"?> <notification xmlns="urn:ietf:params:xml:ns:netconf:notification:1.0"> <eventTime>2016-11-17T16:31:26.934+00:00</eventTime> <vmlcEvent xmlns="http://www.cisco.com/nfvis/vm_lifecycle"> <status>SUCCESS</status> <status_code>200</status_code> <status_message>Successfully recovered VM [SystemAdminTena_ROUTER_0_40ae18be-5930-4d94-95ff-dbb0b56ef12b]. status_message> <depname>1479328919</depname> <tenant>admin</tenant> <tenant_id>AdminTenantId</tenant_id> <svcid>NULL</svcid> <depid>9e7fe4f8-a5f4-4a6d-aad7-121405be4ba4</depid> <vm_group>ROUTER</vm_group> <vm_source> <vmid>000883fc-77f3-4b9e-aaf6-0f31d88a8f67</vmid> <hostid>NFVIS</hostid> <hostname>NFVIS</hostname> </vm_source> <vm_target> <vmid>000883fc-77f3-4b9e-aaf6-0f31d88a8f67</vmid> <hostid>NFVIS</hostid> <hostname>NFVIS</hostname> <interfaces> <interface> <nicid>0</nicid> <port_id>vnet0</port_id> <network>int-mgmt-net</network> <subnet>N/A</subnet> <ip_address>10.20.0.2</ip_address> <mac_address>52:54:00:7b:3f:de</mac_address> <netmask>255.255.255.0</netmask> <gateway>10.20.0.1</gateway> </interface> <interface> <nicid>1</nicid> <port_id>vnet1</port_id> <network>wan-net</network> <subnet>N/A</subnet> <mac_address>52:54:00:96:8a:4d</mac_address> <netmask>255.255.255.0</netmask> <gateway>172.19.181.152</gateway> </interface> </interfaces> </vm_target> <event> <type>**VM_RECOVERY_COMPLETE**</type> </event> </vmlcEvent> </notification>

Event Type	Notification Trigger	Notification Output Example
VM_MONITOR_UNSET	Monitoring is disabled per VM action request.	<?xml version="1.0" encoding="UTF-8"?> <notification xmlns="urn:ietf:params:xml:ns:netconf:notification:1.0"> <eventTime>2016-11-18T13:36:43.613+00:00</eventTime> <vmlcEvent xmlns="http://www.cisco.com/nfvis/vm_lifecycle"> <status>SUCCESS</status> <status_code>200</status_code> <status_message>Unset monitor completed successfully</status_message> <depname>1479413090</depname> <tenant>admin</tenant> <tenant_id>AdminTenantId</tenant_id> <svcid>NULL</svcid> <depid>742dd335-330c-4bf0-a75d-a44003c645c5</depid> <vm_group>ROUTER</vm_group> <vm_source> <vmid>23ec3793-37ab-4ec2-a978-a10e08585fdd</vmid> <hostid>NFVIS</hostid> <hostname>NFVIS</hostname> </vm_source> <vm_target></vm_target> <event> <type>**VM_MONITOR_UNSET**</type> </event> </vmlcEvent> </notification>
VM_MONITOR_SET	Monitoring is enabled per VM action request.	<?xml version="1.0" encoding="UTF-8"?> <notification xmlns="urn:ietf:params:xml:ns:netconf:notification:1.0"> <eventTime>2016-11-18T13:40:15.276+00:00</eventTime> <vmlcEvent xmlns="http://www.cisco.com/nfvis/vm_lifecycle"> <status>SUCCESS</status> <status_code>200</status_code> <status_message>Set monitor completed successfully</status_message> <depname>1479413090</depname> <tenant>admin</tenant> <tenant_id>AdminTenantId</tenant_id> <svcid>NULL</svcid> <depid>742dd335-330c-4bf0-a75d-a44003c645c5</depid> <vm_group>ROUTER</vm_group> <vm_source> <vmid>23ec3793-37ab-4ec2-a978-a10e08585fdd</vmid> <hostid>NFVIS</hostid> <hostname>NFVIS</hostname> </vm_source> <vm_target></vm_target> <event> <type>**VM_MONITOR_SET**</type> </event> </vmlcEvent> </notification>

Event Type	Notification Trigger	Notification Output Example
VM_UPDATED	VM's flavor is changed.	

Event Type	Notification Trigger	Notification Output Example
		<?xml version="1.0" encoding="UTF-8"?> <notification xmlns="urn:ietf:params:xml:ns:netconf:notification:1.0"> <eventTime>2017-12-08T00:50:39.034+00:00</eventTime> <vmlcEvent xmlns='http://www.cisco.com/nfvis/vm_lifecycle'> <status>SUCCESS</status> <status_code>200</status_code> <status_message>VM is resized with flavor [ISRv-medium].</status_message> <user_name>admin</user_name> <depname>1512766000</depname> <tenant>admin</tenant> <tenant_id>adminUUID</tenant_id> <depid>92c11aa1-f6dd-47d1-948f-c8c65b9ef70f</depid> <vm_group>ROUTER</vm_group> <vm_source> <vmid>1a6f587e-2779-4087-b84d-c0a2c8a481b1</vmid> <vmname>1512766000_ROUTER_0_60d15064-0c6d-49b9-aa4a-80587f626004</vmname> <hostid>NFVIS</hostid> <hostname>nfvis</hostname> <interfaces> <interface> <nicid>0</nicid> <type>virtual</type> <port_id>vnic0</port_id> <network>int-mgmt-net</network> <subnet>N/A</subnet> <ip_address>10.20.0.3</ip_address> <mac_address>52:54:00:3c:ee:5b</mac_address> <netmask>255.255.255.0</netmask> <gateway>10.20.0.1</gateway> </interface> <interface> <nicid>1</nicid> <type>virtual</type> <port_id>vnic1</port_id> <network>wan-net</network> <subnet>N/A</subnet> <mac_address>52:54:00:70:06:4a</mac_address> <netmask>255.255.255.0</netmask> <gateway>172.19.181.152</gateway> </interface> <interface> <nicid>2</nicid> <type>virtual</type> <port_id>vnic2</port_id> <network>lan-net</network> <subnet>N/A</subnet> <mac_address>52:54:00:c7:30:1c</mac_address> <netmask>255.255.255.0</netmask> <gateway>192.168.1.1</gateway> </interface> </interfaces> </vm_source> <event> <type>VM_UPDATED</type> </event>

Event Type	Notification Trigger	Notification Output Example
		</vmlcEvent></notification>]]>]]> <?xml version="1.0" encoding="UTF-8"?> <notification xmlns="urn:ietf:params:xml:ns:netconf:notification:1.0"> <eventTime>2017-12-08T00:50:39.06+00:00</eventTime> <vmlcEvent xmlns='http://www.cisco.com/nfvis/vm_lifecycle'> <status>SUCCESS</status> <status_code>200</status_code> <status_message>Service group update completed successfully</status_message> <user_name>admin</user_name> <depname>1512766000</depname> <tenant>admin</tenant> <tenant_id>adminUUID</tenant_id> <depid>92c11aa1-f6dd-47d1-948f-c8c65b9ef70f</depid> <event> <type>SERVICE_UPDATED</type> </event> </vmlcEvent>

Event Type	Notification Trigger	Notification Output Example
VM_UPDATED	VNIC is added, deleted or updated.	

Event Type	Notification Trigger	Notification Output Example

```
<?xml version="1.0" encoding="UTF-8"?>
<notification
 xmlns="urn:ietf:params:xml:ns:netconf:notification:1.0">
 <eventTime>2017-12-08T02:10:56.184+00:00</eventTime>
 <vmlcEvent
  xmlns='http://www.cisco.com/nfvis/vm_lifecycle'>
  <status>SUCCESS</status>
  <status_code>200</status_code>
  <status_message>Added 1 interface: [managed, net=my-net-1, nicid=3]
Updated 2 interface: [managed, net=lan-net, nicid=1],[managed,
net=wan-net, nicid=2]</status_message>
  <user_name>admin</user_name>
  <depname>1512766000</depname>
  <tenant>admin</tenant>
  <tenant_id>adminUUID</tenant_id>
  <depid>92c11aa1-f6dd-47d1-948f-c8c65b9ef70f</depid>
  <vm_group>ROUTER</vm_group>
  <vm_source>
   <vmid>1a6f587e-2779-4087-b84d-c0a2c8a481b1</vmid>

<vmname>1512766000_ROUTER_0_60d15064-0c6d-49b9-aa4a-80587f626004</vmname>

   <hostid>NFVIS</hostid>
   <hostname>nfvis</hostname>
   <interfaces>
    <interface>
     <nicid>0</nicid>
     <type>virtual</type>
     <port_id>vnic0</port_id>
     <network>int-mgmt-net</network>
     <subnet>N/A</subnet>
     <ip_address>10.20.0.3</ip_address>
     <mac_address>52:54:00:3c:ee:5b</mac_address>
     <netmask>255.255.255.0</netmask>
     <gateway>10.20.0.1</gateway>
    </interface>
    <interface>
     <nicid>1</nicid>
     <type>virtual</type>
     <port_id>vnic1</port_id>
     <network>lan-net</network>
     <subnet>N/A</subnet>
     <mac_address>52:54:00:70:06:4a</mac_address>
     <netmask>255.255.255.0</netmask>
     <gateway>192.168.1.1</gateway>
    </interface>
    <interface>
     <nicid>2</nicid>
     <type>virtual</type>
     <port_id>vnic2</port_id>
     <network>wan-net</network>
     <subnet>N/A</subnet>
     <mac_address>52:54:00:c7:30:1c</mac_address>
     <netmask>255.255.255.0</netmask>
     <gateway>172.19.181.152</gateway>
    </interface>
    <interface>
     <nicid>3</nicid>
     <type>virtual</type>
     <port_id>vnic3</port_id>
```

Event Type	Notification Trigger	Notification Output Example
		<network>my-net-1</network> <subnet>N/A</subnet> <mac_address>52:54:00:66:b5:c1</mac_address> </interface> </interfaces> </vm_source> <event> <type>VM_UPDATED</type> </event> </vmlcEvent>